

Auction Algorithms for Network Flow Problems: A Tutorial Introduction

DIMITRI P. BERTSEKAS

Laboratory For Information and Decision Systems, Massachusetts Institute of Technology, Cambridge, MA 02139

Received January 31, 1992, Revised February 3, 1992.

Abstract. This paper surveys a new and comprehensive class of algorithms for solving the classical linear network flow problem and its various special cases such as shortest path, max-flow, assignment, transportation, and transshipment problems. The prototype method, from which the other algorithms can be derived, is the auction algorithm for the assignment problem. This is an intuitive method that operates like a real auction where persons compete for objects by raising their prices through competitive bidding; the prices can be viewed as dual variables. Conceptually, auction algorithms represent a significant departure from the cost improvement idea that underlies primal simplex and dual ascent methods; at any one iteration, they may deteriorate both the primal and the dual cost. Auction algorithms perform very well for several important types of problems, both in theory and in practice, and they are also well suited for parallel computation.

Keywords: Network programming, auction, assignment, transportation, shortest path.

1. Introduction

The classical algorithms for solving linear network flow problems are primal cost improvement methods, including simplex methods, which iteratively improve the primal cost by moving flow around simple cycles, and dual ascent methods, which iteratively improve the dual cost by changing the prices of a subset of nodes by equal amounts.

Auction algorithms, the subject of this tutorial paper, are fundamentally different; they depart from the cost improvement idea and at any one iteration, they may deteriorate both the primal and the dual cost, although in the end they find an optimal primal solution. Their origin lies in nondifferentiable optimization (where nonmonotonic subgradient-based algorithms are common), and in the ϵ -subgradient method of Bertsekas and Mitter ([21]) in particular (where the progress of the method depends on gradually reducing an ϵ -parameter to an acceptable tolerance level). Auction algorithms are also highly intuitive and easy to understand; they can be explained in terms of economic competition concepts, which are couched on everyday experience.

Auction algorithms originated in 1979 with an algorithm proposed by the author for the classical assignment problem ([7]). This algorithm was further

developed in [6], [8], and [20]. The motivation was to solve the problem by using parallelism in a natural way. It turned out, however, that the resulting method was very fast in a serial environment as well. Subsequent work extended the auction algorithm to other linear network flow problems. In particular, an extension to the minimum cost flow problem, the ϵ -relaxation method, was given by the author in [9] and [10]. An auction algorithm for transportation problems was given by the author in collaboration with D. Castañón in [14]. An auction algorithm for shortest paths was given by the author in [5]. Finally, a max-flow algorithm, developed by Goldberg ([31]) from an entirely different point of view, fundamentally involves auction ideas and can be viewed as the natural extension of the auction algorithm to the max-flow problem.

All of the algorithms just mentioned can be derived from the original 1979 auction algorithm for the assignment problem. The derivation is based on well-known transformations of the general linear minimum cost flow problem and its various special cases (e.g., shortest path, max-flow) to equivalent assignment problems. Once the auction algorithm is applied to the corresponding equivalent assignment problem and the computations are streamlined, one obtains an auction algorithm for the original problem.

The purpose of this paper is to provide a tutorial introduction to auction algorithms. For textbook presentations, we refer the reader to [11] and [23]. We will initially focus on the basic algorithm for the assignment problem and then discuss various extensions to other problems. In Sections 2 and 3, we develop the auction algorithm for the symmetric assignment problem. In particular, we discuss a number of issues that are common to all types of auction algorithms, including ϵ -scaling and dealing with infeasibility. We also discuss an alternative form of the auction algorithm, called *reverse auction*; while in regular auction persons bid for objects by raising their prices, in reverse auction objects compete for persons by essentially offering discounts. In Section 4 we develop auction algorithms for shortest path problems. In Section 5 we discuss the extension of the auction algorithm to transportation problems, and in Section 6 we give a generic auction algorithm for the minimum cost network flow problem. In Section 7 we develop the ϵ -relaxation method for the minimum cost flow problem as a special case of the generic algorithm. The ϵ -relaxation method is specialized to the max-flow problem in Section 8, yielding algorithms similar to those of Goldberg and Tarjan ([31] and [32]). In Section 9 we discuss specialized auction algorithms for inequality-constrained problems such as the asymmetric assignment problem. Finally, in Section 10 we discuss the computational aspects of auction algorithms, including issues of parallel and asynchronous implementation.

2. Assignment by Naive Auction

In the classical *symmetric assignment problem* there are n persons and n objects that we have to match on a one-to-one basis. There is a benefit a_{ij} for matching person i with object j and we want to assign persons to objects so as to maximize the total benefit. We are given a set \mathcal{A} of pairs (i, j) that can be matched. For each person i , we denote by $A(i)$ the set of objects that can be matched with i

$$A(i) = \{j \mid (i, j) \in \mathcal{A}\}$$

and for each object j , we denote by $B(j)$ the set of persons that can be matched with j

$$B(j) = \{i \mid (i, j) \in \mathcal{A}\}$$

By an *assignment* we mean a set S of person-object pairs (i, j) , such that each person i and each object j is involved in at most one pair from S . If the number of pairs in S is n , so that every person is assigned to a distinct object, we say that S is *feasible*; otherwise S is said to be *infeasible*. If a feasible assignment exists, the problem is said to be feasible, and otherwise it is said to be infeasible. We seek a feasible assignment [a set of person-object pairs $(1, j_1), \dots, (n, j_n)$ from \mathcal{A} , such that the objects j_1, \dots, j_n are all distinct], which is optimal in the sense that it maximizes the total benefit $\sum_{i=1}^n a_{ij_i}$.

The symmetric assignment problem should be distinguished from its *asymmetric* version where the number of objects is larger than the number of persons, and while there is a requirement to assign every person, there is no such requirement for the objects. Auction algorithms for asymmetric and other related assignment problems will be discussed in Section 9.

The assignment problem is important in many practical contexts, but it is also of great theoretical importance. Despite its simplicity, it embodies a fundamental linear programming structure. One of the most important types of linear programming problems the minimum cost network flow problem, can be reduced to the assignment problem by means of a simple reformulation (see, e.g., [11], p. 17; [23], p. 335; [44], p. 149; and Section 6). Thus, any method for solving the assignment problem can be generalized to solve the minimum cost flow problem. For this reason, the assignment problem has served as a convenient starting point for important algorithmic ideas in linear programming. For example, the primal-dual method ([29] and [42]) was motivated and developed through Kuhn's Hungarian method ([36]), the first specialized method for the assignment problem.

To develop an intuitive understanding of the auction algorithm, it is helpful to introduce an economic equilibrium problem that turns out to be equivalent to the assignment problem.

Consider the possibility of matching the n objects with the n persons through a market mechanism, viewing each person as an economic agent acting in his/her

own best interest. Suppose that object j has a price p_j and that the person who receives the object must pay the price p_j . Then, the (net) value of object j for person i is $a_{ij} - p_j$ and each person i would logically want to be assigned to an object j_i with maximal value, that is, with

$$a_{ij_i} - p_{j_i} = \max_{j \in A(i)} \{a_{ij} - p_j\} \quad (1)$$

The economic system would then be at equilibrium, in the sense that no person would have an incentive to act unilaterally, seeking another object.

Equilibrium assignments and prices are naturally of great interest to economists, but there is also a fundamental relation with the assignment problem; it turns out that *an equilibrium assignment offers maximum total benefit (and, thus, solves the assignment problem), while the corresponding set of prices solves an associated dual problem*. This is a consequence of the duality theorem of linear programming (see, e.g., [27], [39], and [44]). In the terminology of linear programming, relation (1) is known as *complementary slackness* (CS). A simple, first-principles proof of the relation of equilibria to optimal assignments and dual optimization is also developed in Appendix 1.

2.1. The naive auction algorithm

Let us consider a natural process for finding an equilibrium assignment and price vector. We will call this process the *naive auction algorithm* because it has a serious flaw, as will be seen shortly. Nonetheless, this flaw will help motivate a more sophisticated and correct algorithm.

The naive auction algorithm proceeds in iterations and generates a sequence of price vectors and assignments. At the beginning of each iteration, the CS condition

$$a_{ij_i} - p_{j_i} = \max_{j \in A(i)} \{a_{ij} - p_j\} \quad (2)$$

is satisfied for all pairs (i, j_i) of the assignment. If all persons are assigned, the algorithm terminates. Otherwise, a nonempty subset I of persons that are unassigned is selected, and the following computations are performed.

Typical iteration of naive auction algorithm

Let I be a nonempty subset of persons that are unassigned.

- **Bidding phase:** Each person $i \in I$ finds an object j_i which offers maximal value, that is,

$$j_i \in \arg \max_{j \in A(i)} \{a_{ij} - p_j\} \quad (3)$$

and computes a bidding increment

$$\gamma_i = v_i - w_i \quad (4)$$

where v_i is the best object value

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

and w_i is the second-best object value

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

[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 .]

- **Assignment phase:** Each object j that is selected as best object by a nonempty subset $P(j)$ of persons in I , determines the highest bidder

$$i_j = \arg \max_{i \in P(j)} \gamma_i \quad (7)$$

raises its prices by the highest bidding increment $\max_{i \in P(j)} \gamma_i$ and gets assigned to the highest bidder i_j ; the person that was assigned to j at the beginning of the iteration (if any) becomes unassigned.


The algorithm continues with a sequence of iterations until all persons have an assigned object.

Note that γ_i cannot be negative since $v_i \geq w_i$ [compare (5) and (6)], so the object prices tend to increase. In fact, when i is the only bidder, γ_i is the largest bidding increment for which CS is maintained following the assignment of i to his/her preferred object. Just as in a real auction, bidding increments and price increases spur competition by making the bidder's own preferred object less attractive to other potential bidders.

Note also that there is some freedom in choosing the subset of persons I that bid during an iteration. One possibility is to let I consist of a single unassigned person. This version, known as the *Gauss-Seidel version* because of its similarity with Gauss-Seidel methods for solving systems of nonlinear equations, usually works best in a serial computing environment. The version where I consists of all unassigned persons, is the one best suited for parallel computation; it is known as the *Jacobi version* because of its similarity with Jacobi methods for solving systems of nonlinear equations.

3. ϵ -Complementary Slackness and the Auction Algorithm

Unfortunately, the naive auction algorithm does not always work (although it is an excellent initialization procedure for other methods that are based on price adjustment, e.g., primal-dual or relaxation). The difficulty is that the bidding increment γ_i is zero when more than one object offers maximum value for the



At Start of Iteration #	Object Prices	Assigned Pairs	Bidder	Preferred Object	Bidding Increment
1	0,0,0	(1, 1), (2, 2)	3	2	0
2	0,0,0	(1, 1), (3, 2)	2	2	0
3	0,0,0	(1, 1), (2, 2)	3	2	0

Figure 1. Illustration of how the naive auction algorithm may never terminate for a three-person and three-object problem. Here objects 1 and 2 offer benefit $C > 0$ to all persons, and object 3 offers benefit 0 to all persons. The algorithm cycles as persons 2 and 3 alternately bid for object 2 without changing its price because they prefer equally object 1 and object 2 ($\gamma_i = 0$).

bidder i . As a result, a situation may be created where several persons contest a smaller number of equally desirable objects without raising their prices, thereby creating a never-ending cycle; see Figure 1.

To break such cycles, we introduce a perturbation mechanism, motivated by real auctions where each bid for an object must raise the object’s price by a minimum positive increment, and bidders must, on occasion, take risks to win their preferred objects. In particular, let us fix a positive scalar ϵ , and say that an assignment and a price vector p satisfy ϵ -complementary slackness (ϵ -CS) if

$$a_{ij_i} - p_{j_i} \geq \max_{j \in A(i)} \{a_{ij} - p_j\} - \epsilon \tag{8}$$

for all assigned pairs (i, j_i) . In other words, to satisfy ϵ -CS, all assigned persons must be assigned to objects that are within ϵ of being best.

3.1. The auction algorithm

We now reformulate the previous auction process so that the bidding increment is always at least equal to ϵ . The resulting method, the *auction algorithm*, is the same as the naive auction algorithm, except that the bidding increment γ_i is

$$\gamma_i = v_i - w_i + \epsilon \quad (9)$$

[rather than $\gamma_i = v_i - w_i$ as in (4)]. With this choice, the ϵ -CS condition is satisfied. The particular increment $\gamma_i = v_i - w_i + \epsilon$ used in the auction algorithm is the maximum amount with this property. Smaller increments γ_i would also work as long as $\gamma_i \geq \epsilon$, but using the largest possible increment accelerates the algorithm. This is consistent with experience from real auctions, which tend to terminate faster when the bidding is aggressive.


It can be shown that this reformulated auction process terminates in a finite number of iterations, necessarily with a feasible assignment and a set of prices that satisfy ϵ -CS. To see this for the case of a fully dense problem (\mathcal{A} consists of all person-object pairs), note that if an object receives a bid in k iterations, its price must exceed its initial price by at least $k\epsilon$. Thus, for sufficiently large k , the object will become “expensive” enough to be judged “inferior” to some object that has not received a bid so far. It follows that an object can receive a bid in a limited number of iterations while some other object still has not yet received any bid. On the other hand, once all objects receive at least one bid, the auction terminates. Thus, the auction algorithm must terminate, and, in fact, the preceding argument shows that, for the case of zero initial prices, the total number of iterations in which an object receives a bid is no more than

$$\frac{\max_{(i,j)} a_{ij} - \min_{(i,j)} a_{ij}}{\epsilon}$$

If each iteration involves a bid by a single person, the total number of iterations is no more than n times the preceding quantity, and since each bid requires $O(n)$ operations, the running time of the algorithm is $O(n^2 \max_{(i,j)} |a_{ij}| / \epsilon)$.

The preceding termination proof can be generalized for the case of a sparse problem (one where the set of person-object pairs that can be assigned is limited), as long as the problem is feasible; see Appendix 2. Figure 2 shows how the auction algorithm, based on the bidding increment $\gamma_i = v_i - w_i + \epsilon$ [cf., (9)] overcomes the cycling problem of the example of Figure 1.

When the auction algorithm terminates, we have an assignment satisfying ϵ -CS, but is this assignment optimal? The answer here depends strongly on the size of ϵ . In a real auction, a prudent bidder would not place an excessively high bid for fear that he/she might win the object at an unnecessarily high price. Consistent with this intuition, we can show that if ϵ is small, then the final assignment will be “almost optimal.” In particular, the following proposition (proved in Appendix 1) shows that *the total cost of the final assignment is within $n\epsilon$ of being optimal.*



At Start of Iteration #	Object Prices	Assigned Pairs	Bidder	Preferred Object	Bidding Increment
1	0, 0, 0	(1, 1), (2, 2)	3	2	ϵ
2	0, ϵ , 0	(1, 1), (3, 2)	2	1	2ϵ
3	2ϵ , ϵ , 0	(2, 3), (3, 1)	1	2	2ϵ
4	2ϵ , 3ϵ , 0	(1, 2), (2, 1)	3	1	2ϵ
5	4ϵ , 3ϵ , 0	(1, 3), (3, 2)	2	2	2ϵ
6

Figure 2. Illustration of how the auction algorithm overcomes the cycling problem for the example of Figure 1 by making the bidding increment at least equal to ϵ . The table shows one possible sequence of bids and assignments generated by the auction algorithm, starting with all prices equal to 0. At each iteration except the last, the person assigned to object 3 bids for either object 1 or 2, increasing its price by ϵ in the first iteration and by 2ϵ in each subsequent iteration. In the last iteration, after the prices 1 and 2 rise to or above C , object 3 receives a bid and the auction terminates.

The idea is that when a feasible assignment and a set of prices satisfy ϵ -CS, they also satisfy CS (and are, therefore, primal and dual optimal, respectively) for a *slightly perturbed* problem where all costs a_{ij} are the same as before, except for the costs of the n assigned pairs, which are modified by an amount no more than ϵ .

PROPOSITION 1. *A feasible assignment satisfying ϵ - CS together with some price vector is within $n\epsilon$ of being optimal.*

Suppose now that the costs a_{ij} are all integer, which is the typical practical case (if a_{ij} are rational numbers, they can be scaled up to integer by multiplication with a suitable common number). Then, the total benefit of any assignment is integer, so if $n\epsilon < 1$, any complete assignment that is within $n\epsilon$ of being optimal must be optimal. It follows, that if $\epsilon < 1/n$, and the benefits a_{ij} are all integer, then the assignment obtained upon termination of the auction algorithm is optimal. We state this result as a proposition; the proof is given in Appendix 2.

PROPOSITION 2. *Consider a feasible assignment problem with integer benefits a_{ij} . If*

$$\epsilon < \frac{1}{n}$$

the auction algorithm terminates in a finite number of iterations with an optimal assignment.

Figure 3 shows the sequence of generated object prices for the example of Figures 1 and 2 in relation to the contours of the dual cost function of the assignment problem, which is given in Appendix 1. It can be seen from this figure that with each bid, the dual cost is approximately minimized (within ϵ) with respect to the price of the object receiving the bid. This observation can be established in generality; see [6] or [11]. Successive minimization of a cost function along single coordinates is a central feature of coordinate descent and relaxation methods, which are popular for unconstrained minimization of smooth functions and for solving systems of smooth equations. Thus, the auction algorithm can be interpreted as an approximate coordinate descent method.

3.2. ϵ -scaling

The amount of work needed for the auction algorithm to terminate can depend strongly on the value of ϵ and on the maximum absolute object benefit C given by

$$C = \max_{(i,j) \in \mathcal{A}} |a_{ij}| \quad (10)$$




Figure 3. A sequence of prices p_1 and p_2 generated by the auction algorithm for the example of Figures 1 and 2. The figure shows the equal dual cost surfaces in the space of p_1 and p_2 , with p_3 fixed at 0.

Basically, for many types of problems, the number of iterations up to termination tends to be proportional to C/ϵ as argued earlier for fully dense problems. This can also be seen from the example of Figure 3, where the number of iterations up to termination is roughly C/ϵ , starting from zero initial prices. For small ϵ , the method is 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.

Note also that there is a dependence on the initial prices; if these prices are near optimal, we expect that the number of iterations to solve the problem will be relatively small. This can be seen from the example of Figure 3; if the initial prices satisfy $p_1 \approx p_3 + C$ and $p_2 \approx p_3 + C$, the number of iterations up to termination is quite small.

The preceding observations suggest the idea of 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 critical value (for example, $1/n$, when the benefits a_{ij} are integer). Typical ϵ -reduction factors after each scaling phase are of the order of 4–10. Each application of the algorithm

provides good initial prices for the next application. ϵ -scaling was suggested in the original proposal of the auction algorithm ([7]), based on extensive experimentation, which established its effectiveness for many types of assignment problems. In particular, ϵ -scaling is typically beneficial for sparse problems. The cost structure of the problem is also important in determining whether ϵ -scaling is needed.

For integer data, it can be shown that the worst-case running time of the auction algorithm using scaling and appropriate data structures is $O(nA \log(nC))$, where A is the cardinality of the person-object pair set \mathcal{A} ; see [20] and [23]. Based on experiments, the running time of the algorithm for randomly generated problems seems to grow proportionally to something like $A \log n$ or $A \log n \log(nC)$. This is also supported by an approximate analysis in [49]. The practical computational aspects of the auction algorithm will be discussed in Section 10.

3.3. Dealing with infeasibility

Since termination can only occur with a feasible assignment, when the problem is infeasible, the auction algorithm will keep on iterating, as the user is wondering whether the problem is infeasible or just hard to solve. Thus, for problems where existence of a feasible assignment is not known a priori, one must supplement the auction algorithm with a mechanism to detect infeasibility. There are several such mechanisms, which we will now discuss.

A basic result about (symmetric and asymmetric) assignment problems is that given an infeasible assignment S , there are two mutually exclusive possibilities.

1. S has maximal cardinality, that is, there is no assignment having more assigned persons than S .
2. There exists some unassigned person i and some unassigned object j and an *augmenting path* with respect to S that starts at i and ends at j , that is, a sequence of the form

$$(i, j_1, i_1, \dots, j_m, i_m, j)$$

such that $j_1 \in A(i)$, $j \in A(i_m)$, $j_{k+1} \in A(i_k)$ for $k = 1, \dots, m-1$, and j_k is assigned to i_k under S for $k = 1, \dots, m$.

This result can be proved in a number of ways. For example, by introducing in the assignment problem graph a supersource node s connected to all the person nodes and a supersink node t connected to all the object nodes, we can view the problem of finding an assignment of maximal cardinality as the problem of finding a maximum flow from s to t . The result then follows by using the max-flow/min-cut theorem and the related analysis (eg., [11], Chapter 1, Propositions 2.3 and 2.4).

A corollary of the preceding result is that given a feasible assignment problem and an infeasible assignment S , for every person i that is unassigned under S , there exists an augmenting path with respect to S that starts at i . (To prove this, modify S by assigning every person that is unassigned under S , except i , to a fictitious object, and then apply the preceding result.)

One criterion that can be used to detect infeasibility is based on the maximum values

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

It turns out that in the course of the auction algorithm, all of these values will be bounded from below by a precomputable bound when the problem is feasible, but some of these values will be eventually reduced below this bound if the problem is infeasible. In particular, suppose that the auction algorithm is applied with initial object prices $\{p_j^0\}$. Then it can be shown that for any i , if person i is unassigned with respect to the current assignment S , and there is an augmenting path with respect to S that starts at i , we have

$$v_i \geq -(2n-1)C - (n-1)\epsilon - \max_j \{p_j^0\} \quad (11)$$

where $C = \max_{(i,j) \in A} |a_{ij}|$. The proof is obtained by adding the ϵ -CS condition along the augmenting path. If the problem is feasible, then as discussed earlier, there exists an augmenting path starting at each unassigned person at all times, so the lower bound (11) on v_i will hold for all unassigned persons i throughout the auction algorithm. On the other hand, if the problem is infeasible, some persons i will be submitting bids infinitely often, and the corresponding values v_i will be decreasing towards $-\infty$. Thus, we can apply the auction algorithm and keep track of the values v_i as they decrease. Once some v_i gets below its lower bound, we know that the problem is infeasible.

Unfortunately, it may take many iterations for some v_i to reach its lower bound. An alternative method to detect infeasibility is to convert the problem to a feasible problem by adding a set of artificial pairs $\bar{\mathcal{A}}$ to the original set \mathcal{A} . The benefits of these pairs should be very small, so that none of them participates in an optimal assignment unless the problem is infeasible. In particular, it can be shown that if the original problem was feasible, no pair $(i, j) \in \bar{\mathcal{A}}$ will participate in the optimal assignment, provided that

$$a_{ij} < -(2n-1)C, \quad \forall (i, j) \in \bar{\mathcal{A}} \quad (12)$$

where $C = \max_{(i,j) \in A} |a_{ij}|$. To prove this by contradiction, assume that by adding to the set \mathcal{A} the set of artificial pairs $\bar{\mathcal{A}}$, we create an optimal assignment S^* that contains a nonempty subset \bar{S} of artificial pairs. Then, for every assignment S consisting exclusively of pairs from the original set \mathcal{A} we must have

$$\sum_{(i,j) \in \bar{S}} a_{ij} + \sum_{(i,j) \in S^* - \bar{S}} a_{ij} \geq \sum_{(i,j) \in S} a_{ij}$$

from which

$$\sum_{(i,j) \in \bar{S}} a_{ij} \geq \sum_{(i,j) \in S} a_{ij} - \sum_{(i,j) \in S^* - \bar{S}} a_{ij} \geq -(2n-1)C$$

This contradicts (12). Note that if $a_{ij} \geq 0$ for all $(i, j) \in \mathcal{A}$, the preceding argument can be modified to show that it is sufficient to have $a_{ij} < -(n-1)C$ for all artificial pairs (i, j) .

On the other hand, the addition of artificial pairs with benefit $-(2n-1)C$, as per (12), expands the cost range of the problem by a factor of $(2n-1)$. In the context of ϵ -scaling, this necessitates a much larger starting value for ϵ and a correspondingly large number of ϵ -scaling phases. If the problem is feasible, these extra scaling phases are wasted. Thus, for problems that are normally expected to be feasible, it may be better to introduce artificial pairs with benefits that are of the order of $-C$ and then gradually scale downward these benefits towards the $-(2n-1)C$ threshold if artificial pairs persist in the assignments obtained by the auction algorithm. This procedure of scaling downward the benefits of the artificial pairs can be embedded in a number of ways within the ϵ -scaling procedure.

Still another method to detect infeasibility is based on the following property: Even when the problem is infeasible, the auction algorithm will find an assignment of maximal cardinality in a finite number of iterations, provided that unassigned persons are selected for bidding in an order that is either cyclic, or else ensures that each person will get a chance to submit a bid at least once within some fixed number of iterations. The proof of this is based on the lower bound (11) and the property of assignments of maximal cardinality stated earlier. In particular, if the current assignment never reached maximal cardinality, there would always exist an unassigned person i , and a path that starts at i and is augmenting with respect to the current assignment. By adding the ϵ -CS condition along this path, we see that v_i will always satisfy the lower bound (11), which is a contradiction because v_i will tend to $-\infty$ for all i that submit a bid infinitely often.

Now suppose that we periodically interrupt the auction algorithm and check whether there exists an augmenting path from some unassigned person to some unassigned object (a simple search of the breadth-first type can be used for this; see, e.g., [11], p. 26). Then, once the cardinality of the current assignment becomes maximal but is less than n , this check will establish that the problem is infeasible. Thus, this modified auction algorithm is guaranteed to either find a feasible assignment and a set of prices satisfying ϵ -CS, or to establish that the problem is infeasible and simultaneously obtain an assignment of maximal cardinality. In the latter case, it can be shown that the original (symmetric) problem will separate (through a saturated cut) into two components corresponding to (asymmetric) assignment problems. One may then use auction algorithms for asymmetric problems (see Section 9) to optimize the assignment within each component and obtain an optimal assignment within the class of all assignments with maximal cardinality.

3.4. Profits and reverse auction

Since the assignment problem is symmetric, it is possible to exchange the roles of persons and objects. This leads to the idea of *reverse auction*, where the objects compete for persons by essentially offering discounts (lowering their prices). Roughly, given a price vector p , we can view the net value of the best object for person i

$$\max_{(i,j) \in A(i)} \{a_{ij} - p_j\} \quad (13)$$

as a profit for person i . When objects lower their prices they tend to increase the profits of the persons, so that profits play for persons a role analogous to the role prices play for objects. We can thus describe reverse auction in two different ways; one where unassigned objects lower their prices as much as possible to attract a person without violating ϵ -CS, and another where unassigned objects select a best person and raise his/her profit as much as possible without violating ϵ -CS. The second description is analytically more convenient, since, with this description, forward and reverse auctions will be seen to be mathematically equivalent.

Let us introduce a profit variable π_i for each person i , and consider the following ϵ -CS condition for an assignment S and a profit vector π

$$a_{ij} - \pi_i \geq \max_{k \in B(j)} \{a_{kj} - \pi_k\} - \epsilon, \quad \forall (i, j) \in S \quad (14)$$

where $B(j) = \{i | (i, j) \in \mathcal{A}\}$ is the set of persons that can be assigned to object j (assumed nonempty). The relation between the profit variable π_i and the expression $\max_{(i,j) \in A(i)} \{a_{ij} - p_j\}$ [cf. (13)] will become apparent later when we discuss a somewhat different ϵ -CS condition, which involves both prices and profits; see (17a), (17b), and (20). Note the symmetry of the ϵ -CS condition (14) for profits with the corresponding one for prices; cf. (8).

The reverse auction algorithm maintains at the beginning of each iteration an assignment S and a profit vector π satisfying the ϵ -CS condition (14). It terminates when the assignment is feasible.

Typical iteration of reverse auction

Let J be a nonempty subset of objects j that are unassigned.

- **Bidding phase:** For each object $j \in J$, find a best person i_j such that

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

and the corresponding value

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

and find

$$\omega_j = \max_{i \in B(j), i \neq i_j} \{a_{ij} - \pi_i\} \quad (15)$$

[If i_j is the only person in $B(j)$, we define ω_j to be $-\infty$ or, for computational purposes, a number that is much smaller than β_j .]

- **Assignment phase:** Each person i that is selected as best person by a nonempty subset $P(i)$ of objects in J , determines the highest bidder

$$j_i = \arg \max_{j \in P(i)} \{\beta_j - \omega_j + \epsilon\} \quad (16)$$

increases π_i by the highest bidding increment $\max_{j \in P(i)} \{\beta_j - \omega_j + \epsilon\}$ and gets assigned to the highest bidder j_i ; the object that was assigned to i at the beginning of the iteration (if any) becomes unassigned.

Note that reverse auction is identical to (forward) auction with the roles of persons and objects as well as profits and prices interchanged. Thus, by using the corresponding (forward) auction results (cf. Propositions 1 and 2), we have:

PROPOSITION 3. *If at least one feasible assignment exists, the reverse auction algorithm terminates in a finite number of iterations. The feasible assignment obtained upon termination is within $n\epsilon$ of being optimal (and is optimal if the problem data are integer and $\epsilon < 1/n$).*

3.5. Combined forward and reverse auction

One of the reasons we are interested in reverse auction is to construct algorithms that switch from forward to reverse auction and back. Such algorithms must simultaneously maintain a price vector p satisfying the ϵ -CS condition (8) and a profit vector π satisfying the ϵ -CS condition (14). To this end we introduce an ϵ -CS condition for the pair (π, p) , which as we will see, implies the other two. Maintaining this condition is essential for switching gracefully between forward and reverse auction.

We say that an assignment S and a pair (π, p) satisfy ϵ -CS if

$$\pi_i + p_j \geq a_{ij} - \epsilon, \quad \forall (i, j) \in \mathcal{A} \quad (17a)$$

$$\pi_i + p_j = a_{ij}, \quad \forall (i, j) \in S \quad (17b)$$

We have the following proposition.

PROPOSITION 4. *Suppose that an assignment S together with a profit-price pair (π, p) satisfies ϵ -CS. Then:*

- (a) S and π satisfy the ϵ -CS condition

$$a_{ij} - \pi_i \geq \max_{k \in B(j)} \{a_{kj} - \pi_k\} - \epsilon, \quad \forall (i, j) \in S \quad (18)$$

(b) S and p satisfy the ϵ -CS condition

$$a_{ij} - p_j \geq \max_{k \in A(i)} \{a_{ik} - p_k\} - \epsilon, \quad \forall (i, j) \in S \quad (19)$$

(c) If S is feasible, then S is within $n\epsilon$ of being an optimal assignment.

Proof.

- (a) In view of (17b), for all $(i, j) \in S$, we have $p_j = a_{ij} - \pi_i$, so (17a) implies that $a_{ij} - \pi_i \geq a_{kj} - \pi_k - \epsilon$ for all $k \in B(j)$. This shows (18).
- (b) The proof is the same as the one of part (a) with the roles of π and p interchanged.
- (c) Since by part (b), the ϵ -CS condition (19) is satisfied, by Proposition 1, S is within $n\epsilon$ of being optimal. **Q.E.D.**

We now introduce a combined forward/reverse algorithm. The algorithm maintains an assignment S and a profit-price pair (π, p) satisfying the ϵ -CS conditions (17). It terminates when the assignment is feasible. A common way to initialize the algorithm so that the ϵ -CS conditions are satisfied is to take S to be empty, to choose p arbitrarily, and to select π as a function of p via the relation $\pi_i = \max_{k \in A(i)} \{a_{ik} - p_k\}$ for each person i .

Combined forward reverse auction algorithm

- **Step 1 (run forward auction):** Execute several iterations of the forward auction algorithm (subject to the termination condition), and at the end of each iteration (after increasing the prices of the objects that received a bid), set

$$\pi_i = a_{i,j_i} - p_{j_i} \quad (20)$$

for every person-object pair (i, j_i) that entered the assignment during the iteration. Go to step 2.

- **Step 2 (run reverse auction):** Execute several iterations of the reverse auction algorithm (subject to the termination condition), and at the end of each iteration (after increasing the profits of the persons that received a bid), set

$$p_j = a_{i_j,j} - \pi_{i_j}, \quad (21)$$

for every person-object pair (i_j, j) that entered the assignment during the iteration. Go to step 1.

Note that the additional overhead of the combined algorithm over the forward or the reverse algorithm is minimal; just one update of the form (20) or (21) is

required per iteration for each object or person that received a bid during the iteration. An important property is that the updates of (20) and (21) maintain the ϵ -CS conditions (17) for the pair (π, p) and, therefore, by Proposition 4, maintain the required ϵ -CS conditions (18) and (19) for π and p , respectively. This is stated in the following proposition, which was proved in [18] (see also [11]).

PROPOSITION 5. *If the assignment and profit-price pair available at the start of an iteration of either the forward or the reverse auction algorithm satisfy the ϵ -CS conditions (17), the same is true for the assignment and profit-price pair obtained at the end of the iteration, provided (20) is used to update π (in the case of forward auction), and (21) is used to update p (in the case of reverse auction).*

Note that during forward auction, the object prices p_j increase, while the profits π_i decrease, but exactly the opposite happens in reverse auction. For this reason, the termination proof used for forward auction (see Appendix 2) does not apply to the combined method. Indeed, it is possible to construct examples of feasible problems where the combined method never terminates if the switch between forward and reverse auctions is done arbitrarily. However, it is easy to guarantee that the combined algorithm terminates finitely for a feasible problem; it is sufficient to ensure that some “irreversible progress” is made before switching between forward and reverse auction. One easily implementable possibility is to refrain from switching until at least one more person-object pair has been added to the assignment. In this way there can be a switch at most $(n - 1)$ times between the forward and reverse steps of the algorithm. For a feasible problem, forward and reverse auction by themselves have guaranteed finite termination, so the final step will terminate with a feasible assignment satisfying ϵ -CS.

The combined forward/reverse auction algorithm typically works substantially (and often dramatically) faster than the forward version, as shown experimentally in the original paper [18] and in the extensive computational study by D. Castañon [24]. It seems to be affected less by the price war phenomenon that motivated ϵ -scaling. Price wars can still occur in the combined algorithm, but they arise through more complex and unlikely problem structures than in the forward algorithm. For this reason the combined forward/reverse auction algorithm depends less on ϵ -scaling for good performance than its forward counterpart. One consequence of this is that starting with $\epsilon = 1/(n + 1)$ and bypassing ϵ -scaling is often the best choice. Another consequence is that a larger ϵ -reduction factor can typically be used with no price war effects in ϵ -scaled forward/reverse auction than in ϵ -scaled forward auction. As a result, fewer ϵ -scaling phases are typically needed in forward/reverse auction to deal effectively with price wars.

4. Auction Algorithms for Shortest Path Problems

We now turn our attention to other types of network flow problems. Our approach for constructing auction algorithms for such problems is to convert them to assignment problems, and then to suitably apply the auction algorithm and streamline the computations. We start with the classical shortest path problem.

Suppose that we are given a graph with node set \mathcal{N} , arc set \mathcal{A} , and a length a_{ij} for each arc (i, j) . In this section, by a *path* we mean a sequence of nodes (i_1, i_2, \dots, i_k) such that (i_m, i_{m+1}) is an arc for all $m = 1, \dots, k-1$. If, in addition, the nodes i_1, i_2, \dots, i_k are distinct, the sequence (i_1, i_2, \dots, i_k) is called a *simple path*. The length of a path is defined to be the sum of its arc lengths. Assuming that all cycles have positive length, we want to find a path of minimum length over all paths that start at a given origin (node 1) and end at a given destination (node t).

There is a well-known transformation that converts this problem to a particular type of assignment problem as shown in Figure 4. We can apply the auction algorithm to solve the equivalent problem, but it turns out that the structure of this problem is such that the *naive* auction algorithm ($\epsilon = 0$) works. Assuming that all arc lengths are nonnegative, we can start the naive auction algorithm with the zero price vector and with the assignment that assigns, for $i \neq 1, t$, each person i' to object i (refer to Figure 4); this assignment-price pair satisfies CS [since all arc lengths are nonnegative and the assigned arcs (i', i) have zero cost], but is not feasible because it leaves person $1'$ and object t unassigned. It can be seen by using induction or by tracing the steps of the naive auction algorithm that each assignment generated consists of a (possibly empty) sequence of the form

$$(1', i_1), (i'_1, i_2), \dots, (i'_{k-1}, i_k)$$

together with the additional arcs

$$(i', i), \quad \text{for } i \neq i_1, \dots, i_k, t$$

This sequence corresponds to a path $P = (1, i_1, \dots, i_k)$. As long as $i_k \neq t$, the (unique) unassigned person in the naive auction algorithm is person i'_k , corresponding to the terminal node of the path. If $i_k = t$ a feasible assignment results, in which case the naive auction algorithm terminates. Otherwise, the unassigned person i'_k submits a bid and there are two possibilities.

1. i_k is the unique best object, in which case i'_{k-1} becomes unassigned and the path $(1, i_1, \dots, i_{k-1})$ corresponding to the new assignment is obtained from the previous path $(1, i_1, \dots, i_k)$ via contraction by one node.
2. There is an object $i_{k+1} \neq i_k$ which is best, in which case the path $(1, i_1, \dots, i_k, i_{k+1})$ corresponding to the new assignment is obtained from the previous path $(1, i_1, \dots, i_k)$ via extension by one node.




Figure 4. A shortest path problem (the origin is 1, the destination is $t = 4$) and its corresponding assignment problem. The arc lengths and the assignment costs are shown next to the arcs. A shortest path can be associated with an optimal augmenting path that starts at $1'$ and ends at $t = 4$. Generally, for each node $i \neq 1$ we introduce an object node i , and for each node $i \neq t$ we introduce a person node i' . For every arc (i, j) of the shortest path problem with $i \neq t$ and $j \neq 1$, we introduce the arc (i', j) with cost a_{ij} in the assignment problem. We also introduce the zero cost arc (i', i) for each $i \neq 1, t$. Given the assignment that assigns object i to person i' for $i \neq 1, t$, but leaves person $1'$ and object t unassigned, paths from 1 to t can be associated with augmenting paths that start at $1'$ and end at t . A shortest path from 1 to t corresponds to a shortest augmenting path from the unassigned person $1'$ to the unassigned object t . It can be shown that an augmentation along such a shortest augmenting path gives an optimal assignment.

We will now describe the naive auction algorithm directly in terms of the original shortest path problem, properly translating (a) the preceding operations of one-node contraction or extension of a path, and (b) the use of prices and the associated price changes of the terminal node of the path, while maintaining a CS condition.

4.1. The auction/shortest path algorithm

The auction algorithm for shortest paths maintains at all times a simple path $P = (1, i_1, i_2, \dots, i_k)$. If i_{k+1} is a node that does not belong to a path $P = (1, i_1, i_2, \dots, i_k)$ and (i_k, i_{k+1}) is an arc, *extending* P by i_{k+1} means replacing P by the path $P = (1, i_1, i_2, \dots, i_k, i_{k+1})$ called the *extension of* P by i_{k+1} . If P does not consist of just the origin node 1, *contracting* P means replacing P by the path $P = (1, i_1, i_2, \dots, i_{k-1})$.

The algorithm maintains also a price vector p satisfying together with P the following property

$$p_i \leq a_{ij} + p_j, \quad \forall (i, j) \in \mathcal{A} \quad (22a)$$

$$p_i = a_{ij} + p_j, \text{ for all pairs of successive nodes } i \text{ and } j \text{ of } P \quad (22b)$$

which is referred to as *complementary slackness* (CS). This condition can be related to the CS condition for the equivalent assignment problem as well as to CS conditions for a formulation of the shortest path problem as a minimum cost flow problem (see [11], Section 1.3).

It can be shown that if a pair (P, p) satisfies the CS conditions, then the portion of P between node 1 and any node $i \in P$ is a shortest path from 1 to i , while $p_1 - p_i$ is the corresponding shortest distance. To see this, note that by (22b), $p_i - p_k$ is the length of the portion of P between i and k , and every path connecting i to k must have length at least equal to $p_i - p_k$ [add (22a) along the arcs of the path]. We assume that an initial simple path P and price vector p satisfying CS are available. This is not a restrictive assumption when all arc lengths are nonnegative, since then one can use the default choices

$$P = (1), \quad p_i = 0, \quad \forall i$$

The algorithm proceeds in iterations, transforming a pair (P, p) satisfying CS into another pair satisfying CS. At each iteration, the path P is either extended by a new node or else contracted by deleting its terminal node. In the latter case, the price of the terminal node is increased strictly. A degenerate case occurs when the path consists of just the origin node 1; in this case the path is either extended or is left unchanged with the price p_1 being strictly increased. The iteration is as follows.

Typical iteration of the auction/shortest path algorithm

Let i be the terminal node of P . If

$$p_i < \min_{(i,j) \in \mathcal{A}} \{a_{ij} + p_j\}$$

go to step 1; else go to step 2.

• Step 1 (contract path): Set

$$p_i := \min_{(i,j) \in \mathcal{A}} \{a_{ij} + p_j\}$$

and if $i \neq 1$ contract P . Go to the next iteration.

• Step 2 (extend path): Extend P by node j_i where

$$j_i = \arg \min_{(i,j) \in \mathcal{A}} \{a_{ij} + p_j\}$$

If j_i is the destination t , stop; P is the desired shortest path. Otherwise, go to the next iteration.

Note that following an extension (step 2), P is a simple path from 1 to j_i ; if this were not so, then adding j_i to P would create a cycle, and for every arc (i, j) of this cycle we would have $p_i = a_{ij} + p_j$. By adding this condition along the cycle, we see that the cycle should have zero length, which is not possible by our assumptions.

There is an interesting interpretation of the CS conditions in terms of a mechanical model [42]. Think of each node as a ball, and for every arc $(i, j) \in \mathcal{A}$, connect i and j with a string of length a_{ij} . (This requires that $a_{ij} = a_{ji} > 0$, which we assume for the sake of the interpretation.) Let the resulting balls-and-strings model be at an arbitrary position in three-dimensional space, and let p_i be the vertical coordinate of node i . Then the CS condition $p_i - p_j \leq a_{ij}$ clearly holds for all arcs (i, j) , as illustrated in Figure 5(b). If the model is picked up and left to hang from the origin node (by gravity—strings that are tight are perfectly vertical), then for all the tight strings (i, j) we have $p_i - p_j = a_{ij}$, so any tight chain of strings corresponds to a shortest path between the end nodes of the chain, as illustrated in Figure 5(c). In particular, the length of the tight chain connecting the origin node 1 to any other node i is $p_1 - p_i$ and is also equal to the shortest distance from 1 to i . (This result is essentially the well-known min-path/max-tension theorem; see e.g., [11] and [48].)


The auction/shortest path algorithm can also be interpreted in terms of the balls-and-strings model; it can be viewed as a process whereby nodes are raised in stages as illustrated in Figure 6. Initially all nodes are resting on a flat surface. At each stage, we raise the last node in a tight chain that starts at the origin to the level at which at least one more string becomes tight.

When there is a single origin and multiple destinations, the algorithm can be applied with virtually no change. We simply stop the algorithm when all destinations have become the terminal node of the path P at least once.

4.2. The reverse algorithm

There are a number of ways to speed up the basic algorithm, which are described in detail in [5], [11], and [22]; see also Section 10. The most significant of these relates to a two-sided version of the algorithm that maintains, in addition to the path P , another path R that *ends at the destination*. To understand this version, we first note that in shortest path problems, one can exchange the role of origins and destinations by reversing the direction of all arcs. It is, therefore, possible to use a destination-oriented version of our algorithm that maintains a path R that *ends* at the destination and changes at each iteration by means of a contraction or an extension. This algorithm, called the *reverse algorithm*, is mathematically equivalent to the earlier forward algorithm, and parallels the reverse auction algorithm for the assignment problem discussed in the previous section.

Initially, in the reverse algorithm, R is any path ending at the destination,



Shortest path problem with arc lengths shown next to the arcs. Node 1 is the origin. Node 4 is the destination.

Figure 5. Illustration of the CS conditions for the shortest path problem. If each node is a ball, and for every arc $(i, j) \in \mathcal{A}$ nodes i and j are connected with a string of length a_{ij} , the vertical coordinates p_i of the nodes satisfy $p_i - p_j \leq a_{ij}$ as shown in (b) for the problem given in (a). If the model is picked up and left to hang from the origin node 1, then $p_1 - p_i$ gives the shortest distance to each node i , as shown in (c).

and p is any price vector satisfying the CS conditions (22) together with R ; for example

$$R = (t), \quad p_i = 0, \quad \forall i$$

if all arc lengths are nonnegative.

Typical iteration of the reverse algorithm

Let j be the starting node of R . If


$$p_j > \max_{(i,j) \in \mathcal{A}} \{p_i - a_{ij}\}$$

go to step 1; else go to step 2.

• **Step 1(contract path):** Set

$$p_j := \max_{(i,j) \in \mathcal{A}} \{p_i - a_{ij}\}$$

and if $j \neq t$ contract R , (that is, delete the starting node j of R). Go to the next iteration.



Shortest path problem with arc lengths shown next to the arcs. Node 1 is the origin. Node 4 is the destination.




Figure 6. Illustration of the auction/shortest path algorithm in terms of the balls-and-strings model for the problem shown in (a). The model initially rests on a flat surface, and various balls are then raised in stages. At each stage we raise a single ball $i \neq t$ (marked by gray), which is at lower level than the origin 1 and can be reached from 1 through a sequence of tight strings; i should not have any tight string connecting it to another ball at a lower level, that is, i should be the last ball in a tight chain hanging from 1. (If 1 does not have any tight string connecting it to another ball at a lower level, we use $i=1$.) We then raise i to the first level at which one of the strings connecting it to a ball at a lower level becomes tight. Each stage corresponds to a contraction. The ball i , which is being raised, corresponds to the terminal node of the path.

- **Step 2 (extend path):** Extend R by node i_j (that is, make i_j the starting node of R , preceding j), where

$$i_j = \arg \max_{(i,j) \in \mathcal{A}} \{p_i - a_{ij}\}$$

If i_j is the origin 1, stop; R is the desired shortest path. Otherwise, go to the next iteration.

The reverse algorithm is most helpful when it is combined with the forward algorithm. In a combined algorithm, initially we have a price vector p , and two simple paths P and R , satisfying CS together with p , where P starts at the origin and R ends at the destination. The paths P and R are extended and contracted according to the rules of the forward and the reverse algorithms, respectively, and the combined algorithm terminates when P and R have a common node. Both P and R satisfy CS together with p throughout the algorithm, so when P and R meet, say at node i , the composite path consisting of the portion of P from 1 to i followed by the portion of R from i to t will be shortest.

Combined forward reverse auction/shortest path algorithm

- **Step 1 (run forward algorithm):** Execute several iterations of the forward algorithm (subject to the termination condition), at least one of which leads to an increase of the origin price p_1 . Go to step 2.
- **Step 2 (run reverse algorithm):** Execute several iterations of the reverse algorithm (subject to the termination condition), at least one of which leads to a decrease of the destination price p_t . Go to step 1.

The combined forward/reverse algorithm can also be interpreted in terms of the balls-and-strings model of Figure 5. Again, all nodes are resting initially on a flat surface. When the forward part of the algorithm is used, we raise nodes in stages as illustrated in Figure 6. When the reverse part of the algorithm is used, we lower nodes in stages; at each stage, we lower the *top* node in a tight chain that ends at the destination to the level at which at least one more string becomes tight.

Note that the case of multiple destinations can be handled by using a separate reverse path for each destination. One then alternates between a forward step and a reverse step as in the preceding algorithm while taking up cyclically different destinations in different reverse steps.

Based on experiments with randomly generated problems on a serial machine ([5]), the combined forward/reverse auction/shortest path algorithm substantially outperforms its closest competitors for problems with few destinations and a single origin (the computation time is reduced often by an order of magnitude or more). For the case of many (or all) destinations, the algorithm apparently runs slower than state-of-the-art label-setting and label-correcting methods (typical slowdown factors are of the order of two or three). On the other hand, for

the case of multiple destinations, the algorithm is better suited for parallel computation than other shortest path algorithms. Also there is a variation of the algorithm that has substantially improved performance for problems with many destinations. This variation is described next.

4.3. *The auction algorithm with graph reduction*

Despite its excellent practical performance for problems with few destinations, the auction algorithm has pseudopolynomial complexity; for an example see [5] and [11]. Weakly polynomial versions of the algorithm were developed in [5] using the idea of scaling the arc lengths, but these versions did not prove effective in practice. Strongly polynomial versions of the algorithm were obtained by Pallottino and Scutella' ([43]) by adding to the extension and contraction operations a *reduction* operation. Here, each time a node becomes the terminal node of the path for the first time, all its incoming arcs except the one of the path are deleted, since they cannot be used to improve the distance to the node. The auction algorithm thus obtained can be shown to have an $O(m^2)$ running time, where m is the number of arcs. By using the idea of presorting the outgoing arcs of each node in order of increasing length, the running time was reduced further to $O(mn)$, where n is the number of nodes.

An additional advantage of graph reduction is that it allows the relaxation of the positivity assumption on all cycle lengths to nonnegativity. The reason for requiring positive cycle lengths was to ensure that no cycle could be formed through the process of path extension. On the other hand, with graph reduction, every node already visited by the path has a unique (undeleted) incoming arc except for node 1, which has no incoming arc at all. With a little thought, it can be seen that this precludes the extension of the path to a node that is already on the path.

In subsequent work by Bertsekas, Pallottino, and Scutella' ([22]), the graph reduction idea was strengthened by using certain upper bounds to the node shortest distances to delete arcs more effectively. Two algorithms were developed. The first maintains the basic simplicity of the auction algorithm given earlier, and has $O(n \min\{m, n \log n\})$ running time. The second algorithm is somewhat more complex, but has an $O(n^2)$ running time. These theoretical improvements, in conjunction with efficient implementation techniques, have resulted in substantially faster practical performance for single origin/all destination problems. In particular, for fully dense randomly generated problems, the auction algorithm with graph reduction has outperformed its closest competitors ([22]).

4.4. k shortest path problems

Consider the generalization of the single-origin/single-destination shortest path problem, where, instead of a single path, we seek the k best paths. In other words, we want to find the shortest path, say P_1 , between an origin s and a destination t , then find the shortest path P_2 between s and t subject to the condition that it be different by at least one arc from P_1 , then find the shortest path P_3 subject to the condition that it be different than P_1 and P_2 , and so on, up to the k th shortest path P_k . There are two versions of this problem depending on whether the paths P_i are allowed to contain cycles or not ([37]).

Important algorithmic ideas for solving k shortest path problems are based on solving a sequence of ordinary shortest path problems involving graphs that differ slightly from each other by a few arcs and/or nodes ([3], [4], [28], [37], and [40]). The auction/shortest path algorithm is particularly well suited for solving these shortest path problems for two reasons. First, it can conveniently transfer favorable initial price information from one shortest path problem to the next, so that the solution of each subproblem after the first is greatly expedited. Second, it requires the solution of single-origin/single-destination problems, which it can solve (in forward/reverse mode) much faster than the single-origin/all-destination problem that some of the existing methods must solve ([3] and [4]). At present, there are no computational studies exploring the use of auction algorithms for k shortest path problems. This is an interesting subject for further investigation.

4.5. k node-disjoint shortest path problems

Let us consider another generalization of the single-origin/single-destination shortest path problem, where instead of a single-path, we seek k node-disjoint paths with minimum sum of lengths. Once a supersource node s and a supersink node t are added, the $n \times n$ assignment problem can be viewed as a special case of this problem, where $k = n$. Another special case is the separable version of the three-dimensional (asymmetric) assignment problem, involving the optimal choice of k disjoint ordered triplets, where the cost of a triplet (i, j, m) is of the separable form $a_{ij} + a_{jm}$.

An auction algorithm for the k node-disjoint shortest path problem that bears similarity to the preceding shortest path algorithm has been developed in [16]. In particular, the algorithm maintains a price vector p , together with k node-disjoint paths starting at the origin, that satisfy an ϵ -CS condition. If all the paths end at the destination, the algorithm stops. Otherwise one of the paths, say P , that does not end at the destination is either extended or contracted at its terminal node. If, as a result of an extension, P joins another path P' at a node j , P' is truncated to the portion up to the node preceding j ; the portion following j is added to P . Favorable computational results for this auction algorithm are given in [16]. Interestingly, when this algorithm is specialized to the $n \times n$

assignment problem, it becomes identical to the original auction algorithm of Section 4, thereby illustrating vividly the relation of the bidding mechanism of the assignment auction algorithm and the extension/contraction mechanism of the shortest path auction algorithm.

5. Extension to Transportation Problems

We now consider the extension of the auction algorithm to the uncapacitated transportation problem. Here we have a bipartite graph with m sources and n sinks. The problem is the same as the assignment problem except that the node supplies need not be 1 or -1 . It has the form

$$\begin{aligned}
 &\text{maximize } \sum_{(i,j) \in \mathcal{A}} a_{ij}x_{ij} \\
 &\text{subject to} \\
 &\quad \sum_{j \in \mathcal{A}(i)} x_{ij} = \alpha_i, \quad \forall i = 1, \dots, m \\
 &\quad \sum_{i \in \mathcal{B}(j)} x_{ij} = \beta_j, \quad \forall j = 1, \dots, n \\
 &\quad 0 \leq x_{ij}, \quad \forall (i, j) \in \mathcal{A}
 \end{aligned} \tag{23}$$

Here α_i and β_j are positive integers, which for feasibility must satisfy

$$\sum_{i=1}^m \alpha_i = \sum_{j=1}^n \beta_j$$

We can easily convert this problem to an assignment problem by replacing each source (or sink) node with a collection of “duplicate” persons (or objects, respectively). In particular, a source node i with supply α_i is replaced by α_i persons, and a sink node j with demand β_j is replaced by β_j objects. Furthermore, for each arc (i, j) we must create an arc of benefit a_{ij} connecting each person corresponding to i with each object corresponding to j . An example is given in Figure 7.

It is possible to solve the equivalent assignment problem by straightforward application of the auction algorithm but there are two difficulties.

1. The dimension of the problem is greatly increased (both the number of persons and the number of objects become $\sum_{i=1}^m \alpha_i$).
2. The structure of the equivalent assignment problem is such that protracted price wars are inevitable due to the duplicate objects and persons created by the transformation. Figure 8 provides an example.

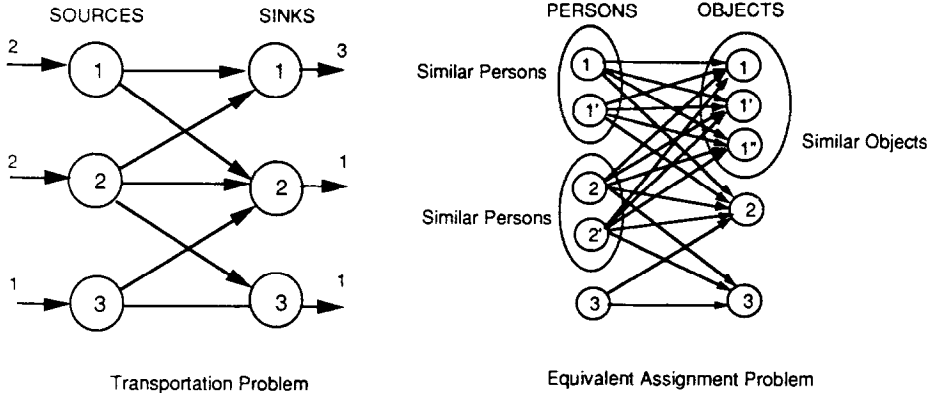


Figure 7. Illustration of the conversion of a transportation problem into an assignment problem. Each source i is transformed into α_i persons (for example source 1 is transformed into persons 1 and 1'), and each sink j is transformed into β_j objects (for example sink 1 is transformed into objects 1, 1', and 1''). If (i, j) is an arc of benefit a_{ij} in the transportation problem, there is an assignment arc of benefit a_{ij} connecting each of the α_i persons corresponding to i with each of the β_j objects corresponding to j .

It is possible to address both of these difficulties by modifying the auction algorithm so that it takes advantage of the special structure of assignment problems created by duplicate objects and persons ([14]). In particular, given an assignment problem, we say that two objects j and j' are *similar*, and write $j \sim j'$, if they can be matched with the same persons and at equal values, that is


$$B(j) = B(j'), \quad \text{and} \quad a_{ij} = a_{ij'} \quad \forall i \in B(j)$$

We say that two persons i and i' are similar, and write $i \sim i'$ if

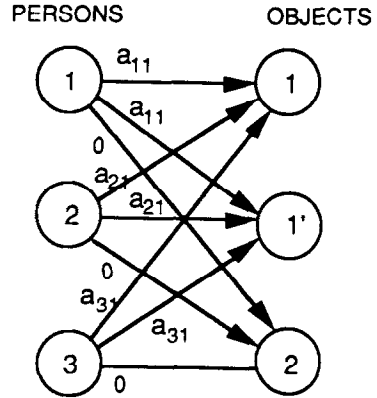
$$A(i) = A(i'), \quad \text{and} \quad a_{ij} = a_{i'j} \quad \forall i \in A(i)$$

For each object j , the set of all objects similar to j is called the *similarity class* of j and is denoted $M(j)$. For each person i , the set of all persons similar to i is called the similarity class of i and is denoted $M(i)$.

To get a sense of how one can deal with similarity classes, consider the auction algorithm of Section 3 applied to the assignment problem of Figure 8. Here, person i ($i = 1, 2, 3$) will continue bidding for the similar objects 1 and 1' until both prices p_1 and p'_1 become at least equal to a_{i1} , and object 2 becomes no less attractive for person i than objects 1 and 1'. Thus a_{i1} may be viewed as a *contention threshold* to be reached by the prices of *all* the similar objects before person i becomes interested in an object outside their similarity class. The contention threshold of an object in some similarity class can, of course, be much



Transportation Problem



Equivalent Assignment Problem

Figure 8. Typical example where a transportation problem is converted into an equivalent assignment problem and the duplicate objects induce price wars in the auction algorithm. Here a_{11} , a_{21} , and a_{31} are much larger than zero. The persons 1, 2, and 3 will keep on bidding up the prices of objects 1 and 1' by 2ϵ increments until one of the prices p_1 or $p_{1'}$ reaches or exceeds the minimum of a_{11} , a_{21} , and a_{31} . (For $a_{11} = a_{21} = a_{31} = C$, this example is the same as the one of Figures 1-3.)

higher than the price of the object (which, by the rules of the auction algorithm, is constrained to be within ϵ of the price of any other object of the same class). Suppose now that when a person bids for an object of a similarity class he/she not only updates its price, but also updates the corresponding contention threshold. Suppose further that the contention thresholds of *all* the objects of a similarity class rise above the (essentially common) price of these objects. Then it can be seen that, without violating ϵ -CS, we can simultaneously raise the prices of all the objects of the class to the minimum contention threshold, thereby resolving the corresponding price war.

5.1. The auction algorithm with similar objects

The preceding idea can be implemented very simply in an algorithm that we call *auction algorithm with similar objects*. For each object j , we maintain a contention threshold \hat{p}_j . If j is unassigned, the contention threshold \hat{p}_j is equal to the initial price p_j ; each time j is assigned to a new person i , the contention threshold \hat{p}_j is set to ϵ plus the minimum level to which p_j should rise before i finds an object from a *different similarity class* equally attractive to j . The prices are determined by the contention thresholds. In particular, the price of an object j is the minimum contention threshold over the objects of the similarity class $M(j)$ of j

$$p_j = \min_{k \in M(j)} \hat{p}_k \quad (24)$$

Thus, *all objects in a similarity class have the same price but possibly different contention thresholds.*

More formally, the auction algorithm with similar objects is the same as the auction algorithm of Section 4 except for one difference: In the bidding phase, an unassigned person i finds his/her best object j_i and corresponding best value on the basis of the contention thresholds

$$j_i = \arg \max_{j \in A(i)} \{a_{ij} - \hat{p}_j\}, \quad v_i = \max_{j \in A(i)} \{a_{ij} - \hat{p}_j\} \quad (25)$$

Then i gets assigned to j_i and sets the contention threshold \hat{p}_{j_i} to p_{j_i} plus a bidding increment $v_i - w_i + \epsilon$, which is based on the second-best object from a different similarity class, that is, w_i is defined now as

$$w_i = \max_{j \in A(i), j \notin M(j_i)} \{a_{ij} - \hat{p}_j\} \quad (26)$$

[instead of $w_i = \max_{j \in A(i), j \neq j_i} \{a_{ij} - p_j\}$]. When the contention thresholds of all objects of a similarity class have been raised, the price of the class is also raised to the minimum contention threshold according to (24). Note that because of the relation $p_j = \min_{k \in M(j)} \hat{p}_k$ between prices and contention thresholds, \hat{p}_j can be replaced by p_j in (25). Thus, an unassigned person bids for the best object of the best similarity class, but the bidding increment is now based on the contention thresholds and can be significantly higher than the bidding increment based on prices; see Figure 9.

Example. For the problem of Figure 8, if initially the object prices and the contention thresholds are zero, in the first iteration the bidding person 1 will bid the contention threshold of object 1 to $\hat{p}_1 = a_{11} + \epsilon$ rather than ϵ ; in the second iteration the bidding person 2 will bid the contention threshold of object 1' to $\hat{p}_{1'} = a_{21} + \epsilon$, so the price of the similarity class $M(1)$ will increase to $\min\{a_{11}, a_{21}\} + \epsilon$; in the third iteration, if $a_{31} \leq \min\{a_{11}, a_{21}\} + \epsilon$, the bidding person 3 will bid for object 2 and the auction will terminate; otherwise person 3 will bid for either object 1 or 1' depending on which one has the smallest contention threshold, and the person that is left unassigned will bid for object 2 at the next iteration, thereby terminating the auction. It can thus be seen that the auction algorithm with similar objects very quickly resolves the price wars due to the duplicate objects corresponding to the sinks of the transportation problem (23).

Regarding the validity of the algorithm, it can be shown that the object prices still satisfy ϵ -CS and by essentially repeating the proof of Proposition 2, it follows that the algorithm will terminate for a feasible problem. However, it can be shown that for integer a_{ij} , the final assignment will be optimal if $\epsilon \in 1/m$, where

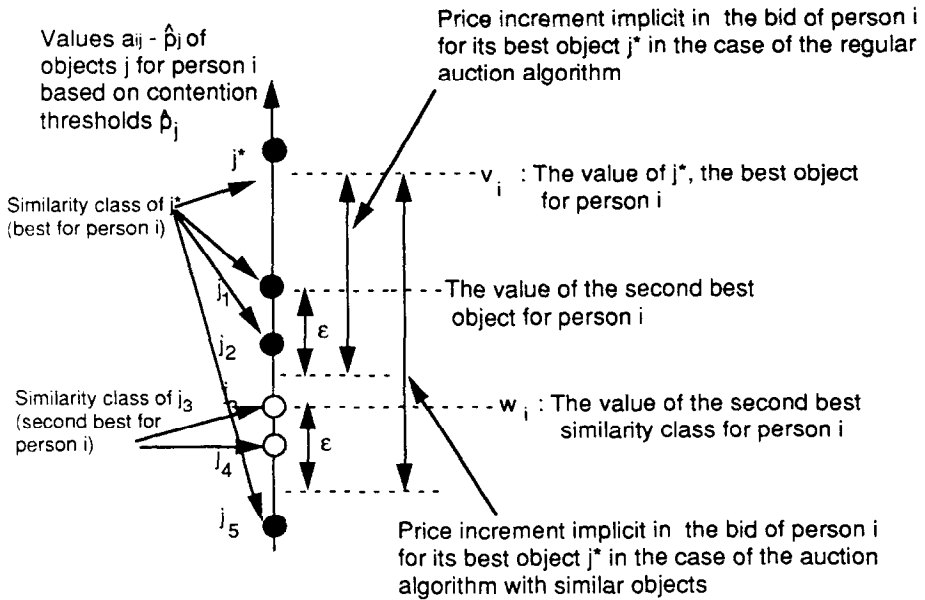


Figure 9. Illustration of the bid of person i in the auction algorithm with similar objects. The object j^* offers the best value $a_{ij} - \hat{p}_j$ for person i , based on the contention thresholds \hat{p}_j . However, here w_i is the value offered by the second-best similarity class, rather than the value of the second-best object as in the regular auction algorithm. When the second-best object belongs to the similarity class of j^* (as in this figure), the bid of person i will be higher in the auction algorithm with similar objects than in the regular auction algorithm.

m is the number of object similarity classes rather than $\epsilon < 1/n$ as is the case for the auction algorithm of Section 3; see [11], pp. 230–232, or [14], or use Proposition 10, which is proved in Appendix 1. Thus, the increased dimensionality due to the duplicate objects created by transforming the transportation problem (23) to an assignment problem is not reflected in a corresponding reduction of the threshold value of ϵ to obtain an optimal solution.

5.2. The auction algorithm with similar persons

One can further enhance the performance of the auction algorithm by taking into account the presence of similar persons. The idea is to submit a *common bid* for all persons in a similarity class if at least one person in the class is unassigned. As a result, persons in the same class do not compete against each other for the same objects, and the bids submitted are higher than they would otherwise be. In effect, the persons of the same similarity class cooperate to compress a price war and resolve it within one iteration. We illustrate this idea and the corresponding auction algorithm in Figure 10.

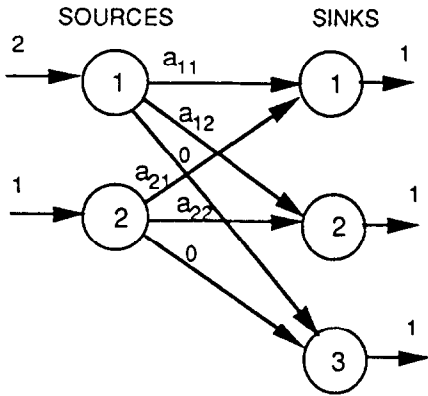
The idea of a common bid for a person similarity class can be combined with the previously discussed idea of contention thresholds and the corresponding auction algorithm for similar objects. In particular, the cooperative bid of a person similarity class is based on the contention thresholds of the objects. With proper streamlining of the computations, one obtains an algorithm for the transportation problem (23), which is, in effect, the auction algorithm for the corresponding equivalent assignment problem (cf. Figure 7), but with similar persons and objects properly taken into account. In this auction/transportation algorithm, each sink j has a price p_j and each arc (i, j) that carries positive flow x_{ij} has a contention threshold y_{ij} associated with it. The price p_j is equal to the original price of the sink j if the demand of the sink has not yet been exhausted ($\sum_{i \in B(j)} x_{ij} < \beta_j$) and is equal to the minimum contention threshold of the incoming arcs that carry positive flow ($\min_{i \in B(j), x_{ij} > 0} y_{ij}$) otherwise. The sources bid for sinks by increasing the contention threshold of the corresponding arcs, and a sink price increases once all of its demand is exhausted and the associated contention thresholds have increased. As the sink prices increase, the corresponding profits of the sources defined by

$$\pi_i = \max_{j \in A(i)} \{a_{ij} - p_j\} \quad (27a)$$


decrease, while the ϵ -CS condition

$$\pi_i + p_j \leq a_{ij} + \epsilon, \quad \forall (i, j) \in \mathcal{A} \text{ with } x_{ij} > 0 \quad (27b)$$

is maintained. The ϵ -CS conditions (27a) and (27b) extend the corresponding condition (8) for the assignment problem. It is shown in [14] that for an integer



Transportation Problem



Equivalent Assignment Problem with Similar Persons

Figure 10. Illustration of the auction algorithm with similar persons. In this algorithm, all the persons in a similarity class with k persons submit bids for the k best objects for the class. The levels of the bids are such that once the bids are accepted, the profit of each person in the class is set at ϵ below the profit offered by the $(k + 1)$ st best object. In the example of the figure the similar persons result from the conversion of the transportation source 1 into the two similar persons 1 and 1'. Consider first the case where the initial assignment is empty, all initial prices are 0, and $a_{11} = a_{12} = a_{21} = a_{22} = C > 0$. Then the regular auction algorithm will operate as in Figure 2, involving a long price war if ϵ is small relative to C . In the first iteration of the auction algorithm with similar persons, persons 1 and 1' will submit a common bid of $C + \epsilon$ for the two best objects 1 and 2. Then person 2 will submit a bid for object 3 at the next iteration and the algorithm will terminate, thus avoiding the price war. If, instead, we have $a_{11} > a_{12} > 0$, at the initial iteration persons 1 and 1' will submit a bid of $a_{11} + \epsilon$ for object 1 and a bid of $a_{12} + \epsilon$ for object 2, bringing the net value of these objects to the value of the third object minus ϵ . If $a_{12} - \epsilon > a_{22} > a_{21} > 0$, person 2 will then submit a bid for object 3 at the next iteration and the algorithm will terminate. If $a_{22} > a_{12} - \epsilon$ and $a_{22} > a_{21} > 0$, person 3 will submit a bid for object 2 at the second iteration, raising its price to $a_{22} + \epsilon$; persons 1 and 1' will then submit another common bid for objects 1 and 3 at the third iteration, and the algorithm will terminate.

problem, a feasible flow vector is optimal provided that there exists a price vector p and corresponding profit vector π satisfying ϵ -CS with $\epsilon < 1/\min\{m, n\}$.

6. A Generic Auction Algorithm for Minimum Cost Flow Problems

We now discuss a general algorithm of the auction type for linear minimum cost flow problems, given by D. Castañón and the author in [13] and [16]; it includes as special cases the auction algorithm for assignment and transportation problems, as well as the ϵ -relaxation method to be discussed in Section 7. By specializing this general algorithm to network flow problems with special structure, one may obtain efficient methods that exploit the structure; for example, the general algorithm is the basis for the k node-disjoint shortest path algorithm briefly described in Section 4.

We are given a directed graph with set of nodes \mathcal{N} and set of arcs \mathcal{A} . The number of nodes is denoted by N and the number of arcs is denoted by A . For each arc (i, j) there is an optimization variable x_{ij} called the *flow of arc* (i, j) . We denote by x the flow vector $\{x_{ij} \mid (i, j) \in \mathcal{A}\}$. The minimum cost flow problem that we consider is

$$\text{minimize} \quad \sum_{(i,j) \in \mathcal{A}} a_{ij} x_{ij}$$

subject to

$$\sum_{j|(i,j) \in \mathcal{A}} x_{ij} - \sum_{\{j|(j,i) \in \mathcal{A}\}} x_{ji} = s_i, \quad \forall i \in \mathcal{N} \quad (28)$$

$$b_{ij} \leq x_{ij} \leq c_{ij}, \quad \forall (i, j) \in \mathcal{A} \quad (29)$$

where a_{ij} , b_{ij} , c_{ij} , and s_i are given integers. The problem is said to be feasible if there is a flow vector x satisfying the constraints (28) and (29); otherwise it is said to be infeasible.

The problem can be converted into an equivalent transportation problem as shown in Figure 11. It is thus possible to use this equivalence to transcribe the auction/transportation ideas of the previous section into the minimum cost flow context. In particular, let us first derive the appropriate form of ϵ -CS.

In reference to the equivalent transportation problem of Figure 11, let us denote by λ_{ij} the price of the sink corresponding to arc (i, j) , and let us denote by p_i the profit of the source corresponding to node i . By (27a) (with a_{ij} replaced by $-a_{ij}$ to account for the change from maximization to minimization), we have

$$p_i + \lambda_{ij} \geq 0, \quad p_j + \lambda \geq -a_{ij}, \quad \text{for all } (i, j) \in \mathcal{A} \quad (30)$$

Let us assume that throughout the auction the total incoming flow to the sink corresponding to arc (i, j) along the two transportation arcs $(j, (i, j))$ and $(i, (i, j))$ is equal to the demand $c_{ij} - b_{ij}$; for any initial set of prices p_i it is




Figure 11. Transformation of a minimum cost flow problem into a transportation problem of the form (23). We take as sinks of the transportation problem the arcs of the original network, and as sources of the transportation problem the nodes of the original network. Each transportation problem sink has two incoming arcs with cost coefficients as shown. The demand of each transportation problem sink is the feasible flow range length of the corresponding original network arc. The demand of each transportation problem source is the sum of the feasible flow range lengths of the outgoing arcs from the corresponding original network node minus the supply of that node as shown. An arc flow x_{ij} in the minimum cost flow problem corresponds to flows equal to $x_{ij} - b_{ij}$ and $c_{ij} - x_{ij}$ on the transportation problem arcs $(j, (i, j))$ and $(i, (i, j))$, respectively.




Figure 12. Illustration of ϵ -CS. All pairs of arc flows x_{ij} and price differences $p_i - p_j$ should either lie on the thick lines or in the shaded area between the thick lines.

possible to enforce this condition while preserving ϵ -CS by selecting the initial λ_{ij} to be $\lambda_{ij} = \min \{-p_i, -p_j - a_{ij}\}$. Then, by (27b) we also have

$$p_i + \lambda_{ij} \leq \epsilon \quad \text{for all } (i, j) \in \mathcal{A} \text{ with } x_{ij} < c_{ij} \quad (31a)$$

$$p_j + \lambda_{ij} \leq -a_{ij} + \epsilon \quad \text{for all } (i, j) \in \mathcal{A} \text{ with } b_{ij} < x_{ij} \quad (31b)$$

From (30) and (31a), we obtain $p_j + a_{ij} \geq -\lambda_{ij} \geq p_i - \epsilon$ if $x_{ij} < c_{ij}$, so that

$$p_i - p_j \leq a_{ij} + \epsilon \quad \text{for all } (i, j) \in \mathcal{A} \text{ with } x_{ij} < c_{ij} \quad (32a)$$

Also from (30) and (31b), we obtain $p_i \geq -\lambda_{ij} \geq a_{ij} + p_j - \epsilon$ if $b_{ij} < x_{ij}$, so that

$$p_i - p_j \geq a_{ij} - \epsilon \quad \text{for all } (i, j) \in \mathcal{A} \text{ with } b_{ij} < x_{ij} \quad (32b)$$

We say that a flow-price pair (x, p) satisfies ϵ -complementary slackness (ϵ -CS) if x satisfies the arc flow bounds (29), and (32a) and (32b) hold; see Figure 12. The following proposition parallels Proposition 1 for the assignment problem and is proved in more general form in Appendix 1 (cf. Proposition 10).

PROPOSITION 6. *If x is feasible, and for some price vector p , the pair (x, p) satisfies ϵ -CS with $\epsilon < 1/N$ (N is the number of nodes), then x is optimal for the minimum cost flow problem.*

To develop auction algorithms for the minimum cost flow problem, we draw motivation from auction algorithms for the equivalent transportation problem and the associated assignment problem with similar objects and persons. The notion of a bid by a single source and the associated increase of flow along the best outgoing arc can be transcribed in the context of the minimum cost problem as a flow increase by a single node along the best outgoing arcs or a flow reduction by a node along the best incoming arcs (cf. Figure 11). These flow modifications could also be followed by an increase of the node price, while maintaining the ϵ -CS condition. Since in the context of the transportation algorithm we used notions of cooperative bids by several nodes, it also makes sense in the corresponding minimum cost flow context to consider simultaneous price increases by several nodes.

Motivated by these associations, we now introduce some computational operations that will constitute the building blocks for a general auction algorithm for the minimum cost flow problem. Each of the following definitions assumes that (x, p) is a flow-price vector pair satisfying ϵ -CS.

DEFINITION 1. *An arc (i, j) is said to be ϵ^+ -unblocked if*

$$p_i = p_j + a_{ij} + \epsilon, \quad x_{ij} < c_{ij}$$

An arc (j, i) is said to be ϵ^- -unblocked if

$$p_i = p_j - a_{ji} + \epsilon, \quad b_{ji} < x_{ji}$$

The push list of a node i is the (possibly empty) set of the arcs (i, j) that are ϵ^+ -unblocked, and the arcs (j, i) that are ϵ^- -unblocked

In the subsequent algorithms, flow is allowed to increase only along ϵ^+ -unblocked arcs and is allowed to decrease only along ϵ^- -unblocked arcs. The next definition specifies the type of flow changes considered.

DEFINITION 2. *For an arc (i, j) [or arc (j, i)] of the push list of node i , let δ be a scalar such that $0 < \delta \leq c_{ij} - x_{ij}$ ($0 < \delta \leq x_{ji} - b_{ji}$, respectively). A δ -push at node i on arc (i, j) [(j, i) , respectively] consists of increasing the flow x_{ij} by δ (decreasing the flow x_{ji} by δ , respectively), while leaving all other flows, as well as the price vector unchanged.*

It is evident from the definitions that a δ -push preserves ϵ -CS. The next operation consists of raising the prices of a subset of nodes by the maximum common increment γ for which ϵ -CS will not be violated.

DEFINITION 3. *A price rise of a nonempty, strict subset of nodes I (i.e., $I \neq \emptyset$, $I \neq \mathcal{N}$), consists of leaving unchanged the flow vector x and the prices of nodes not belonging to I , and of increasing the prices of the nodes in I by the amount γ given by*

$$\gamma = \begin{cases} \min \{S^+ \cup S^-\}, & \text{if } S^+ \cup S^- \neq \emptyset \\ 0, & \text{if } S^+ \cup S^- = \emptyset \end{cases} \quad (33)$$

where S^+ and S^- are the sets of scalars given by

$$S^+ = \{p_j + a_{ij} + \epsilon - p_i \mid (i, j) \in \mathcal{A} \text{ such that } i \in I, j \notin I, x_{ij} < c_{ij}\} \quad (34)$$

$$S^- = \{p_j - a_{ji} + \epsilon - p_i \mid (j, i) \in \mathcal{A} \text{ such that } i \in I, j \notin I, x_{ji} < b_{ji}\} \quad (35)$$

It can be verified using the definitions that a price rise maintains ϵ -CS. Note that every scalar in the sets S^+ and S^- of (34) and (35) is nonnegative by the ϵ -CS conditions (32a) and (32b), respectively, so the scalar γ is nonnegative. If $\gamma > 0$ the price rise is said to be *substantive* and if $\gamma = 0$, the price rise is said to be *trivial*. (A trivial price rise changes neither the flow vector nor the price vector; it is introduced to facilitate the presentation.)

In the case where the subset I consists of a single node i , a price rise of the singleton set $\{i\}$ is also referred to as a *price rise of node i* . A price rise of a single node i is substantive if and only if the set $S^+ \cup S^-$ is nonempty but the push list of i is empty. It can be shown that, for a feasible problem, if the push list of a node i with $g_i > 0$ is empty, then the set $S^+ \cup S^-$ must be nonempty.

6.1. The generic algorithm

The generic algorithm to be described shortly consists of a sequence of δ -push and price rise operations. The order of execution of these operations is subject to very mild restrictions, thus allowing a great deal of flexibility to exploit the structure of the problem at hand.

Suppose that the minimum cost flow problem is feasible, and consider a pair (x, p) satisfying ϵ -CS. For a given flow vector x , let us define the *surplus g_i of node i* as the difference of the total flow coming into i and the total flow coming out of i , that is,

$$g_i = \sum_{\{j \mid (j, i) \in \mathcal{A}\}} x_{ji} - \sum_{\{j \mid (i, j) \in \mathcal{A}\}} x_{ij} + s_i$$

Consider a node i with $g_i > 0$. There are two possibilities:

1. The push list of i is nonempty, in which case a δ -push at node i is possible.
2. The push list of i is empty, in which case, as mentioned earlier, the price rise of node i will be substantive.

Thus, if $g_i > 0$ for some i and the problem is feasible, then either a δ -push or a substantive price rise is possible at node i . Furthermore, since following a price rise at a node i , the push list of i will be nonempty [see (33)–(35)], for a

feasible problem, a δ -push is always possible at a node i with $g_i > 0$, perhaps following a price rise at i .

The preceding observations form the basis for a method, called *generic algorithm*, which uses a fixed positive value of ϵ and starts with a pair (x, p) satisfying ϵ -CS. The algorithm terminates when $g_i \leq 0$ for all nodes i ; otherwise it continues to perform iterations. Each iteration consists of a sequence of δ -pushes and price rises, including at least one δ -push, as described below.

Typical iteration of the generic algorithm

Perform in sequence and in any order a finite number of δ -pushes and price rises; there should be at least one δ -push but not necessarily at least one price rise. Furthermore:

1. Each δ -push should be performed at some node i with $g_i > 0$ and the flow increment δ must satisfy $\delta \leq g_i$.
2. Each price rise should be performed on a set I with $g_i \geq 0$ for all $i \in I$.

The following proposition, proved in [11], [13], and [16], establishes the validity of the generic algorithm.

PROPOSITION 7. *Assume that the minimum cost flow problem is feasible. If the increment δ of each δ -push is integer, then the generic algorithm terminates with a pair (x, p) satisfying ϵ -CS. The flow vector x is feasible, and is optimal if $\epsilon < 1/N$.*

The idea of the proof is that δ -pushes cannot be indefinitely performed without some intermediate substantive price rises. Thus, if the algorithm does not terminate, the prices of some nodes with positive surplus must increase to infinity, while the prices of some other nodes with negative surplus remain at their initial level. With some thought, it can be seen that this implies that for a feasible problem, ϵ -CS must be violated, leading to a contradiction.

If the problem is infeasible, the algorithm may not terminate. One way to deal with infeasibility is to convert the problem to one that is guaranteed to be feasible by introducing artificial high-cost arcs. The appropriate level of cost can be quantified similar to the case of assignment problems; see (11) and (12), and [11].

7. The ϵ -Relaxation Method

The price rise operations of the generic algorithm may involve several nodes. When we require that only one node is involved in each price rise, we obtain the ϵ -relaxation method first proposed by the author in [9] and [10], and described in this section.

We use a fixed positive value of ϵ and we start with a pair (x, p) satisfying

ϵ -CS. Furthermore, the starting arc flows are integer; the algorithm preserves the integrality of the arc flows. At the start of a typical iteration we have a flow-price pair (x, p) satisfying ϵ -CS and we select a node i with $g_i > 0$; if no such node can be found, the algorithm terminates. During the iteration we perform several δ -pushes and price rises involving node i .

Typical iteration of the ϵ -relaxation method

Select a node i with $g_i > 0$. If no such node exists, terminate the algorithm; else go to step 1.

Step 1: If the push list of node i is empty go to step 3; else select an arc a from the push list of i and go to step 2.

Step 2: Let j be the end node of arc a , which is opposite to i . Let

$$\delta = \begin{cases} \min\{g_i, c_{ij} - x_{ij}\} & \text{if } a = (i, j) \\ \min\{g_i, x_{ji} - b_{ji}\} & \text{if } a = (j, i) \end{cases} \quad (36)$$

Perform a δ -push of i on arc a . If, as a result of this operation, we obtain $g_i = 0$ go to step 3; else go to step 1.

Step 3: Perform a price rise of node i . If $g_i = 0$ stop; else go to step 1.

It is straightforward to verify that the ϵ -relaxation method is a special case of the generic algorithm of Section 6 to which Proposition 7 applies. As a result, for a feasible problem, the algorithm terminates with a pair (x, p) satisfying ϵ -CS and the flow x is optimal if $\epsilon < 1/N$. Similar to assignment problems, it is possible to use ϵ -scaling in conjunction with the method, and this is frequently essential for good practical performance.

The computational complexity of the unscaled version of the ϵ -relaxation method was first studied in [9]. The complexity of the scaled version was first studied in [30], where particularly favorable polynomial running time estimates were derived; see [19], [20], and [33] for additional results along these lines.

8. Application of the ϵ -Relaxation Method to Max-Flow Problems

The ϵ -relaxation method can be applied to the max-flow problem once this problem is converted to the minimum cost flow format, involving a feedback arc connecting the sink with the source, and having cost -1 as shown in Figure 13. Because of the small cost range one can forego ϵ -scaling here, while still maintaining polynomial complexity [$O(N^3)$, or even better with appropriate implementation]. In practice one can also typically forego ϵ -scaling with impunity




Figure 13. The max-flow problem and its minimum cost flow representation. Here there are two special nodes: the source (s) and the sink (t). The objective is to push as much flow as possible from s into t while observing the arc capacity constraints. We introduce an artificial arc (t, s) with cost -1 and very large upper flow bound and we assign cost zero to all other arcs. At the optimum, the flow x_{ts} equals the maximum flow that can be sent from s to t through the original graph.

but some additional computational tricks are needed to nullify the effects of price wars; see Section 10.5.

If ϵ -scaling is not used, one can change the cost of the feedback arc (t, s) to $a_{ts} = -(N + 1)$ and use $\epsilon = 1$ throughout. When this is done, the ϵ -relaxation method bears a close resemblance with a max-flow algorithm first proposed in [31] and [32], and often called “push-relabel” algorithm. This latter max-flow algorithm was derived from a different point of view that is unrelated to duality or ϵ -CS. It uses node “labels,” which in the context of the ϵ -relaxation approach can be viewed as prices. The max-flow version of the ϵ -relaxation method, first given in [9] and [10], is simpler than the algorithm of [31] and [32] in that it obtains a maximum flow in one phase rather than two (single-phase versions were also given later by several other authors; see [1]). It can also be initialized with arbitrary prices, whereas in the max-flow algorithm of [31] and [32], the initial prices must satisfy $p_i \leq p_j + 1$ for all arcs (i, j) . This can be significant, for example if one wants to use scaling or more generally, in a reoptimization setting where one wants to capitalize on the results of the solution of a slightly different max-flow problem. Related max-flow algorithms and their computational complexity are discussed in [1], [21], [26], and [41].

9. Extension to Asymmetric Assignment Problems

A slight variation of the auction algorithm can be used for asymmetric assignment problems, where the number of persons is m and the number of objects is n with $m < n$, while there is a requirement that each person be assigned to some object. To solve this problem, the auction algorithm need only be modified in the choice of initial conditions; it is sufficient to require that all initial prices be zero, as was noted in the original paper ([7]).

Unfortunately, this approach to the asymmetric assignment problems has a deficiency: Since the initial object prices must be zero, the use of ϵ -scaling is impossible. As a result the method is susceptible to price wars and the associated long running times. To address this difficulty, one may convert the asymmetric problem to a symmetric one by adding $n - m$ artificial persons that can be assigned to any object at zero cost. This introduces an undesirable increase in the problem's dimension. One could use the auction algorithm with similar persons of Section 5 to allow ϵ -scaling while taking advantage of the structure induced by the artificial persons. Nonetheless, the indirect approach of converting an asymmetric assignment problem into a symmetric one has not seen much use. We now discuss a direct and probably more effective auction approach for asymmetric assignment problems.

9.1. Reverse auction for asymmetric assignment problems

It is possible to use reverse auction in conjunction with forward auction to provide asymmetric assignment algorithms that use ϵ -scaling. We first introduce the proper form of the ϵ -CS condition for an assignment S and a pair (π, p) . The following proposition, proved in [11] and [18], parallels Proposition 1 for the symmetric assignment problem.

PROPOSITION 8. *If a feasible assignment S and a pair (π, p) satisfy the conditions*

$$\pi_i + p_j \geq a_{ij} - \epsilon, \quad \forall (i, j) \in \mathcal{A} \quad (37a)$$

$$\pi_i + p_j = a_{ij}, \quad \forall (i, j) \in S \quad (37b)$$

$$p_j \leq \min_{k: \text{assigned under } S} p_k, \quad \forall j : \text{unassigned under } S \quad (37c)$$

then S is within $m\epsilon$ of being optimal for the asymmetric assignment problem.

Now consider trying to solve the asymmetric assignment problem by means of auction. We can start with any assignment S and pair (π, p) satisfying the first two ϵ -CS conditions (37a) and (37b), and perform a forward auction (as defined earlier for the symmetric assignment problem) up to the point where

each person is assigned to a distinct object. For a feasible problem, it can be seen that this will yield, in a finite number of iterations, a feasible assignment S satisfying the first two conditions (37a) and (37b). This assignment may not be optimal because the prices of the unassigned objects may not be minimal, that is, they may not satisfy the third ϵ -CS condition (37c). Roughly, what is happening here is that forward auction cannot resolve whether the objects that were left unassigned upon termination are intrinsically undesirable because they offer relatively low benefit to the persons, or whether they were left unassigned because their initial prices were high relative to the initial prices of the assigned objects.

To resolve this dilemma, we use a modified form of reverse auction to lower the prices of the objects that were left unassigned upon termination of the forward auction. After several reverse auction iterations in which persons may be reassigned to other objects, the third condition (37c) is satisfied. The assignment thus obtained can be shown to satisfy all the ϵ -CS conditions (37a)–(37c) and by Proposition 8, is optimal within $m\epsilon$ (and strictly optimal if the problem data are integer and $\epsilon < 1/m$).

The modified reverse auction starts with a feasible assignment S and with a pair (π, p) satisfying the first two ϵ -CS conditions (37a) and (37b). [For a feasible problem, such an S and (π, p) can be obtained by regular forward or reverse auction, as discussed earlier.] Let us denote by λ the minimal assigned object price under the initial assignment,

$$\lambda = \min_{j: \text{ assigned under the initial assignment } S} p_j$$

The typical iteration of modified reverse auction is the same as the one of reverse auction, except that only unassigned objects j with $p_j > \lambda$ participate in the auction. In particular, the algorithm maintains a feasible assignment S and a pair (π, p) satisfying (37a) and (37b), and terminates when all unassigned objects j satisfy $p_j \leq \lambda$, in which case it will be seen that the third ϵ -CS condition (37c) will be satisfied as well. The scalar λ will be kept fixed throughout the algorithm.

Typical iteration of modified reverse auction for asymmetric assignment

Select an object j that is unassigned under the assignment S , and satisfies $p_j > \lambda$ (if no such object can be found, the algorithm terminates). Find a best person i_j such that

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

and the corresponding value

$$\beta_j = \max_{i \in B(j)} \{a_{ij} - \pi_i\} \tag{38}$$

and find

$$\omega_j = \max_{i \in B(j), i \neq i_j} \{a_{ij} - \pi_i\} \quad (39)$$

[If i_j is the only person in $B(j)$, we define ω_j to be $-\infty$.] If $\lambda \geq \beta_j - \epsilon$, set $p_j := \lambda$ and go to the next iteration. Otherwise, let

$$\delta = \min \{\beta_j - \lambda, \beta_j - \omega_j + \epsilon\} \quad (40)$$

set

$$p_j := \beta_j - \delta \quad (41)$$

$$\pi_{i_j} := \pi_{i_j} + \delta \quad (42)$$

add to the assignment S the pair (i_j, j) , and remove from S the pair (i_j, j') , where j' is the object that was assigned to i_j under S at the start of the iteration.

Note that the formula (40) for the bidding increment δ is such that the object j enters the assignment at a price which is no less than λ [and is equal to λ if and only if the minimum in (40) is attained by the first term]. Furthermore, we have $\delta \geq \epsilon$ (when δ is calculated, that is, when $\lambda < \beta_j - \epsilon$), so it can be seen from (41) and (42) that throughout the algorithm, prices are monotonically decreasing and profits are monotonically increasing. The following proposition, first proved in [18] (see also [11]), asserts the validity of the method.

PROPOSITION 9. *The modified reverse auction algorithm for the asymmetric assignment problem terminates in a finite number of iterations and the assignment obtained is within $m\epsilon$ of being optimal.*

As mentioned earlier, forward auction followed by modified reverse auction can start with arbitrary initial prices. As a result, one can use ϵ -scaling, performing a sequence of auctions with decreasing values of ϵ . Out of several possible variations of the method, the one found most effective in [18] is to use the modified reverse auction only in the last ϵ -scaling phase. In all other ϵ -scaling phases, forward auction is used exclusively. There are also alternative algorithms for the asymmetric assignment problems that switch between the forward and the reverse methods within each ϵ -scaling phase; see [15], which gives particularly favorable computational results for difficult asymmetric assignment problems.

9.2. Forward/reverse auction for other types of inequality-constrained problems

Reverse auction can also be used to solve the variation of the *two-sided inequality-constrained assignment problem*, where persons (as well as objects) need not be assigned if this degrades the assignment's total benefit. This problem can be converted to an asymmetric assignment problem where all persons must be

assigned by introducing for each person i an artificial object i' and a zero benefit pair (i, i') . One can then use the algorithm given earlier to solve this problem. The algorithm can be streamlined so that the calculations involving the artificial objects and pairs are handled efficiently.

Finally, a class of interesting assignment-like problems, called *multiassignment problems*, can be solved efficiently by using combined forward/reverse auction ideas [18]. In these problems, persons can be assigned to more than one object, and objects can be assigned to more than one person simultaneously (the number of assignments for each person and object may or may not be subject to an upper bound).

10. Practical Computational Aspects of Auction Algorithms

In this section we provide an overview of the computational aspects of auction algorithms, concentrating primarily on the case of the assignment problem. Generally, at present, it appears that for simple problems such as assignment, shortest path, and max-flow, auction algorithms are at least competitive and often far superior to the classical methods. For problems with more complex structures, auction algorithms have yet to match the performance of their closest competitors such as primal simplex, primal-dual, and relaxation methods. Generally, however, auction algorithms are better suited for parallel computation.

10.1. Assignment problems

In practice, the auction algorithm has proved very effective for assignment problems. Computational studies ([18] and [20]) have shown that it is often substantially faster than its closest competitors, particularly for sparse problems. The speedup factor typically increases with the problem size. As the problem density increases, the speedup factor tends to decrease, and for fully dense problems, the auction algorithm is roughly competitive to methods that combine the primal-dual (or sequential shortest path method) with an extensive initialization procedure based on the naive auction algorithm of Section 2 ([12] and [34]). A recent extensive computational study [24] explores the behavior of various auction algorithms for randomly generated problems with a broad variety of different structures.

10.1.1. Adaptive ϵ -scaling. There are a number of variations of the ϵ -scaling technique aimed at improving computational efficiency. There are many types of problems where ϵ -scaling is essential because of the high likelihood of price wars. There are also other types of problems for which price wars are highly unlikely and it is better to either forego ϵ -scaling altogether or implement it aggressively, that is, reduce ϵ very quickly to its ultimate value. This motivates

scaling techniques, known collectively as *adaptive scaling*, where ϵ is reduced aggressively or conservatively based on the method's progress.

One type of adaptive scaling uses two values of ϵ within each scaling phase, the *ceiling value* denoted by $\bar{\epsilon}$, and the *current value* denoted by ϵ . The ceiling value $\bar{\epsilon}$ is kept constant within each scaling phase and is reduced at the end of the phase [up to an ultimate value of $1/(n+1)$]. The current value ϵ is the one used in the bidding increment $\gamma_i = w_i - v_i + \epsilon$ [cf. (9)], and may change within the range $[1/(n+1), \bar{\epsilon}]$. If the current value coincides with the ceiling value, the standard form of (nonadaptive) scaling is obtained. Scaling becomes truly adaptive if ϵ is started at a relatively small value [such as $1/(n+1)$] and is gradually increased according to various criteria towards the ceiling value $\bar{\epsilon}$. For example, ϵ may be kept constant as long as a new pair is added to the assignment within a fixed number of iterations, and may be gradually increased otherwise.

Adaptive scaling may also be useful when high-cost artificial pairs are introduced to guarantee that the problem is feasible as described earlier. Because of these pairs, the cost range of the problem is greatly expanded, thereby potentially increasing the number of scaling phases. On the other hand, the high-cost pairs are superfluous when the problem is feasible, and one may use a form of adaptive scaling to recognize this case and accordingly reduce ϵ to the appropriate levels.

Generally, the effective use of adaptive scaling requires some ingenuity and trial and error, as well as a fair amount of intuition into the structure of the problem at hand.

10.1.2. The problem of integer overflow. If the benefits a_{ij} are integer, the auction algorithm can, in most cases, be executed using integer arithmetic. The most common way to do this is to use integer starting prices, to multiply a_{ij} with $(n+1)$, and then to use integer values of ϵ and a final value of $\epsilon=1$. The generated prices and other quantities will then be integer. There is a potential difficulty however; some prices may become large enough to exceed the integer range of the computer. In particular, from (11) we see that an upper bound for p_j is of the order of $n^2 \max_{(i,j) \in \mathcal{A}} |a_{ij}|$ [since a_{ij} has been multiplied by $(n+1)$ to work with integer arithmetic], which is large enough to lie outside the integer range of many computers even for moderate values of n . To remedy this situation, one may store the prices in floating-point format (preferably double precision) and use floating-point arithmetic to execute the price update calculations, while using integer arithmetic for all other computations. In some experiments reported in [24], it was observed that this involves a slowdown of the practical running time of the algorithm by a factor of no more than two. Thus, it appears that the difficulty of integer overflow can be effectively addressed at a relatively small cost.

10.1.3. The third-best implementation. Frequently in the auction algorithm, the two best objects for a given person do not change between two successive bids of that person. This motivates an implementation idea (developed by the author

in collaboration with H. Tsaknakis—unpublished) that attempts to exploit this fact by using a test to check whether the two best objects from the previous bid continue to be best. If the test is passed, the computation of the values $a_{ij} - p_j$ of the remaining objects j is unnecessary.

Consider an auction algorithm iteration when we calculate the bid of the person i on the basis of a price vector p . Suppose that in addition to the best value $v_i = \max_{j \in A(i)} \{a_{ij} - p_j\}$, the best object $j_1 = \arg \max_{j \in A(i)} \{a_{ij} - p_j\}$, and the second-best value $w_i = \max_{j \in A(i), j \neq j_1} \{a_{ij} - p_j\}$, we compute:

1. The second-best object

$$j_2 = \arg \max_{j \in A(i), j \neq j_1} \{a_{ij} - p_j\}$$

2. The third-best value

$$y_i = \max_{j \in A(i), j \neq j_1, j \neq j_2} \{a_{ij} - p_j\}$$

Consider now a subsequent iteration where person i bids based on an updated price vector $\bar{p} \geq p$. If we have

$$a_{ij_1} - \bar{p}_{j_1} \geq y_i \tag{43}$$

and

$$a_{ij_2} - \bar{p}_{j_2} \geq y_i \tag{44}$$

it can be seen that j_1 and j_2 continue to be the two best objects for i . The reason is that the object prices cannot decrease in the course of the algorithm, implying that $y_i \geq a_{ij} - \bar{p}_j$ of all j other than j_1 and j_2 . As a result, if the tests (43) and (44) are passed, we can forego the calculation of the values $a_{ij} - p_j$ for the objects j other than j_1 and j_2 . If, on the other hand, one of the tests (43) and (44) is violated, we cannot make any inference regarding the best and second-best objects; we must then exhaustively compare the values of all objects $j \in A(i)$, and compute the best, second-best, and third-best values. Computational experience has shown that the savings in the calculation of object values whenever the tests (43) and (44) are passed far outweigh the overhead for maintaining the third-best values y_i , and for performing the tests (43) and (44). This is particularly so when the problem is fairly dense, so that the set $A(i)$ has many objects.

10.1.4. Parallel and asynchronous implementation. Both the bidding and the assignment phases of the auction algorithm are highly parallelizable. In particular, the bids can be computed simultaneously and in parallel for all persons participating in the auction. Similarly, the subsequent awards to the highest bidders

can be computed in parallel by all objects that received a bid. Furthermore, the bid of a single person can be computed cooperatively by several processors in parallel. There have been several implementations of the auction algorithm in parallel shared memory machines ([17], [35], and [52]), and in SIMD machines ([25], [45], [50], and [51]). If there is strict temporal separation between the bid calculation phase and the highest bidder award phase, the implementation is said to be *synchronous*. In such an implementation, there are two basic methods to parallelize the bidding phase for the set I of unassigned persons that submit a bid, and a third method that is a combination of the other two. Let m be the number of processors:

1. *Parallelization across bids (or Jacobi parallelization)*: Here the calculations involved in the bid of each person $i \in I$ are carried out by a single processor. If the number of persons in I , call it $|I|$, exceeds the number of processors m , some processors will execute the calculations involved in more than one bid. If $|I| < m$ then $m - |I|$ processors will be idle during the bidding phase, thereby reducing efficiency. (This will typically happen in the late stages of the algorithm.)
2. *Parallelization within a bid (or Gauss-Seidel parallelization)*: Here the set I consists of a single person i . The calculations involved in the bid of person i are shared by the m processors of the system. Thus, the set of admissible objects $A(i)$ is divided in m groups of objects $A_1(i), A_2(i), \dots, A_m(i)$ [assuming the number of processors is less than the number of objects in $A(i)$; otherwise some of the processors will be left idle]. The best object, best value, and second-best value are calculated within each group in parallel by a separate processor. After these calculations are completed (a synchronization of the processors is required to check this) the results are merged by one of the processors who finds the best value over all best group values, while simultaneously computing the corresponding best object and bid increment. (It is possible to do the merging in parallel using several processors, but in the absence of special SIMD-type hardware, this may be inefficient, particularly when the number of processors is small, because of the extra synchronization and other overhead involved.) The drawback of this method over the preceding one is that it typically requires a larger number of iterations, since each iteration involves a single person. This is significant because even though each Gauss-Seidel iteration may take less time because it is executed by multiple processors in parallel, the synchronization overhead is roughly proportional to the number of iterations.
3. *Hybrid approach (or block Gauss-Seidel parallelization)*: In this approach, the bid calculations of each person are parallelized as in the preceding method, but the number of processors used per bid is s , where $1 < s < m$. Thus, we can compute the bids of roughly m/s persons in parallel, assuming enough unassigned persons are available for the iteration. With proper choice of s , this method combines the best features and alleviates the drawbacks of the preceding two.

Once the bidding phase of an iteration is completed (a synchronization point), the assignment phase is executed. This phase may be carried out by a single processor. It is also possible to consider using multiple processors to execute the assignment phase in parallel, but the potential gain from parallelization may be modest while the associated overhead may more than offset this gain, as suggested by one computational study [17] that used a shared memory machine.

There are also *totally asynchronous* implementations of the auction algorithm, which are interesting because they are quite flexible and may also result in faster solution. To our knowledge, all parallel asynchronous implementations of the auction algorithm to date have used a shared memory machine. In one such implementation ([17]), the bidding and merging calculations of the Gauss-Seidel method are divided in tasks, which are organized in a first in–first out queue. When a processor becomes free it starts executing the top task of the queue if the queue is nonempty, and otherwise it checks whether a termination condition is satisfied. The algorithm stops when all processors encounter the termination condition.

Similar to the synchronous block Gauss-Seidel implementation, each set of admissible objects $A(i)$ is divided in s groups of objects $A_1(i), \dots, A_s(i)$. The calculation of the bid of a person i is divided in s tasks. The first $s - 1$ tasks are search tasks involving the groups of objects $A_1(i), \dots, A_{s-1}(i)$. To perform one of these tasks, a processor must calculate and store in memory the best value, second-best value, and best object within the corresponding object group. The s th task starts with a search and memory storage of the best value, second-best value, and best object within the group $A_s(i)$, and following this, it completes the bid of person i by merging the individual group search results, that is, by finding the best object and bid for person i based on the currently stored group results. The s th task also includes raising the price of the best object and changing the assignment of the object (assuming the calculated bid is larger than the best object's price by at least ϵ).

There are two sources of asynchronism here. First, it is possible for some prices to be changed between the time a search task is completed and the time the results of that task are used to calculate a person's bid. Second, it is possible that the merging task of a person's bid is carried out before some of the search tasks associated with that bid are completed. In both cases the bid may reflect out-of-date price information and may prove ineffective in that it yields a bid that does not exceed the corresponding best object's price by at least ϵ (if this occurs, one should simply cancel the bid and forego the corresponding update of the object's price and assignment).

The advantage of the asynchronous implementation is that processors do not remain idle, waiting to get synchronized with other processors or waiting for merging tasks to be completed. A careful formulation of the totally asynchronous model, and a proof of its validity is given in [17], which also includes extensive computational results on a shared-memory machine, confirming the advantage of asynchronous over synchronous implementations.

Computational experience so far suggests that generally, because of the typically prolonged “endgame” of the auction algorithm, where only a small percentage of the persons remains unassigned, the speedup that can be obtained by Jacobi parallelization is relatively modest (in the order of three to six). The attainable speedup for Gauss-Seidel parallelization or hybrid schemes is potentially higher, particularly when the problem is dense, and also when special hardware with vector processing capabilities are used.

10.2. Shortest path problems

Aside from combining the forward auction/shortest path algorithm with its reverse counterpart, there are a number of useful implementation ideas.

The main computational bottleneck of the algorithm is the calculation of

$$\min_{(i,j) \in \mathcal{A}} \{a_{ij} + p_j\}$$

which is done every time node i becomes the terminal node of the path. We can reduce the number of these calculations using the following observation. Since the CS condition $p_i \leq a_{ij} + p_j$ is maintained at all times for all arcs (i, j) , if some (i, j_i) satisfies

$$p_i = a_{ij_i} + p_{j_i}$$

it follows that

$$a_{ij_i} + p_{j_i} = \min_{(i,j) \in \mathcal{A}} \{a_{ij} + p_j\}$$

so the path can be extended by j_i if i is the terminal node of the path. This suggests the following implementation strategy: each time a path contraction occurs with i being the terminal node, we calculate

$$\min_{(i,j) \in \mathcal{A}} \{a_{ij} + p_j\}$$

together with an arc (i, j_i) such that

$$j_i = \arg \min_{(i,j) \in \mathcal{A}} \{a_{ij} + p_j\}$$

At the next time node i becomes the terminal node of the path, we check whether the condition $p_i = a_{ij_i} + p_{j_i}$ is satisfied, and if it is we extend the path by node j_i without going through the calculation of $\min_{(i,j) \in \mathcal{A}} \{a_{ij} + p_j\}$. In practice, this device is very effective, typically saving from one-third to one-half of the calculations of the preceding expression. The reason is that the test $p_i = a_{ij_i} + p_{j_i}$ rarely fails; the only way it can fail is if the price p_{j_i} is increased between the two

successive times i became the terminal node of the path. For some theoretical substantiation of this point, see [22].

The preceding idea can be strengthened further. Suppose that whenever we compute the best neighbor

$$j_i = \arg \min_{(i,j) \in \mathcal{A}} \{a_{ij} + p_j\}$$

we also compute the second-best neighbor k_i given by

$$k_i = \arg \min_{(i,j) \in \mathcal{A}, j \neq j_i} \{a_{ij} + p_j\}$$

and the corresponding second-best level

$$w_i = a_{ik_i} + p_{k_i}$$

Then, at the next time node, i becomes the terminal node of the path, we can check whether the condition $a_{ij_i} + p_{j_i} \leq w_i$ is satisfied, and if it is we know that j_i still attains the minimum in the expression

$$\min_{(i,j) \in \mathcal{A}} \{a_{ij} + p_j\}$$

thereby obviating the calculation of this minimum. If, on the other hand, we have $a_{ij_i} + p_{j_i} > w_i$ (due to an increase of p_{j_i} subsequent to the calculation of w_i), we can check to see whether we still have $w_i = a_{ik_i} + p_{k_i}$; if this is so, then k_i becomes the best neighbor,

$$k_i = \arg \min_{(i,j) \in \mathcal{A}} \{a_{ij} + p_j\}$$

thus, again obviating the calculation of the minimum.

With proper implementation the devices outlined above can typically reduce the number of calculations of the expression $\min_{(i,j) \in \mathcal{A}} \{a_{ij} + p_j\}$ by a factor that is typically in the range from 3 to 5, thereby dramatically reducing the total computation time. Both of the above devices, together with graph reduction, are used in the auction/shortest path codes that are presently the fastest for dense single-origin/all-destination randomly generated problems; see [22].

10.2.1. Parallel implementation. When there is a single destination and multiple origins, several interesting parallel computation possibilities arise. The idea is to maintain a different path P^i for each origin i , and possibly, a reverse path R for the destination. Different paths may be handled by different processors, and price information can be shared by the processors in some way. There are several possible implementations of this idea. We will describe a shared-memory implementation, and we refer to [5], and the MS thesis [46] for discussions of message-passing implementations. For simplicity, we will not consider the possibility of using the reverse path R ; the thesis [46] and the paper [47] discuss parallel two-sided algorithms.

Here, there is a common price vector p stored in the shared memory that is accessible by all processors. For each origin i , there is a path P^i satisfying CS together with p . In a synchronous implementation of the algorithm, an iteration is executed simultaneously for some origins (possibly all origins, depending on the availability of processors). At the end of an iteration, the results corresponding to the different origins are coordinated. To this end, we note that if a node is the terminal node of the path of several origins, the result of the iteration will be the same for all these origins, i.e., a path extension or a path contraction and corresponding price change will occur simultaneously for all these origins. The only potential conflict arises when a node i is the terminal path node for some origin and the path of a different origin is extended by i as a result of the iteration. Then, if p_i is increased due to a path contraction for the former origin, the path extension of the latter origin is canceled. An additional important detail is that an origin i can stop its computation once the terminal node of its path P^i is an origin that has already found its shortest path to the destination. Thus, the processor handling this origin may be diverted to handle the path of another origin.

It is reasonable to speculate that the parallel time to solve the multiple-origins problem is closer to the smallest time over all origins to find a single-origin shortest path, rather than to the longest time. However, this conjecture needs to be tested experimentally on a shared-memory machine ([47] contains some experimental results).

The parallel implementation outlined above is synchronous, that is, all origins iterate simultaneously, and the results are communicated and coordinated at the end of the iteration to the extent necessary for the next iteration. An asynchronous implementation is also possible principally because of the monotonicity of the mapping

$$p_i := \min_{(i,j) \in \mathcal{A}} \{a_{ij} + p_j\}$$

see [23]. We refer to [46] and [47] for a discussion of such an asynchronous implementation.

10.3. Transportation problems

A serial version of the auction/transportation algorithm of Section 6 has been implemented and tested in [14]. The algorithm uses adaptive ϵ -scaling and integer arithmetic; a version using floating-point arithmetic, which would be useful for a large cost range, has not been tested. Computational results show that this auction algorithm is considerably faster than its chief competitors for important classes of transportation problems. Generally, these problems are characterized by two properties, *homogeneity* and *asymmetry*. A homogeneous problem is one for which there are only a few levels of supply and demand. An

asymmetric problem is one for which the number of sources is much larger than the number of sinks. For other types of transportation problems, computational experimentation suggests that the auction algorithm is outperformed by state-of-the-art codes based on relaxation methods; see [14].

We know also of unpublished studies of parallel and asynchronous implementations of the transportation/auction algorithm by D. A. Castañon (unpublished). These studies indicate that with Jacobi parallelization, a modest speedup (of the order of about five) is possible on shared-memory machines.

10.4. *Minimum cost flow problems*

Serial implementations of the ϵ -relaxation method for general minimum cost flow problems are not yet competitive in practice with implementations of other methods, although they are not overwhelmingly worse and further research may change the situation in their favor; see [11]. However, the ϵ -relaxation method is well suited for parallel computation, so in some parallel computing environments it may be faster than its closest competitors. The method admits a totally asynchronous implementation, as shown in the original paper, [9]; see also [20] and [23]. A synchronous massively parallel implementation of the ϵ -relaxation method is presented in [38].

The ϵ -relaxation method can also be implemented in a reverse version where a node price is *decreased*, while flow is “pulled” along incident arcs toward the node rather than “pushed” away from the node. The idea here is fairly similar to the one of reverse auction. In computational tests [11] and [38], a combined forward/reverse version of ϵ -relaxation seems to perform better than the forward version, but the evidence is not conclusive. The forward/reverse version also has greater parallelism potential, since nodes with positive as well as negative surplus can iterate simultaneously (with some safeguards to preserve desirable termination properties). Parallel implementations of this type have not yet been tested.

10.5. *Max-flow problems*

As mentioned in Section 8, the max-flow version of the ϵ -relaxation method can be implemented efficiently without using ϵ -scaling. It is important to understand, however, that intense price wars can still occur, particularly for very sparse problems, despite the small cost range and the polynomial complexity of the method. To alleviate the detrimental effects of price wars, some computational tricks are essential. In particular, it has been observed that for problems where price wars occur, the ϵ -relaxation method finds a minimum cut very quickly and may then spend a great deal of additional time to resolve subsequent price wars. Thus, if one is interested in just a minimum cut or just the value of the maximum

flow, it is worth testing periodically to see whether a minimum cut has already been obtained. One method for doing this is based on breadth-first search and is described in [11] (Exercise 5.4); it is used in a max-flow code given in Appendix A.6 of [11]. A maximum flow can still be obtained after a minimum cut is found by using a variant of the Ford-Fulkerson method, and this can typically be done very quickly. Related procedures that aim at ameliorating the effects of price wars have been discussed in [30] and [41].

It should be noted, however, that despite the practical success of schemes such as the one described above, worst-case max-flow examples have been constructed where the unscaled version of the ϵ -relaxation method can take a very large number of iterations (proportional to N^2); see [23], p. 387. These examples are quite artificial, and it is unknown whether the method can exhibit such behavior for some important types of practical problems.

11. Conclusions

Much progress has been made in the last few years to extend and apply auction algorithms to a variety of linear network flow problems, and to place them on an equal footing with the classical primal cost and dual cost improvement methods. Still, auction algorithms are new and not yet fully developed. With further research, they should become more broadly applicable and more competitive with the classical methods.

The auction algorithms discussed in this paper have been implemented in computer codes. Several of these codes appear in the author's textbook [11]; their latest versions are available from the author.

Appendix 1. ϵ -CS, Primal Optimality, and Dual Optimality

Assignment problems

Let us fix $\epsilon \geq 0$. We show that given a feasible assignment $\{(i, j_i) \mid i = 1, \dots, n\}$ and a set of prices $\{\bar{p}_j \mid j = 1, \dots, n\}$, which satisfy ϵ -CS (if $\epsilon > 0$) or complementary slackness (if $\epsilon = 0$), the assignment is within $n\epsilon$ of maximizing the total benefit, and is optimal if $\epsilon=0$. Furthermore, the set of prices is within $n\epsilon$ of minimizing a certain dual cost function.

Let $\epsilon > 0$. We first note that the total benefit of any feasible assignment $\{(i, k_i) \mid i = 1, \dots, n\}$ satisfies

$$\sum_{i=1}^n a_{ik_i} \leq \sum_{j=1}^n p_j + \sum_{i=1}^n \max_j \{a_{ij} - p_j\}$$

for any set of prices $\{p_j \mid j = 1, \dots, n\}$, since the second term of the right-hand side is no less than

$$\sum_{i=1}^n (a_{ik_i} - p_{k_i})$$

while the first term is equal to $\sum_{i=1}^n p_{k_i}$. Therefore

$$A^* \leq D^*$$

where A^* is the optimal total assignment benefit

$$A^* = \max_{\substack{k_i \in A(i), i = 1, \dots, n \\ k_i \neq k_m \text{ if } i \neq m}} \sum_{i=1}^n a_{ik_i}$$

and

$$D^* = \min_{p_j} \left\{ \sum_{j=1}^n p_j + \sum_{i=1}^n \max_{j \in A(i)} \{a_{ij} - p_j\} \right\}$$

On the other hand, since the given assignment $\{(i, j_i) \mid i = 1, \dots, n\}$ satisfies ϵ -CS together with the given set of prices $\{\bar{p}_j \mid j = 1, \dots, n\}$, we have

$$\max_{j \in A(i)} \{a_{ij} - \bar{p}_j\} - \epsilon \leq a_{ij_i} - \bar{p}_{j_i}$$

and by adding this relation over all i , we see that

$$D^* \leq \sum_{i=1}^n \left(\bar{p}_{j_i} + \max_{j \in A(i)} \{a_{ij} - \bar{p}_j\} \right) \leq \sum_{i=1}^n a_{ij_i} + n\epsilon \leq A^* + n\epsilon$$

Since we showed earlier that $A^* \leq D^*$, it follows that the total assignment benefit $\sum_{i=1}^n a_{ij_i}$ is within $n\epsilon$ of the optimal value A^* .

Note that the function

$$\sum_{j=1}^n p_j + \sum_{i=1}^n \max_{j \in A(i)} \{a_{ij} - p_j\}$$

appearing in the definition of D^* may be viewed as a *dual function* of the price variables p_j , and its minimization may be viewed as a *dual problem* in the standard linear programming duality context; see [11], [23], [44], and [48]. It is seen from the preceding analysis that the prices \hat{p}_j attain within $n\epsilon$ the dual optimal value D^* .

If we let $\epsilon=0$ in the preceding argument, we see that $A^* = D^*$ and that an assignment and a set of prices that are at equilibrium, maximize the total benefit and minimize the dual function, respectively.

Minimum cost flow problems

We now consider the minimum cost flow problem and we prove the following generalized version of Proposition 6, which holds even if the problem data are not integer.

PROPOSITION 10. *Let the flow-price pair (x, p) satisfy ϵ -CS and suppose that x is feasible. Then x is optimal for the minimum cost flow problem provided*

$$\epsilon < \min_{\text{All simple cycles } Y} \left\{ -\frac{\text{cost of } Y}{\text{number of arcs of } Y} \mid \text{cost of } Y < 0 \right\}$$

where

$$\text{cost of } Y = \sum_{(i,j) \in Y^+} a_{ij} - \sum_{(i,j) \in Y^-} a_{ij}$$

and Y^+ and Y^- are the sets of forward and backward arcs of Y , respectively. In particular, x is optimal if the problem data are integer and $\epsilon < 1/N$.

Proof. If x is not optimal, then (see e.g., [11], p. 24) there exists a simple cycle Y that has negative cost, i.e.

$$\sum_{(i,j) \in Y^+} a_{ij} - \sum_{(i,j) \in Y^-} a_{ij} < 0 \tag{45}$$

and is unblocked with respect to x , i.e.

$$x_{ij} < c_{ij}, \quad \forall (i, j) \in Y^+$$

$$b_{ij} < x_{ij}, \quad \forall (i, j) \in Y^-$$

By ϵ -CS [cf. (32)], the preceding relations imply that

$$p_i \leq p_j + a_{ij} + \epsilon, \quad \forall (i, j) \in Y^+$$

$$p_j \leq p_i - a_{ij} + \epsilon, \quad \forall (i, j) \in Y^-$$

By adding these relations over all arcs of Y (whose number is no more than N), and by using the hypothesis $\epsilon < 1/N$ we obtain

$$\sum_{(i,j) \in Y^+} a_{ij} - \sum_{(i,j) \in Y^-} a_{ij} \geq -N\epsilon > -1$$

Since the arc costs a_{ij} are integer, we obtain a contradiction of (45). **Q.E.D.**

Appendix 2. Finite Termination of the Auction Algorithm

In this appendix we show that for a feasible problem and for any positive value of ϵ , the auction algorithm terminates with a feasible assignment that is within $n\epsilon$ of being optimal (and is optimal if the problem data are integer and $\epsilon < 1/n$). This shows in particular Proposition 2.

The proof relies on the following facts:

1. Once an object is assigned, it remains assigned throughout the remainder of the algorithm's duration. Furthermore, except at termination, there will always exist at least one object that has never been assigned, and has a price equal to its initial price. The reason is that a bidding and assignment phase can result in a reassignment of an already assigned object to a different person, but cannot result in the object becoming unassigned.
2. Each time an object receives a bid, its price increases by at least ϵ [see (9)]. Therefore, if the object receives a bid an infinite number of times, its price increases to ∞ .
3. Every $|A(i)|$ bids by person i , where $|A(i)|$ is the number of objects in the set $A(i)$, the best object value v_i defined by

$$v_i = \max_{j \in A(i)} \{a_{ij} - p_j\} \tag{46}$$

decreases by at least ϵ . The reason is that a bid by person i either decreases v_i by at least ϵ , or else leaves v_i unchanged because there is more than one object j attaining the maximum in (46). However, in the latter case, the price of the object j_i receiving the bid will increase by at least ϵ , and object j_i will not receive another bid by person i until v_i decreases by at least ϵ . The conclusion is that if a person i bids an infinite number of times, v_i must decrease to $-\infty$.

We now argue by contradiction. If termination did not occur, the subset J^∞ of objects that received an infinite number of bids is nonempty. Also, the subset of persons I^∞ that bid an infinite number of times is nonempty. As argued in 2 above, the prices of the objects in J^∞ must tend to ∞ , while as argued in 3 above, the scalars $v_i = \max_{j \in A(i)} \{a_{ij} - p_j\}$ must decrease to $-\infty$ for all persons $i \in I^\infty$. Therefore, $a_{ij} - p_j$ tends to $-\infty$ for all $j \in A(i)$, implying that

$$A(i) \subset J^\infty, \quad \forall i \in I^\infty \quad (47)$$

The ϵ -CS condition (8) states that $a_{ij} - p_j \geq v_i - \epsilon$ for every assigned pair (i, j) , so after a finite number of iterations, each object in J^∞ can only be assigned to a person from I^∞ . Since after a finite number of iterations at least one person from I^∞ will be unassigned at the start of each iteration, it follows that the number of persons in I^∞ is strictly larger than the number of objects in J^∞ . This contradicts the existence of a feasible assignment, since by (47), persons in I^∞ can only be assigned to objects in J^∞ . Therefore, the algorithm must terminate. The feasible assignment obtained upon termination satisfies ϵ -CS (since the algorithm preserves ϵ -CS throughout), so by Proposition 1, this assignment is within $n\epsilon$ of being optimal.

Acknowledgments

This research was supported by the NSF under grants, DDM-8903385 and CCR-9103804, and by the ARO under grant DAAL 03-86-K-D171

References

1. R.K. Ahuja, T.L. Magnanti, and J.B. Orlin, "Network flows," Sloan W.P. No. 2059-88. M.I.T., Cambridge, MA, 1989 (also in *Handbooks in Operations Research and Management Science Vol. 1. Optimization* (G.L. Nemhauser, A.H.G. Rinnooy-Kan, and M.J. Todd (eds.), North-Holland: Amsterdam, 1989, pp. 211-369).
2. R.K. Ahuja and J.B. Orlin, "A fast and simple algorithm for the maximum flow problem," working paper, M.I.T., Cambridge, MA, 1986 (also in *Oper. Res.*, vol. 37, 1989, pp. 748-759).
3. J.A. Ajevedo, M.E.S. Costa, J.J.S. Madeira, and E.V. Martins, "An algorithm for the ranking of shortest paths," working paper, Universidade de Coimbra, Coimbra, Portugal, 1991.
4. J.A. Ajevedo, J.J.S. Madeira, E.V. Martins, and F.M. Pires, "A computational improvement for a shortest paths ranking algorithm," working paper, Universidade de Coimbra, Coimbra, Portugal, 1991.
5. D.P. Bertsekas, "The auction algorithm for shortest paths," *SIAM J. on Optimization*, vol. 1, 1991, pp. 425-447.
6. D.P. Bertsekas, "The auction algorithm: A distributed relaxation method for the assignment problem," *Annals of Oper. Res.*, vol. 14, 1988, pp. 105-123.
7. D.P. Bertsekas, "A distributed algorithm for the assignment problem," working paper, Lab. for Information and Decision Systems, M.I.T., Cambridge, MA, March 1979.

8. D.P. Bertsekas, "A distributed asynchronous relaxation algorithm for the assignment problem," in *Proc. 24th IEEE Conf. Dec. and Contr.*, 1985, pp. 1703–1704.
9. D.P. Bertsekas, "Distributed asynchronous relaxation methods for linear network flow problems," Lab. for Information and Decision Systems, M.I.T., Report P-1606, Cambridge, MA, November 1986.
10. D.P. Bertsekas, "Distributed relaxation methods for linear network flow problems," in *Proc. of 25th IEEE Conf. Dec. and Contr.*, 1986, pp. 2101–2106.
11. D.P. Bertsekas, *Linear Network Optimization: Algorithms and Codes*, M.I.T. Press, Cambridge, MA, 1991.
12. D.P. Bertsekas, "A new algorithm for the assignment problem," *Math. Programming*, vol. 21, 1981, pp. 152–171.
13. D.P. Bertsekas and D.A. Castañon, "The auction algorithm for minimum cost network flow Problem," Laboratory for Information and Decision Systems, M.I.T., Cambridge, MA, November, Report LIDS-P-1925, 1989.
14. D.P. Bertsekas and D.A. Castañon, "The auction algorithm for transportation problems," *Annals of Oper. Res.*, vol. 20, pp. 67–96.
15. D.P. Bertsekas and D.A. Castañon, "A forward/reverse auction algorithm for the asymmetric assignment problem," Alphatech Inc. Report, Burlington, MA, April, 1992, *J. of Computational Optimization and its Applications* (to appear).
16. D.P. Bertsekas and D.A. Castañon, "A generic auction algorithm for the minimum cost network flow problem," Alphatech Inc. Report, Burlington, MA, Sept., 1991.
17. D.P. Bertsekas and D.A. Castañon, "Parallel synchronous and asynchronous implementations of the auction algorithm," Alphatech Inc. Report, Burlington, MA, Nov. 1989, (also in *Parallel Computing*, vol. 17, 1991, pp. 707–732).
18. D.P. Bertsekas, D.A. Castañon, and H. Tsaknakis, "Reverse auction and the solution of inequality constrained assignment problems," unpublished report, 1991, *SIAM J. on Optimization* (to appear).
19. D.P. Bertsekas and J. Eckstein, "Distributed asynchronous relaxation methods for linear network flow problems," *Proc. of IFAC '87*, Munich, Germany, July, 1987.
20. D.P. Bertsekas and J. Eckstein, "Dual coordinate step methods for linear network flow problems," *Math. Progr., Series B*, vol. 42, 1988, pp. 203–243.
21. D.P. Bertsekas and S.K. Mitter, "Descent numerical methods for optimization problems with nondifferentiable cost functions," *SIAM J. on Control*, vol. 11, pp. 637–652.
22. D.P. Bertsekas, S. Pallottino, and M.G. Scutella, "Polynomial auction algorithms for shortest paths," *Lab. for Information and Decision Systems*, Report p. 2107, M.I.T, 1992, submitted for publication.
23. D.P. Bertsekas and J.N. Tsitsiklis, *Parallel and Distributed Computation: Numerical Methods*, Prentice-Hall: Englewood Cliffs, NJ, 1989.
24. D.A. Castañon, "Reverse auction algorithms for assignment problems," DIMACS Series in Discrete Mathematics and Theoretical Computer Science, 1992.
25. D. Castañon, B. Smith, and A. Wilson, "Performance of parallel assignment algorithms on different multiprocessor architectures," Alphatech Inc. Report, Burlington, MA, 1989.
26. J. Cheriyan and S.N. Maheshwari, "Analysis of preflow push algorithms for maximum network flow," *SIAM J. Comput.*, vol. 18, 1989, pp. 1057–1086.
27. G.B. Dantzig, *Linear Programming and Extensions*, Princeton Univ. Press: Princeton, NJ, 1963.
28. S.E. Dreyfus, "An appraisal of some shortest path algorithms," *Oper. Res.*, vol. 17, 1969, pp. 395–412.
29. L.R. Ford, Jr. and D.R. Fulkerson, *Flows in Networks*, Princeton Univ. Press: Princeton, NJ, 1962.
30. A.V. Goldberg, "Efficient graph algorithms for sequential and parallel computers," Laboratory for Computer Science, M.I.T., Cambridge, MA., Tech. Report TR-374, 1987.
31. A.V. Goldberg, "A new max-flow algorithm," Laboratory for Computer Science, M.I.T., Cambridge, MA, Tech. Mem. MIT/LCS/TM-291, 1985.

32. A.V. Goldberg and R.E. Tarjan, "A new approach to the maximum flow problem," in *Proc. 18th ACM STOC*, 1986, pp. 136–146.
33. A.V. Goldberg and R.E. Tarjan, "Solving minimum cost flow problems by successive approximation," *Math. of Oper. Res.*, vol. 15, 1990, pp. 430–466.
34. R. Jonker and A. Volegnant, "A shortest augmenting path algorithm for dense and sparse linear assignment problems," *Computing*, vol. 38, 1987, pp. 325–340.
35. D. Kempa, J. Kennington, and H. Zaki, "Performance characteristics of the Jacobi and Gauss-Seidel versions of the auction algorithm on the alliant FX/8," Dept. of Mech. and Ind. Eng., Univ. of Illinois, Champaign-Urbana, IL, Report OR-89-008, 1989.
36. H.W. Kuhn, "The Hungarian method for the assignment problem," *Nav. Res. Log. Q.*, vol. 2, 1955, pp. 83–97.
37. E. Lawler, *Combinatorial Optimization: Networks and Matroids*, Holt, Reinhart, and Winston: New York, 1976.
38. X. Li and S.A. Zenios, "Data parallel solutions of min-cost network flow problems using ϵ -relaxations," Dept. of Decision Sciences, The Wharton School, Univ. of Pennsylvania, Philadelphia, PA, Report 1991-05-20, 1991.
39. D.G. Luenberger, *Linear and Nonlinear Programming*, Addison-Wesley: Reading, MA, 1984.
40. E.V. Martins, "An algorithm for ranking paths that may contain cycles," *European J. of Oper. Res.*, vol. 18, 1984, pp. 123–130.
41. G. Mazzoni, S. Pallotino, and M.G. Scutella', "The maximum flow problem: A max-preflow approach," *European J. of Oper. Res.*, vol. 53, 1991.
42. G.J. Minty, "A Comment on the Shortest Route Problem," *Operations Research*, vol. 5, p. 724, 1957.
43. S. Pallotino and M.G. Scutella', "Strongly polynomial algorithms for shortest paths," Dipartimento di Informatica Report TR-19/91, University of Pisa, Italy, 1991.
44. C.H. Papadimitriou and K. Steiglitz, *Combinatorial Optimization: Algorithms and Complexity*. Prentice-Hall: Englewood Cliffs, NJ, 1982.
45. C. Phillips and S.A. Zenios, "Experiences with large scale network optimization on the Connection Machine," Dept. of Decision Sciences, The Wharton School, Univ. of Pennsylvania, Philadelphia, PA, Report 88-11-05, Nov. 1988.
46. L. Polymenakos, "Analysis of parallel asynchronous schemes for the auction shortest path algorithm," MS thesis, EECS Dept., M.I.T., Cambridge, MA., Jan. 1991.
47. L. Polymenakos and D.P. Bertsekas, "Parallel shortest path auction algorithms," Lab. for Information and Decision Systems, Unpublished Report, M.I.T., Cambridge, MA, April 1992.
48. R.T. Rockafellar, *Network Flows and Monotropic Programming*, Wiley-Interscience: New York, 1984.
49. B.L. Schwartz, "A computational analysis of the auction algorithm," unpublished manuscript.
50. J. Wein and S.A. Zenios, "Massively parallel auction algorithms for the assignment problem," *Proc. of 3rd Symposium on the Frontiers of Massively Parallel Computation*, MD., 1990.
51. J. Wein and S.A. Zenios, "On the massively parallel solution of the assignment problem," *J. of Parallel and Distributed Computing*, vol. 13, 1991, pp. 228–236.
52. H. Zaki, "A comparison of two algorithms for the assignment problem," Dept. of Mechanical and Industrial Engineering, Univ. of Illinois, Urbana, IL., Report ORL 90-002, 1990.