

A FORWARD/REVERSE AUCTION ALGORITHM FOR ASYMMETRIC ASSIGNMENT PROBLEMS¹

by

Dimitri P. Bertsekas² and David A. Castañon³

Abstract

In this paper we consider the asymmetric assignment problem and we propose a new auction algorithm for its solution. The algorithm uses in a novel way the recently proposed idea of reverse auction, where in addition to persons bidding for objects by raising their prices, we also have objects competing for persons by essentially offering discounts. In practice, the new algorithm apparently deals better with price wars than the currently existing auction algorithms. As a result it frequently does not require ϵ -scaling for good practical performance, and tends to terminate substantially (and often dramatically) faster than its competitors.

¹ This work was supported in part by NSF under Grant CCR-9108058, and in part by the BM/C3 Technology branch of the United States Army Strategic Defense Command. It will be published also by the Journal of Computational Optimization and its Applications.

² Department of Electrical Engineering and Computer Science, M. I. T., Cambridge, Mass., 02139.

³ Department of Electrical Engineering, Boston University, and ALPHATECH, Inc., Burlington, Mass., 01803.

1. INTRODUCTION

We consider the classical asymmetric assignment problem where we want to match m persons with m out of n objects ($m < n$). The benefit for matching a person with an object is given, and we want to assign all the persons to distinct objects so as to maximize the total benefit. There are a number of methods for solving this problem, including primal-simplex and primal-dual (or sequential shortest path) methods [Ber91], [KeH80], [PaS82], [Roc84]. In this paper we will focus on auction algorithms, first proposed in [Ber79] for both symmetric and asymmetric problems, and subsequently developed in several other papers [Ber85], [Ber88], [BeE88], [BCT91]. The textbook [Ber91] contains an extensive discussion of these methods and their extensions to other network flow problems. Recent experimental evidence suggests that auction algorithms outperform their competitors by a substantial margin, particularly for sparse assignment problems [BeE88], [Ber90], [BCT91], and are also well suited for parallel computation [BeC89], [KKZ90], [PhZ88], [WeZ90], [WeZ91], [Zak91].

In the original proposal of the auction algorithm there is a price for each object, and at each iteration, one or more unassigned persons bid simultaneously for their “best” objects (the ones offering maximum benefit minus price), thereby raising the corresponding prices. Objects are then awarded to the highest bidder. The bidding increments must be at least equal to a positive parameter ϵ , and are chosen so as to preserve an ϵ -complementary slackness condition. For good practical (as well as theoretical) performance, it may be important to use ϵ -scaling, which consists of applying the algorithm several times, starting with a large value of ϵ and successively reducing ϵ up to an ultimate value that is less than some threshold ($1/m$ when a_{ij} are integer). Each scaling phase provides good initial prices for the next. The original proposal of the auction algorithm for asymmetric assignment problems had a deficiency: it required that the initial object prices be zero, thereby precluding the use of ϵ -scaling. As a result the method was susceptible to “price wars”, that is, protracted sequences of small price rises resulting from groups of persons competing for a smaller number of roughly equally desirable objects. Thus, in order to use auction algorithms to solve asymmetric problems where price wars are likely, one had to convert the problem to a symmetric one by adding $n - m$ artificial persons that can be assigned to any object at zero cost. There are specialized versions of the auction algorithm (the auction algorithm with similar persons [BeC89]) that can take advantage of the structure induced by the artificial persons. However, the approach of converting the problem to a symmetric problem introduces an undesirable increase in the problem’s dimension and to our knowledge has not seen much use.

In part to address the difficulty with price wars of the original asymmetric auction algorithm,

an alternative algorithm, called *reverse auction*, was recently developed in [BCT91]. Here, roughly speaking, the objects compete for persons by lowering their prices. In particular, objects decrease their prices to a level that is sufficiently low to lure a person away from its currently held object. One can show that forward and reverse auctions are mathematically equivalent, but their combination has resulted in algorithms that can solve various assignment-like problems much faster than forward or reverse auction can by themselves. In particular, an ϵ -scaled version of a combined forward/reverse auction was developed for asymmetric problems that can deal effectively with price wars. This method operates principally as a forward auction and uses reverse auction only near the end to rectify violations in the optimality conditions. According to computational results given in [BCT91], the solution times of this method for $m \times n$ asymmetric problems are quite reasonable and do not exceed the solution times of the original (forward only) auction algorithm for similar symmetric $m \times m$ problems by a factor larger than the natural ratio n/m .

However, as demonstrated in [BCT91], by frequently switching between forward and reverse auction, a substantial performance improvement can be obtained for *symmetric* assignment problems. A natural question therefore arises whether a similar improvement can be realized for asymmetric assignment problems by similarly combining forward and reverse auctions. The purpose of this paper is to develop such a method. Contrary to the method given in [BCT91], it is typically unnecessary to resort to ϵ -scaling, involving the solution of several subproblems, to deal with price wars. Our computational results show that ...

In Section 2, we define the asymmetric assignment problem, and we develop ϵ -complementary slackness conditions in a form suitable for our purposes. In Section 3, we introduce the new combined forward/reverse auction algorithm and we develop its basic properties. Finally, in Section 4 we provide computational results.

2. ASYMMETRIC ASSIGNMENT PROBLEMS

In the asymmetric assignment problem there are m persons and n objects ($m < n$). The benefit or value for assigning person i to object j is a_{ij} . The set of arcs of the underlying bipartite graph is denoted by \mathcal{A}

$$\mathcal{A} = \{(i, j) \mid j \in A(i), i = 1, \dots, m\}.$$

The set of objects to which person i can be assigned is a nonempty set denoted $A(i)$. The set of persons to which object j can be assigned is assumed nonempty and is denoted by $B(j) = \{i \mid j \in A(i)\}$. An *assignment* S is a (possibly empty) set of person-object pairs (i, j) such that $j \in A(i)$

3. A Forward/Reverse Auction Algorithm for Asymmetric Assignment Problems

for all $(i, j) \in S$; for each person i there can be at most one pair $(i, j) \in S$; and for every object j there can be at most one pair $(i, j) \in S$. Given an assignment S , we say that person i is *assigned* if there exists a pair $(i, j) \in S$; otherwise we say that i is *unassigned*. We use similar terminology for objects. An assignment is said to be *feasible* if it contains m pairs, so that every person is assigned; otherwise the assignment is called *partial*. The problem is said to be feasible if there exists at least one feasible assignment. We want to find an assignment $\{(1, j_1), \dots, (m, j_m)\}$ with maximum total benefit $\sum_{i=1}^m a_{ij_i}$.

A dual problem can be defined by introducing a price variable p_j for each object j and a profit variable π_i for each person i . It was shown in [BCT91] (see also [Ber91]) that a corresponding dual problem is

$$\begin{aligned} \text{minimize} \quad & \sum_{i=1}^m \pi_i + \sum_{j=1}^n p_j - (n - m) \min_{j=1, \dots, n} p_j \\ \text{subject to} \quad & \pi_i + p_j \geq a_{ij}, \quad \forall (i, j) \in \mathcal{A}. \end{aligned} \tag{1}$$

We denote by p the vector of prices (p_1, \dots, p_n) , and by π the vector of profits (π_1, \dots, π_m) . The following condition was introduced in [BCT91] for an assignment S and a pair (π, p) .

Definition 1: An assignment S and a pair (π, p) are said to satisfy ϵ -complementary slackness (ϵ -CS for short) if

$$\pi_i + p_j \geq a_{ij} - \epsilon, \quad \forall (i, j) \in \mathcal{A}, \tag{2a}$$

$$\pi_i + p_j = a_{ij}, \quad \forall (i, j) \in S, \tag{2b}$$

$$p_j \leq \min_{k: \text{assigned under } S} p_k, \quad \forall j: \text{ unassigned under } S. \tag{2c}$$

The following proposition, proved in [BCT91], clarifies the significance of ϵ -CS.

Proposition 1: If a feasible assignment S satisfies the ϵ -CS conditions (2a)-(2c) together with a pair (π, p) , then S is within $m\epsilon$ of being optimal for the asymmetric assignment problem. In particular, if the benefits a_{ij} are all integer and $\epsilon < 1/m$, S is an optimal assignment.

3. A FORWARD/REVERSE AUCTION ALGORITHM FOR ASYMMETRIC ASSIGNMENT PROBLEMS

In this section we consider algorithms that use a fixed $\epsilon > 0$, and maintain an assignment S and a pair (π, p) satisfying together with S the first two ϵ -CS conditions (2a) and (2b). They also maintain a scalar λ such that

$$p_j \geq \lambda, \quad \forall j \text{ that are assigned under } S. \tag{3}$$

3. A Forward/Reverse Auction Algorithm for Asymmetric Assignment Problems

The algorithms terminate when S becomes feasible and in addition all unassigned objects j satisfy $p_j \leq \lambda$. Thus upon termination, in view of Eq. (3), the third ϵ -CS condition (2c) is satisfied and by Prop. 1, the assignment S is optimal if $\epsilon < 1/m$ and the benefits a_{ij} are all integer. The level λ may be viewed as a *profitability threshold* below which we cannot drop the price of any assigned object. In the course of the algorithm, λ may be adjusted downward if it is set initially so high that not all persons can be assigned at prices above λ .

Note that we can initially select S to be empty, λ and p to be arbitrary, and π_i to be sufficiently large so that the conditions (2a) and (2b) are satisfied. Thus, in particular, we can try to use a favorable price vector such as one obtained from a scaling phase corresponding to a larger value of ϵ .

3.1 Forward and Reverse Auction Iterations

There are two types of iterations, forward and reverse. Forward iterations can be performed only as long as there is an unassigned person and reverse iterations can be performed only as long as there is an unassigned object j with $p_j > \lambda$. Both types of iterations start with an assignment S , a pair (π, p) , and a scalar λ satisfying conditions (2a), (2b), and (3).

Forward Iteration

Find an unassigned person i , its best object j_i

$$j_i = \arg \max_{j \in A(i)} \{a_{ij} - p_j\}, \quad (4)$$

the corresponding values

$$v_i = \max_{j \in A(i)} \{a_{ij} - p_j\}, \quad (5)$$

and

$$w_i = \max_{j \in A(i), j \neq j_i} \{a_{ij} - p_j\}. \quad (6)$$

[If j_i is the only object in $A(i)$, we define w_i to be $-\infty$ or, for computational purposes, a number that is much smaller than v_i .] Set

$$p_{j_i} := \max\{\lambda, a_{ij_i} - w_i + \epsilon\}, \quad (7)$$

$$\pi_i := w_i - \epsilon. \quad (8)$$

If $\lambda \leq a_{ij_i} - w_i + \epsilon$, add (i, j_i) to S and if j_i was assigned to some i' at the start of the iteration, remove from S the pair (i', j_i) .

3. A Forward/Reverse Auction Algorithm for Asymmetric Assignment Problems

Reverse Iteration

Find an unassigned object j with $p_j > \lambda$, its best person i_j

$$i_j = \arg \max_{i \in B(j)} \{a_{ij} - \pi_i\}, \quad (9)$$

the corresponding values

$$\beta_j = \max_{i \in B(j)} \{a_{ij} - \pi_i\}, \quad (10)$$

and

$$\gamma_j = \max_{i \in B(j), i \neq i_j} \{a_{ij} - \pi_i\}. \quad (11)$$

[If i_j is the only object in $B(j)$, we define γ_j to be ∞ or, for computational purposes, a number that is much larger than β_j .] Proceed according to the following two cases:

(1) $\beta_j \geq \lambda + \epsilon$. In this case set

$$p_j := \max\{\lambda, \gamma_j - \epsilon\}, \quad (12)$$

$$\pi_{i_j} := a_{i_j j} - \max\{\lambda, \gamma_j - \epsilon\}, \quad (13)$$

add (i_j, j) to S , and if i_j was assigned to some j' at the start of the iteration, remove from S the pair (i_j, j') .

(2) $\beta_j < \lambda + \epsilon$. In this case, set

$$p_j := \beta_j - \epsilon \quad (14)$$

and if the number of objects k with $p_k < \lambda$ is now more than $n - m$, set λ to the value

$$\min\{\xi \mid p_k \leq \xi \text{ for } n - m \text{ or more objects } k\}. \quad (15)$$

Note that λ remains unchanged in a forward iteration, and it can either decrease or stay unchanged in a reverse iteration. Note also that the “bidding object” j in the reverse iteration may not be assigned during the iteration; this happens when its “best value” β_j is low relative to λ , in which case its price p_j is reduced below λ [cf. Eq. (14)], and the object cannot bid again until λ decreases from its current level. Figure 1 illustrates the two cases that can arise in the reverse iteration.

The next proposition establishes a basic property of the forward and reverse iterations:

Proposition 2: Suppose that at the beginning of a forward or a reverse iteration, (π, p) satisfies together with S the first two ϵ -CS conditions (2a) and (2b), and λ satisfies condition (3). The same is true for (π, p) , S , and λ at the end of the iteration.

Proof: Suppose that the iteration starts with S , (π, p) , and λ satisfying Eqs. (2a), (2b), and (3). Let \bar{S} , $(\bar{\pi}, \bar{p})$, and $\bar{\lambda}$ be the corresponding quantities at the end of the iteration.

3. A Forward/Reverse Auction Algorithm for Asymmetric Assignment Problems

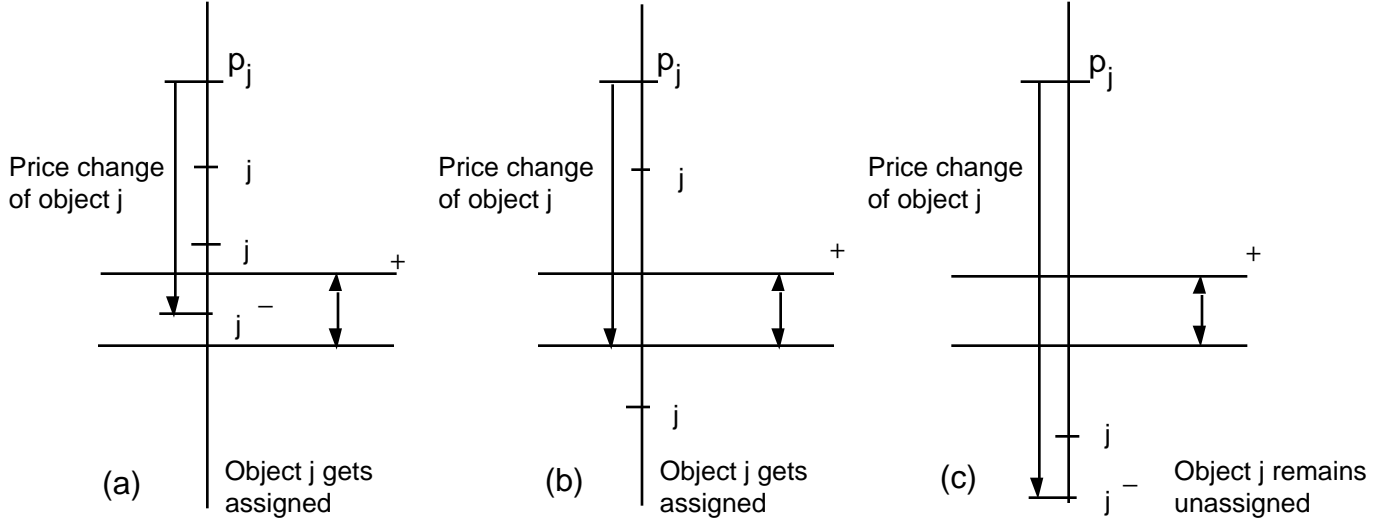


Figure 1: Illustration of the possible cases that can arise in the reverse iteration. These are:

- (a) $\beta_j \geq \lambda + \epsilon$ and $\gamma_j - \epsilon > \lambda$. Then j gets assigned to i_j , p_j is set to $\gamma_j - \epsilon$, and π_{i_j} is set to $a_{i_j j} - \gamma_j + \epsilon$.
- (b) $\beta_j \geq \lambda + \epsilon$ and $\gamma_j - \epsilon \leq \lambda$. Then j gets assigned to i_j , p_j is set to λ , and π_{i_j} is set to $a_{i_j j} - \lambda$.
- (c) $\beta_j < \lambda + \epsilon$. Then j stays unassigned and p_j is set to $\beta_j - \epsilon$, while π_{i_j} remains unchanged.

Consider first a forward iteration and let i be the corresponding unassigned person that submits the bid as per Eqs. (4)-(8). We will first verify that the pair $(\bar{\pi}, \bar{p})$ satisfies Eq. (2a) for each arc. From Eqs. (4)-(6), we see that $a_{i j_i} - w_i \geq p_{j_i}$, so by Eq. (7), we have

$$\bar{p}_{j_i} \geq p_{j_i} + \epsilon. \quad (16)$$

By adding this relation to the relation $\pi_k + p_{j_i} \geq a_{k j_i} - \epsilon$ [cf. Eq. (2a)], and by using the fact $\pi_k = \bar{\pi}_k$ for all $k \neq i$, we obtain

$$\bar{\pi}_k + \bar{p}_{j_i} \geq a_{k j_i} - \epsilon, \quad \forall k \in B(j_i), k \neq i. \quad (17)$$

On the other hand, since $p_j = \bar{p}_j$ for all $j \neq j_i$, we have

$$\bar{\pi}_i = w_i - \epsilon \geq a_{i j} - p_j - \epsilon = a_{i j} - \bar{p}_j - \epsilon, \quad \forall j \in A(i), j \neq j_i,$$

while from Eqs. (7) and (8), we have

$$\bar{\pi}_i = w_i - \epsilon \geq a_{i j_i} - p_{j_i} - \epsilon.$$

Combining the last two relations, we obtain

$$\bar{\pi}_i + \bar{p}_j \geq a_{i j} - \epsilon, \quad \forall j \in A(i). \quad (18)$$

3. A Forward/Reverse Auction Algorithm for Asymmetric Assignment Problems

Finally, for arcs (k, j) with $k \neq i$, $j \neq j_i$, we have $\pi_k = \bar{\pi}_k$ and $p_j = \bar{p}_j$, so by Eq. (2a), we obtain

$$\bar{\pi}_k + \bar{p}_j \geq a_{kj} - \epsilon, \quad \forall (k, j) \in \mathcal{A}, k \neq i, j \neq j_i. \quad (19)$$

By combining Eqs. (17)-(19), we see that

$$\bar{\pi}_k + \bar{p}_j \geq a_{kj} - \epsilon, \quad \forall (k, j) \in \mathcal{A},$$

that is, the ϵ -CS condition (2a) holds at the end of the iteration.

We next show that Eq. (2b) is preserved by the iteration, that is,

$$\bar{\pi}_k + \bar{p}_j = a_{kj}, \quad \forall (k, j) \in \bar{\mathcal{S}}. \quad (20)$$

Note that if $(k, j) \in \bar{\mathcal{S}}$ and $j \neq j_i$, we must have $k \neq i$, $(k, j) \in \mathcal{S}$, and $\pi_k = \bar{\pi}_k$, $p_j = \bar{p}_j$, so by using the hypothesis [cf. Eq. (2b)], we see that Eq. (20) holds. If on the other hand $(k, j_i) \in \bar{\mathcal{S}}$ for some k , we claim that $k = i$, since otherwise, by the rules of the iteration, we would have $\lambda > a_{ij_i} - w_i + \epsilon$, so that by Eqs. (7) and (16),

$$\lambda = \bar{p}_{j_i} \geq p_{j_i} + \epsilon,$$

contradicting the condition (3). Now if $(i, j_i) \in \bar{\mathcal{S}}$, we must have by the rules of the iteration, $\lambda \leq a_{ij_i} - w_i + \epsilon$ and $\bar{p}_{j_i} = a_{ij_i} - w_i + \epsilon$, so that

$$\bar{\pi}_i = w_i - \epsilon = a_{ij_i} - \bar{p}_{j_i}. \quad (21)$$

We see therefore that Eq. (20) holds for the case where $j = j_i$ as well.

Finally, to show that condition (3) is preserved by the iteration, note that $\lambda = \bar{\lambda}$, that $p_j = \bar{p}_j$ for all $j \neq j_i$, and that $\bar{p}_{j_i} \geq \lambda$ [cf. Eq. (7)]. Since the only object that can become assigned during the iteration is j_i , we see that

$$\bar{p}_j \geq \bar{\lambda}, \quad \forall j \text{ that are assigned under } \bar{\mathcal{S}}.$$

The proof of the proposition for the case of a forward iteration is thus complete.

Consider next the case of a reverse iteration. Let j be the corresponding unassigned object that submits the bid as per Eqs. (9)-(13).

In the case where $\beta_j \geq \lambda + \epsilon$, we have by Eqs. (12) and (13)

$$\bar{\pi}_{i_j} = a_{i_j j} - \max\{\lambda, \gamma_j - \epsilon\} \geq a_{i_j j} - \beta_j + \epsilon = \pi_{i_j} + \epsilon \quad \text{if } \beta_j \geq \lambda + \epsilon. \quad (22)$$

By using also the relation $\pi_{i_j} + p_j \geq a_{i_j j} - \epsilon$, we have

$$\bar{p}_j = \max\{\lambda, \gamma_j - \epsilon\} \leq \beta_j - \epsilon = a_{i_j j} - \pi_{i_j} - \epsilon \leq p_j \quad \text{if } \beta_j \geq \lambda + \epsilon. \quad (23)$$

3. A Forward/Reverse Auction Algorithm for Asymmetric Assignment Problems

In the case where $\beta_j < \lambda + \epsilon$ we have $\gamma_j \leq \beta_j < \lambda - \epsilon$ and by using also the fact $p_j > \lambda$, we obtain

$$\bar{p}_j = \beta_j - \epsilon < \lambda < p_j \quad \text{if } \beta_j < \lambda + \epsilon, \quad (24)$$

$$\bar{\pi}_{i_j} = \pi_{i_j} \quad \text{if } \beta_j < \lambda + \epsilon. \quad (25)$$

To prove that the ϵ -CS condition (2a) is preserved by the reverse iteration, consider first the arcs (i, j) with $i \neq i_j$. Since $\bar{\pi}_i = \pi_i$ and $p_j \geq \gamma_j - \epsilon$, we have

$$\bar{\pi}_i + \bar{p}_j \geq \pi_i + \gamma_j - \epsilon \geq \pi_i + a_{ij} - \pi_i - \epsilon = a_{ij} - \epsilon, \quad \forall i \in B(j), i \neq i_j. \quad (26)$$

Consider next the arcs (i_j, k) with $k \neq j$. We have $\pi_{i_j} + p_k \geq a_{i_j k} - \epsilon$ [cf. Eq. (2a)], and since $p_k = \bar{p}_k$ for $k \neq j$, we obtain [cf. Eq. (22)]

$$\bar{\pi}_{i_j} + \bar{p}_k \geq \pi_{i_j} + \epsilon + p_k \geq a_{i_j k} \quad \text{if } \beta_j \geq \lambda + \epsilon, k \neq j, \quad (27)$$

and [cf. Eq. (25)]

$$\bar{\pi}_{i_j} + \bar{p}_k = \pi_{i_j} + p_k \geq a_{i_j k} - \epsilon \quad \text{if } \beta_j < \lambda + \epsilon, k \neq j. \quad (28)$$

Finally for the arc (i_j, j) , we have by Eq. (13),

$$\bar{\pi}_{i_j} + \bar{p}_j = a_{i_j k} \quad \text{if } \beta_j \geq \lambda + \epsilon, \quad (29)$$

and by Eq. (14),

$$\bar{\pi}_{i_j} + \bar{p}_j = \pi_{i_j} + \beta_j - \epsilon = a_{i_j j} - \epsilon \quad \text{if } \beta_j < \lambda + \epsilon. \quad (30)$$

By combining Eqs. (26)-(30), we see that

$$\bar{\pi}_i + \bar{p}_k \geq a_{ik} - \epsilon, \quad \forall (i, k) \in \mathcal{A},$$

so the condition (2a) is preserved by the reverse iteration.

To show that Eq. (2b) is preserved by the reverse iteration, that is,

$$\bar{\pi}_i + \bar{p}_k = a_{ik}, \quad \forall (i, k) \in \bar{\mathcal{S}}, \quad (31)$$

note that if $(i, k) \in \bar{\mathcal{S}}$ and $i \neq i_j$, we must have $\pi_i = \bar{\pi}_i$, $p_k = \bar{p}_k$, and $(i, k) \in S$, so by using the hypothesis [cf. Eq. (2b)], we see that Eq. (31) holds. If on the other hand $(i_j, k) \in \bar{\mathcal{S}}$ for some k , then either $k = j$ in which case we must have $\bar{\pi}_{i_j} + \bar{p}_j = a_{i_j j}$ by Eqs. (12) and (13), or else $k \neq j$, in which case $(i_j, k) \in S$, $\bar{\pi}_{i_j} = \pi_{i_j}$, and $\bar{p}_k = p_k$, so by Eq. (2b) and the induction hypothesis we have

3. A Forward/Reverse Auction Algorithm for Asymmetric Assignment Problems

$\bar{\pi}_{i_j} + \bar{p}_k = a_{i_j k}$. Thus Eq. (31) holds in all cases and the condition (2b) is preserved by the reverse iteration.

Finally to show that Eq. (3) is preserved by the reverse iteration, note that $\lambda \geq \bar{\lambda}$, while the only object that can become assigned during the iteration and whose price can change is j . On the other hand if j becomes assigned, we must have $\bar{p}_j \geq \lambda$ by Eq. (12), so at the end of the iteration, we will have $\bar{p}_j \geq \bar{\lambda}$, thereby preserving Eq. (3). **Q.E.D.**

As a corollary of the preceding proof, we obtain the following proposition.

Proposition 3: Suppose that S , (π, p) , and λ satisfy conditions (2a), (2b), and (3). Then:

- (a) In a forward iteration, p_{j_i} increases by at least ϵ . Furthermore, either j_i is assigned to i during the iteration and p_{j_i} is increased to a level no less than λ , or else p_{j_i} is increased to the level λ .
- (b) In a reverse iteration, either π_{i_j} increases by at least ϵ and j becomes assigned, or else j remains unassigned and p_j decreases to a level below λ .
- (c) If all persons are assigned (S is feasible), the reverse iteration leaves λ unchanged.

Proof: (a) See Eqs. (7) and (16).

(b) See Eqs. (22) and (24).

(c) If all persons are assigned, the number of assigned objects k is m and all these objects satisfy $p_k \geq \lambda$ by Eq. (3). Therefore, the number of objects k with $p_k < \lambda$ cannot become more than $n - m$, which is the only situation where λ can change. **Q.E.D.**

We will now use the results obtained so far to analyse several possible algorithms.

3.2 Purely Forward Algorithm

It is possible to consider a forward auction algorithm that consists exclusively of forward iterations. In such an algorithm it is essential to choose initially $\lambda \geq p_j$ for all unassigned objects j . Then, since λ will remain unchanged, by using Prop. 3(a), it can be seen that in the course of the algorithm, we will have

$$\max_{k: \text{ unassigned under } S} p_k \leq \lambda \leq \min_{k: \text{ assigned under } S} p_k, \quad \forall j : \text{ unassigned under } S, \quad (32)$$

so by using also Prop. 2, we see that all three ϵ -CS conditions (2a)-(2c) will be satisfied. Furthermore, by Prop. 3(a), the price p_{j_i} is increased by at least ϵ at each forward iteration. Using this fact and

3. A Forward/Reverse Auction Algorithm for Asymmetric Assignment Problems

standard arguments (see e.g. [Ber88], [Ber91]), it can be shown that this forward algorithm will terminate with a feasible assignment S that satisfies ϵ -CS together with (π, p) (and is optimal if $\epsilon < 1/m$ and the problem data are integer). Unfortunately, even though this forward algorithm will work with arbitrary initial prices, it is not suitable for use in conjunction with ϵ -scaling because of the initial requirement that $\lambda \geq p_j$ for all unassigned objects j . Since for an object to get assigned its price must rise to at least the level λ , the advantage of approximately optimal initial prices that ϵ -scaling attempts to carry from one ϵ -scaling phase to the next is largely diminished.

3.3 Purely Reverse Algorithm

It is also possible to consider a purely reverse auction algorithm that consists exclusively of reverse iterations, provided that the initial assignment is feasible and the initial λ is such that condition (3) is satisfied ($\lambda \geq p_j$ for all assigned objects j). The following proposition establishes the validity of the algorithm.

Proposition 4: For a feasible problem, the purely reverse algorithm starting from a feasible assignment, a pair (π, p) , and a scalar λ satisfying conditions (2a), (2b), and (3) terminates. The assignment obtained satisfies ϵ -CS together with (π, p) .

Proof: From Prop. 3(c), we have that λ will remain unchanged and that at each iteration there are two possibilities: (1) π_{i_j} will increase by ϵ and the selected unassigned object j will get assigned to i_j , or (2) The number of unassigned objects whose price exceeds λ will decrease by one. Therefore, after some iteration, case (1) will occur exclusively. By Eq. (13) we have

$$\bar{\pi}_{i_j} = a_{i_j j} - \max\{\lambda, \gamma_j - \epsilon\} \leq a_{i_j j} - \lambda, \quad (33)$$

so π_{i_j} cannot exceed $\max_{(i,k) \in \mathcal{A}} a_{ik} - \lambda$. It follows that the algorithm cannot execute an infinite number of iterations and must therefore terminate. **Q.E.D.**

The disadvantage of the purely reverse algorithm is that it requires an initial feasible assignment. The reason is that if the current assignment S is infeasible, we may have $p_j \leq \lambda$ for all unassigned objects j , while we have $p_j < \lambda$ for no more than $n - m$ unassigned objects. Then the purely reverse algorithm will leave λ unchanged and will terminate without finding a feasible solution. A possible remedy is to start with an arbitrary assignment but to reduce λ by some positive increment whenever the difficulty just described occurs. Unfortunately, however, it is not easy to determine the proper size of the increment for fast termination.

Another possibility to circumvent the need for an initial feasible assignment is to combine the forward and reverse algorithms, so that the forward part guarantees that a feasible assignment will

3. A Forward/Reverse Auction Algorithm for Asymmetric Assignment Problems

be obtained, while the reverse part is capable of dealing with essentially arbitrary starting values of λ . In particular, one may use the purely forward algorithm first to obtain a feasible assignment, and then switch to the reverse algorithm after setting

$$\lambda = \min_{k: \text{ assigned under } S}.$$

This is the algorithm proposed in [BCT91]; it is suitable for ϵ -scaling but does not take advantage of the beneficial effect of mixing the forward and the reverse algorithms that was demonstrated for symmetric problems in [BCT91]. The following algorithm switches several times between the two algorithms aiming at less reliance on ϵ -scaling and faster termination.

3.4 Combined Forward/Reverse Algorithm

The combined forward/reverse algorithm that we now introduce switches between forward and reverse auction until all persons are assigned. Then it executes reverse iterations exclusively, aiming to satisfy the final remaining optimality condition ($p_j \leq \lambda$ for all unassigned objects j). The initial S , (π, p) , and λ must satisfy the ϵ -CS conditions (2a), (2b), and the condition $\lambda \leq p_j$ for all assigned objects j . Thus if the initial assignment is empty, any initial p and λ can be used. We assume that initially there is at least one unassigned person (otherwise the forward part of the algorithm is inapplicable and unnecessary).

Combined Forward/Reverse Auction Algorithm

Step 1: (Forward auction cycle) Execute iterations of the forward auction algorithm until at least one more person becomes assigned. If there is an unassigned person left, go to Step 2; else go to Step 3.

Step 2: (Reverse auction cycle) Execute several iterations of the reverse auction algorithm until at least one more object becomes assigned or until we have $p_j \leq \lambda$ for all unassigned objects j . If there is an unassigned person left, go to Step 1; else go to Step 3.

Step 3: (Reverse auction) Execute successive iterations of the reverse auction algorithm until the algorithm terminates with $p_j \leq \lambda$ for all unassigned objects j .

The following proposition establishes the validity of the algorithm.

Proposition 5: For a feasible problem, the combined forward/reverse algorithm terminates with an optimal assignment.

Proof: We will assume that the algorithm does not terminate and will arrive at a contradiction. When the algorithm obtains a feasible assignment, it gets reduced to the purely reverse algorithm and terminates by Prop. 4. Assume therefore that the algorithm never obtains a feasible assignment.

3. A Forward/Reverse Auction Algorithm for Asymmetric Assignment Problems

Since the cardinality of the assignment must increase before switching from a forward to a reverse cycle, there are two possibilities: (1) The algorithm will execute forward iterations exclusively after some iteration, or (2) The algorithm will execute reverse iterations exclusively after some iteration, and we will always have $p_j > \lambda$ for some unassigned object j .

In case (1) the algorithm will be reduced to the purely forward algorithm, and as discussed earlier, it must terminate for a feasible problem. This contradicts our earlier hypothesis that the algorithm does not terminate.

In case (2), since whenever a profit variable increases it increases by at least ϵ , there are two possibilities:

- (a) After some iteration, all π_i stay constant and no object changes assignment.
- (b) Some profit variable increases to ∞ , in which case by the argument given in the proof of Prop. 4 [cf. Eq. (33)], λ decreases to $-\infty$.

In case (a), the variables β_j stay constant after some iteration, so in view of Eq. (14), the object prices cannot change after some iteration. This contradicts Prop. 3(b), which states that $\bar{p}_j < p_j$ at each reverse iteration [see also Eq. (24)].

In case (b), let

$$J_\infty = \{j \mid p_j \rightarrow -\infty\}, \quad \bar{J}_\infty = \{j \mid j \notin J_\infty\},$$

$$I_\infty = \{i \mid \pi_i \rightarrow \infty\}, \quad \bar{I}_\infty = \{i \mid i \notin I_\infty\}.$$

By the ϵ -CS condition (2a), we must have

$$i \in I_\infty \quad \Rightarrow \quad j \in J_\infty \quad \forall (i, j) \in \mathcal{A}, \quad (34)$$

$$i \in \bar{I}_\infty \quad \Rightarrow \quad j \in \bar{J}_\infty \quad \forall (i, j) \in \mathcal{A}. \quad (35)$$

We claim that after some iteration, each of the objects in \bar{J}_∞ must be assigned at all times to the same person from \bar{I}_∞ . To see this, note that if some object $j \in \bar{J}_\infty$ bids an infinite number of times for some person i_j , then π_{i_j} will increase by at least ϵ an infinite number of times, in view of Prop. 3(b), the definition of \bar{J}_∞ , and the fact $\lambda \rightarrow -\infty$. On the other hand by Eq. (35), we must have $i_j \in \bar{I}_\infty$, so π_{i_j} must remain bounded and we have a contradiction.

Thus, \bar{I}_∞ contains the set of persons that are assigned to \bar{J}_∞ plus the nonempty set of persons that are unassigned throughout the last reverse cycle (a person that becomes assigned in a reverse cycle remains assigned for the duration of the cycle). Therefore, the number of persons in \bar{I}_∞ exceeds the number of objects in \bar{J}_∞ . In view of Eq. (35), this contradicts the hypothesis that the problem is feasible. **Q.E.D.**

3. A Forward/Reverse Auction Algorithm for Asymmetric Assignment Problems

A careful examination of the preceding proof shows that there are other valid variations of the combined/forward reverse algorithm, corresponding to variations of the reverse iterations and/or the scheme for switching from a reverse to a forward cycle. What is important is that: (a) λ should remain unchanged at all forward iterations and at all reverse iterations where the current assignment is feasible, (b) λ should not increase during all reverse iterations (again here only unassigned objects j with $p_j > \lambda$ should be allowed to bid), and (c) a mechanism is provided whereby the combined method is guaranteed to eventually exit from a reverse cycle if the current assignment is not feasible.

Consider now what happens if the problem is infeasible. Then, eventually the number of unassigned persons will stop decreasing and the method will get caught in either a forward cycle (Step 1) or in a reverse cycle (Step 2). Infeasibility will then be detected in the standard way for auction algorithms, that is, some price or some profit will exceed a certain precomputable upper bound, as described in [Ber91]. It is also possible to deal with infeasibility by adding a sufficient number of artificial arcs to convert the problem to a feasible problem. These arcs must have sufficiently small values to guarantee that they are not part of an optimal assignment unless the original problem is infeasible; see [Ber91].

3.5 An Alternative Reverse Iteration and Combined Forward/Reverse Algorithm

A variation of the reverse iteration is obtained if we keep λ constant, even if the number of objects k with $p_k < \lambda$ becomes greater than $n - m$. Thus, this alternative iteration is defined to be identical to the one given earlier except that we forego the change of λ in case (2) [cf. Eq. (15)]. For this iteration, Props. 2 and 3 still hold, but the purely reverse algorithm may terminate with some persons still unassigned because λ was set to a value so high that the number of unassigned objects with price less or equal to λ exceeds $n - m$. Nonetheless, if the alternative iteration is combined with the forward iteration as in the algorithm given earlier, the resulting combination is valid because forward iterations will continue as long as there are some unassigned persons, even if no reverse iterations can be executed.

Note that λ remains unchanged throughout this alternative combined forward/reverse algorithm, and can only change at the beginning of each scaling phase. Thus, the choice of λ at each scaling phase is critical for the algorithm's performance. A reasonable scheme is to choose λ at the beginning of each scaling phase except the first as

$$\lambda = \min_{j: \text{assigned under } S} p_j,$$

where S is the assignment obtained at the end of the preceding scaling phase. At the first scaling phase one may start with the empty assignment, zero object prices, and $\lambda = 0$. With these choices,

no reverse iterations will be executed in the first scaling phase, since the prices of the unassigned objects as well as λ will remain at zero throughout the phase.

4. COMPUTATIONAL RESULTS

In order to evaluate the relative performance of the new forward/reverse auction algorithms, we implemented both variations of the combined forward-reverse auction algorithms discussed previously and the auction algorithm for inequality constraints from [BCT91]. The three algorithms were evaluated on the following classes of inequality-constrained assignment problems:

1. Randomly generated problems, generated using the DIMACS assign.c problem generator [Cas92], which also include a number of high-cost arcs.
2. Geometric matching problems, consisting of matching a list of two-dimensional points with a randomly-perturbed copy of the same list.
3. Clustered geometric matching problems, consisting of matching a list of clustered two-dimensional points with a randomly perturbed copy of the same list.

The last two classes are representative of an important class of applications which motivated this research: data association in multi-object tracking. In these problems, new sensor measurements at each time frame must be associated with the predicted position of existing tracks. Due to the presence of false alarms (due to clutter or other effects), missed detections and sensor measurement inaccuracies, the set of measurement values will be a random perturbation of the set of predicted positions. The maximum likelihood problem of determining which measurement-track associations are most likely is equivalent to an inequality-constrained assignment problem.

4.1 Results on Random Problems

Table 1 summarizes the results of our random experiments for inequality-constrained assignment problems with 2000 persons. In these experiments, an initial random problem is generated with 8 arcs per person, with cost range [1,200]. Based on the results of [BCT91], purely random problems are often easy to solve and require no scaling; in order to make scaling necessary, we modified the problems to increase the costs of 20% of the arcs by a factor of 100. The resulting inequality-constrained assignment problems have a difficult structure which requires scaled auction algorithms, as discussed in [BCT91]. The AS algorithm is the forward-reverse algorithm of [BCT91] discussed in Section 3.3, which uses a scaled forward auction algorithm to find a complete assignment of

4. Computational Results

persons into objects and a set of dual prices satisfying conditions (2a), (2b) and (3), and then uses a purely reverse algorithm to find an optimal assignment. The ASFR1 algorithm is the combined forward-reverse algorithm of Section 3.4, and the ASFR2 is the combined forward-reverse algorithm of Section 3.5.

Table 2 summarizes the results of random experiments with 4000 person, degree 8 problems, with cost range [1,200], with 20% of the arc costs increased by an additional factor of 100. The results in these two tables indicate little difference in the performance of the ASFR1 and ASFR2 algorithms. The results also indicate the superiority of the new forward-reverse auction algorithms (ASFR) over the previous algorithm (AS) of [BCT91].

Problem	2000 x 2020	2000 x 2050	2000 x 2100	2000 x 2200
AS	3.05	2.50	2.00	0.53
ASFR1	1.13	0.84	0.69	0.50
ASFR2	1.14	0.83	0.67	0.50

Table 1: Average run time in seconds for 10 problems on the NeXTStation 68040 for 2000 person, degree 8 random problems.

Problem	4000 x 4040	4000 x 4100	4000 x 4200	4000 x 4400
AS	9.86	7.85	6.54	2.39
ASFR1	4.43	3.57	2.90	1.14
ASFR2	4.45	3.56	2.92	1.16

Table 2: Average run time in seconds for 10 problems on the NeXTStation 68040 for 4000 person, degree 8 random problems.

Table 3 contains the results of experiments with 4000 x 4400 person assignment problems as a

4. Computational Results

function of increasing density for two classes of problems: “easy” problems where the arc costs are randomly selected uniformly in $[1,20000]$ and “hard” problems where the arc costs are selected uniformly in $[1,200]$, and 20% of the arc costs are increased by a factor of 100. As the times indicate, the ASFR and ASFR2 algorithms offer little advantage over the old AS algorithm for easy problems; these problems do not create “price wars” among persons, and can be solved without the use of scaling. For the “hard” problems, the new forward-reverse auction algorithms are much faster than the AS algorithm. Interestingly, the advantage of the new algorithms seems to increase with increasing problem density. This result is somewhat surprising, since increasing density often results in shorter price wars because the high-cost arcs can be ignored. These results highlight the performance advantages of the new forward-reverse auction algorithms.

Problem	Degree 8	Degree 16	Degree 32	Degree 64
AS - easy	0.36	0.57	1.17	2.40
ASFR - easy	0.40	0.68	1.46	2.91
ASFR2 - easy	0.42	0.72	1.43	2.91
AS - hard	2.39	8.27	17.32	38.95
ASFR - hard	1.18	2.48	5.17	10.42
ASFR2 - hard	1.16	2.50	5.18	10.40

Table 3: Average run times in seconds for 10 problems on the NeXTStation 68040 for 4000 person, 4400 object random problems with increasing degree.

4.2 Results on Geometric Matching Problems

The geometric matching problems were generated to simulate the structure of data association problems arising in multiobject tracking. An initial number of points was randomly generated using a uniform distribution on the square 10^5 by 10^5 . From these initial points, two lists of points were generated according to the following rules:

1. The first list was generated by accepting each point of the initial list with probability 0.95,

independent across points. This effect was chosen to simulate missed detections. Thus, the first list is a reduction of the initial list of points.

2. The second list was generated by first accepting each point of the initial list with probability 0.95, independent across points and across the events used to generate the first list. This effect was chosen to simulate false alarms in the data set. Then, the locations of all the points in the second list were shifted by a constant bias, which was randomly generated from a bivariate Gaussian distribution with a specified bias standard deviation. Subsequently, each point in the second list was shifted by an independent bivariate Gaussian random variable, representing measurement noise, with a specified measurement standard deviation.
3. The list with the least number of points was selected to be the persons in the inequality assignment problem. The other list was selected to be the objects. Arcs were created between each person-object pair for which the Euclidean distance between the corresponding points was less than 3 times the measurement standard deviation. The costs assigned to each arc were integers between 1 and 1000, proportional to the Euclidean distance of the corresponding person-object pair.
4. In order to guarantee feasibility of the inequality assignment problem, an extra object node was introduced for each person node, with a corresponding arc cost of 20000. This large cost encouraged the problem to find a feasible assignment without using the extra nodes.

Table 4 summarizes our results with random geometric experiments corresponding to 2000 points in the initial list. Four different combinations of bias/measurement standard deviation were tested. For each combination, Table 4 lists the average run times across 10 different problems with similar statistics for each of the three algorithms. For small bias/measurement standard deviations, the points in the person lists and object lists are far apart, and thus the solution of the assignment problem is trivial. As the standard deviation increases, object-person groups form with an unbalanced number of persons or real objects in the group (because of the missed detection and false alarm probability 0.95). This creates long price wars to determine which extra person in the group will be assigned to an artificial object, or objects will remain unassigned. The size of these groups increases with bias/measurement standard deviations, leading to longer price wars.

As the results in Table 4 indicate, the new forward-reverse algorithms are much more efficient than the AS algorithm of [BCT91]. The reason for this efficiency is that forward and reverse iterations are interleaved at each scaling step; in contrast, the AS algorithm performs only forward iterations at most scales, and then switches to an unscaled reverse-only algorithm. Most of the computation time (over 95 %) is spent in this unscaled reverse-only algorithm trying to enforce condition (2c) for groups with more objects than persons. The new forward-reverse algorithms use scaling both for forward and reverse iterations, resulting in more robust performance for this class of problems.

Bias/Measurement SD	5/50	10/100	15/150	20/200
AS	0.26	30.37	378.42	374.03
ASFR	0.32	1.67	3.78	9.34
ASFR2	0.34	1.62	3.04	5.99

Table 4: Average run times in seconds on the NeXTStation 68040 across 10 geometric problems with 2000 points per list, probability of detection 0.95.

4.3 Results on Clustered Geometric Matching Problems

The clustered geometric matching problems were generated to simulate a different type of data association problems arising in multiobject tracking: groups of objects moving close together, but with significant distance among the groups. The principal difference between this class of problems and the geometric class of problems is the location of the initial number of points, which are generated as follows:

1. An initial number of cluster centers are generated with a uniform distribution on the square 10^5 by 10^5 .
2. For each cluster center, a fixed number of points is generated by adding to the cluster center an independent bivariate Gaussian random variable with a specified cluster spread standard deviation.

Once the initial list of points is available, generation of the inequality constrained assignment problem follows identically steps 1-5 of the geometric matching problems of the previous section.

Table 5 summarizes our results with random geometric experiments corresponding to 50 clusters of 40 points each in the initial list. Four different combinations of spread/measurement/bias standard deviation were tested. For each combination, Table 5 lists the average run times across 10 different problems with similar statistics for each of the three algorithms. For small spread/measurement/bias standard deviations, the assignment problem decouples by cluster, and thus corresponds to solution of 50 small problems. As the spread standard deviation increases, the assignment problems become coupled across clusters, and finding an optimal assignment becomes harder, and more susceptible to long price wars.

The results in Table 5 agree closely with the results from Table 4. The performance of the

new forward-reverse algorithms is much more robust to scenario variations than the performance of the AS algorithm of [BCT91]. This is due largely to the use of scaling both in forward and reverse iterations, and the mixing of forward and reverse iterations at each scale. In essence, this mixing of forward and reverse iterations forces the object prices and person profits to satisfy the complementary slackness condition (2c) locally for each cluster and at each scale. In contrast, the AS algorithm tries to enforce this condition only at the finest scale, and then using a single value of λ for all the clusters; if the prices in one cluster rise much higher than the prices in other clusters (because of price wars), the reverse-only part of the AS algorithm may require an excessive number of bids to satisfy condition (2c).

Spread/Measurement/Bias SD	500/50/5	1000/100/10	2500/150/15	2000/200/20
AS	1.04	28.33	225.60	813.92
ASFR	0.45	1.70	2.21	3.68
ASFR2	0.49	1.74	1.94	3.53

Table 5: Average run times in seconds on the NeXTStation 68040 across 10 clustered geometric problems with 2000 points per list, probability of detection 0.95.

REFERENCES

- [Ber79] Bertsekas, D. P., “A Distributed Algorithm for the Assignment Problem,” Lab. for Information and Decision Systems Working Paper, M.I.T., March 1979.
- [Ber85] Bertsekas, D. P., “A Distributed Asynchronous Relaxation Algorithm for the Assignment Problem,” Proc. 24th IEEE Conf. Dec. & Contr., 1985, pp. 1703-1704.
- [Ber88] Bertsekas, D. P., “The Auction Algorithm: A Distributed Relaxation Method for the Assignment Problem,” Annals of Operations Research, Vol. 14, 1988, pp. 105-123.
- [Ber90] Bertsekas, D. P., “The Auction Algorithm for Assignment and Other Network Flow Problems: A Tutorial,” Interfaces, Vol. 20, 1990, pp. 133-149.

- [Ber91] Bertsekas, D. P., *Linear Network Optimization: Algorithms and Codes*, M.I.T. Press, Cambridge, MA., 1991.
- [BCT91] Bertsekas, D. P., Castañon, D. A., and Tsaknakis, H., "Reverse Auction and the Solution of Inequality Constrained Assignment Problems," Alphatech Report, Burlington, MA, March 1991 (submitted for publication).
- [BeC89] Bertsekas, D. P., and Castañon, D. A., "Parallel Synchronous and Asynchronous Implementations of the Auction Algorithm," Alphatech Report, Burlington, MA, Nov. 1989; also *Parallel Computing*, Vol. 17, 1991, pp. 707-732.
- [BeE88] Bertsekas, D. P., and Eckstein, J., "Dual Coordinate Step Methods for Linear Network Flow Problems," *Math. Progr., Series B*, Vol. 42, 1988, pp. 203-243.
- [Cas92] Castañon, D. A., "Reverse Auction Algorithms for Assignment Problems," to appear in *DIMACS Series in Discrete Mathematics and Computer Science*.
- [KKZ89] Kempa, D., Kennington, J., and Zaki, H., "Performance Characteristics of the Jacobi and Gauss-Seidel Versions of the Auction Algorithm on the Alliant FX/8," Report OR-89-008, Dept. of Mech. and Ind. Eng., Univ. of Illinois, Champaign-Urbana, 1989.
- [KeH80] Kennington, J., and Helgason, R., *Algorithms for Network Programming*, Wiley, N. Y., 1980.
- [PaS82] Papadimitriou, C. H., and Steiglitz, K., *Combinatorial Optimization: Algorithms and Complexity*, Prentice-Hall, Englewood Cliffs, N. J., 1982.
- [PhZ88] Phillips, C., and Zenios, S. A., "Experiences with Large Scale Network Optimization on the Connection Machine," Report 88-11-05, Dept. of Decision Sciences, The Wharton School, Univ. of Pennsylvania, Phil., Penn., Nov. 1988.
- [Roc84] Rockafellar, R. T., *Network Flows and Monotropic Programming*, Wiley-Interscience, N. Y., 1984.
- [WeZ90] Wein, J., and Zenios, S. A., "Massively Parallel Auction Algorithms for the Assignment Problem," *Proc. of 3rd Symposium on the Frontiers of Massively Parallel Computation*, Md., pp. 90-99, Nov. 1990.
- [WeZ91] Wein, J., and Zenios, S. A., "On the Massively Parallel Solution of the Assignment Problem," *J. of Parallel and Distributed Computing*, Vol. 13, 1991, pp. 228-236.
- [Zak90] Zaki, H., "A Comparison of Two Algorithms for the Assignment Problem," Report ORL 90-002, Dept. of Mechanical and Industrial Engineering, Univ. of Illinois, Urbana, Ill.