

AN ϵ -RELAXATION METHOD FOR SEPARABLE CONVEX COST NETWORK FLOW PROBLEMS*

DIMITRI P. BERTSEKAS[†], LAZAROS C. POLYMENAKOS[†], AND PAUL TSENG[‡]

Abstract. We propose a new method for the solution of the single commodity, separable convex cost network flow problem. The method generalizes the ϵ -relaxation method developed for linear cost problems and reduces to that method when applied to linear cost problems. We show that the method terminates with a near optimal solution, and we provide an associated complexity analysis. We also present computational results showing that the method is much faster than earlier relaxation methods, particularly for ill-conditioned problems.

Key words. network optimization, ϵ -relaxation, network flows

AMS subject classifications. 90C35, 90C25

PII. S1052623495285886

1. Introduction. We consider a directed graph with node set $\mathcal{N} = \{1, \dots, N\}$ and arc set $\mathcal{A} \subset \mathcal{N} \times \mathcal{N}$. The number of nodes is N and the number of arcs is denoted by A . We denote by x_{ij} the flow of the arc (i, j) , and we refer to the vector $x = \{x_{ij} \mid (i, j) \in \mathcal{A}\}$ as the flow vector. The separable convex cost network flow problem is

$$\begin{aligned} \text{(P)} \quad & \text{minimize} && \sum_{(i,j) \in \mathcal{A}} f_{ij}(x_{ij}) \\ \text{(1)} \quad & \text{subject to} && \sum_{\{j \mid (i,j) \in \mathcal{A}\}} x_{ij} - \sum_{\{j \mid (j,i) \in \mathcal{A}\}} x_{ji} = s_i \quad \forall i \in \mathcal{N}, \end{aligned}$$

where s_i are given scalars and $f_{ij} : \mathfrak{R} \rightarrow (-\infty, \infty]$ are given convex, closed, proper functions (extended real valued, lower semicontinuous, not identically taking the value ∞). We refer to problem (P) as the *primal* problem. We have implicitly assumed that there exists at most one arc in each direction between any pair of nodes, but this assumption is made for notational convenience and easily can be dispensed with. A flow vector x with $f_{ij}(x_{ij}) < \infty$ for all $(i, j) \in \mathcal{A}$, which satisfies the conservation of flow constraint (1), is called *feasible*. For a given flow vector x , the *surplus* of node i is defined as the difference between the supply s_i and the net outflow from i :

$$\text{(2)} \quad g_i = s_i + \sum_{\{j \mid (j,i) \in \mathcal{A}\}} x_{ji} - \sum_{\{j \mid (i,j) \in \mathcal{A}\}} x_{ij}.$$

We will assume that *there exists at least one feasible flow vector x such that*

$$\text{(3)} \quad f_{ij}^-(x_{ij}) < \infty \quad \text{and} \quad f_{ij}^+(x_{ij}) > -\infty \quad \forall (i, j) \in \mathcal{A},$$

where $f_{ij}^-(x_{ij})$ and $f_{ij}^+(x_{ij})$ denote the left and right directional derivatives of f_{ij} at x_{ij} [Roc84, p. 329].

*Received by the editors May 8, 1995; accepted for publication (in revised form) January 18, 1996. This research was supported by the NSF under grants CCR-9103804 and 9300494-DMI.

<http://www.siam.org/journals/siopt/7-3/28588.html>

[†]Department of Electrical Engineering and Computer Science, M.I.T., Rm. 35-210, Cambridge, MA 02139 (dimitrib@mit.edu, lcpolyme@lids.mit.edu).

[‡]Department of Mathematics, University of Washington, Seattle, WA 98195 (tseng@math.washington.edu).

There is a well-known duality framework for this problem, primarily developed by Rockafellar [Roc70] and discussed in several texts; see, e.g., [Roc84], [BeT89]. This framework involves a Lagrange multiplier p_i for the i th conservation of flow constraint (1). We refer to p_i as the *price* of node i and to the vector $p = \{p_i \mid i \in \mathcal{N}\}$ as the *price vector*. The *dual* problem is as follows:

$$(D) \quad \begin{array}{ll} \text{minimize} & q(p) \\ \text{subject to} & \text{no constraint on } p, \end{array}$$

where the dual functional q is given by

$$q(p) = \sum_{(i,j) \in \mathcal{A}} q_{ij}(p_i - p_j) - \sum_{i \in \mathcal{N}} s_i p_i,$$

and q_{ij} is related to f_{ij} by the conjugacy relation

$$q_{ij}(t_{ij}) = \sup_{x_{ij} \in \mathfrak{R}} \{x_{ij} t_{ij} - f_{ij}(x_{ij})\}.$$

We will assume throughout that f_{ij} is such that q_{ij} is real valued for all $(i, j) \in \mathcal{A}$. This is true, for example, if each function f_{ij} takes the value ∞ outside some compact interval.

It is known (see [Roc84, p. 360]) that under our assumptions, both the primal problem (P) and the dual problem (D) have optimal solutions, and their optimal costs are the negatives of each other. The standard optimality conditions for a feasible flow-price vector pair (x, p) to be primal and dual optimal are

$$f_{ij}^-(x_{ij}) \leq p_i - p_j \leq f_{ij}^+(x_{ij}) \quad \forall (i, j) \in \mathcal{A}.$$

These, known as the *complementary slackness conditions* (CS conditions), may be represented explicitly as

$$(x_{ij}, p_i - p_j) \in \Gamma_{ij} \quad \forall (i, j) \in \mathcal{A},$$

where

$$\Gamma_{ij} = \{(x_{ij}, t_{ij}) \in \mathfrak{R}^2 \mid f_{ij}^-(x_{ij}) \leq t_{ij} \leq f_{ij}^+(x_{ij})\}$$

is the *characteristic curve* associated with arc (i, j) as shown in Fig. 1.

The traditional methods for solving the problem of this paper for the case of linear arc cost functions (when each f_{ij} is linear on some closed interval and is ∞ outside the interval) are primal and dual cost improvement methods, which iteratively improve the primal or the dual cost function. Recently, methods based on the auction approach have gained attention, following the original proposal of the auction algorithm for the assignment problem [Ber79] and the ϵ -relaxation method [Ber86a], [Ber86b]. These methods may not improve the primal or the dual cost at any iteration, and they are based on a relaxed version of the CS conditions, called ϵ -complementary slackness (ϵ -CS). Their worst-case computational complexity, when properly implemented, is excellent; see [Gol87] (see also [BeE88], [BeT89], [GoT90]). Their practical performance is also very good, particularly for special classes of problems such as assignment and max-flow. Furthermore, these methods are well suited for parallel implementation (see [BCE95], [LiZ91], [NiZ93]). We will focus on extending one such method, the ϵ -relaxation method, to the general convex cost case.




FIG. 1. A cost function f_{ij} and its corresponding characteristic curve.

One possibility for dealing with the convex cost case is to use efficient ways to reduce the problem to an essentially linear cost problem by piecewise linearization of the arc cost functions; see [Mey79], [KaM84], [Roc84]. Another possibility is to use differentiable unconstrained optimization methods, such as coordinate descent [BHT87], conjugate gradient [Ven91], and adaptations of other nonlinear programming methods [HaH93], [Hag92] or fixed point methods [BeE87], [TBT90]. However, these methods require that the dual cost function be differentiable, which essentially amounts to the primal cost functions being strictly convex. A more general alternative, which applies to nondifferentiable dual cost functions as well, is to use an extension of the primal or dual cost improvement methods developed for the linear cost case. In particular, there have been proposals of primal cost improvement methods in [Wei74] and more recently in [KaM93]. There have also been proposals of dual cost improvement methods: see the fortified descent method [Roc84] that extends the primal–dual method of Ford and Fulkerson [FoF62] and the relaxation method of [BHT87] that extends the corresponding linear cost relaxation method of [Ber85] and [BeT88]. These methods, together with the price vector, maintain a flow vector that satisfies the ϵ -CS conditions and progressively work towards primal feasibility. The flow vector becomes feasible at termination.

In this paper we develop and analyze the first extension of an auction method, the ϵ -relaxation method, to the convex arc cost case.¹ (An analogous extension of the auction/sequential shortest path method given in [Ber92], which has also been incorporated in the latest release of the RELAX code [BeT94], is developed in the Ph.D. thesis of the second author [Pol95].) Our method is based on the ϵ -CS conditions first introduced in [BHT87] for the case of convex arc costs. In particular, we say that the flow vector x and the price vector p satisfy ϵ -CS if and only if (see Fig. 2)

$$(4) \quad f_{ij}(x_{ij}) < \infty \quad \text{and} \quad f_{ij}^-(x_{ij}) - \epsilon \leq p_i - p_j \leq f_{ij}^+(x_{ij}) + \epsilon \quad \forall (i, j) \in \mathcal{A}.$$

¹We have learned that the same method was independently developed and analyzed by De Leone, Meyer, and Zakarian [DMZ95]. The results of their computational tests qualitatively agree with ours.




FIG. 2. A visualization of the ϵ -CS conditions as a cylinder around the characteristic curve (bold line). The shaded area represents flow-price differential pairs that satisfy the ϵ -CS conditions.

It was shown in [BHT87] that if a feasible flow-price vector pair (x, p) satisfies ϵ -CS, then x and p are primal and dual optimal, respectively, within a factor that is proportional to ϵ (see Proposition 6). Our method is similar to the ϵ -relaxation method for linear cost network flow problems. It iteratively modifies the price vector while effecting attendant flow changes that maintain the ϵ -CS conditions. The method terminates with a feasible flow-price vector pair which, however, satisfies ϵ -CS rather than CS. There is a fundamental difference from the other dual descent methods for nondifferentiable dual cost problems: the price changes are made exclusively along coordinate directions, that is, one price at a time, and a price change need not improve the dual cost. However, because the flow-price vector pair (x, p) maintained by the ϵ -relaxation method satisfies ϵ -CS rather than CS, there is more freedom in adjusting the flow-price vector pair towards feasibility, even though the pair obtained when the method terminates is optimal only within a factor proportional to ϵ .

The method of this paper essentially provides a mechanism to move around the ϵ -CS diagram of Fig. 2 while changing one price at a time and works towards primal feasibility. There is a variety of mechanisms for effecting such price changes and Fig. 3 illustrates some of the possibilities. In particular, by starting from a point on the characteristic curve of arc (i, j) we can follow any direction around that point and change the price p_i or the price p_j and/or the flow x_{ij} simultaneously until $(x_{ij}, p_i - p_j)$ is either on the characteristic curve or is within a distance of ϵ above or below the characteristic curve of arc (i, j) . For example, if node i has positive surplus, by increasing the flow of an outgoing arc (i, j) or by decreasing the flow of an incoming arc (j, i) the surplus of i will be decreased, while the surplus of j will be increased by an equal amount. This is the basic mechanism for moving flow from nodes of positive surplus to nodes of negative surplus, thus working towards primal feasibility. It is possible, however, that node i has positive surplus, while the flow of none of the outgoing arcs (i, j) can be increased and the flow of none of the incoming arcs (j, i) can be decreased without violating the ϵ -CS conditions. In this case, the method increases the price of node i in order to “make room” in the ϵ -CS diagram




FIG. 3. Starting from any point on the characteristic curve (dark points) of arc (i, j) , a new point on the characteristic curve can be obtained in a variety of ways. The figure depicts a few such examples where the flow and price differential for arc (i, j) are changed simultaneously according to some linear relation.

for a subsequent flow change.

The paper is organized as follows. In section 2 we present the ϵ -relaxation method extended to solve convex cost problems. In section 3, we show that this method terminates with a near optimal flow-price vector pair, and in section 4 we provide a complexity analysis for the method. In section 5, we describe a version of this method that uses both price increase and decrease steps, and in section 6 we report our computational experience with the methods of sections 2 and 5 on some convex linear/quadratic cost problems. Our test results show that, on problems where some (possibly all) arcs have strictly convex cost, the new method outperforms, often by an impressive margin, earlier relaxation methods. Significantly, our method seems to be minimally affected by ill-conditioning in the dual problem. We do not know of any other method for which this is true.

2. The ϵ -relaxation method. For a flow-price vector pair (x, p) satisfying ϵ -CS, we define for each node $i \in \mathcal{N}$ its *push list* as the set of arcs

$$(5) \quad \{(i, j) \mid \epsilon/2 < p_i - p_j - f_{ij}^+(x_{ij}) \leq \epsilon\} \cup \{(j, i) \mid -\epsilon \leq p_j - p_i - f_{ji}^-(x_{ji}) < -\epsilon/2\}.$$

Figure 4 illustrates when an arc (i, j) is in the push list of i and when it is in the push list of j . We note that a more general definition of the push list can be given by replacing the term $\epsilon/2$ with $\beta\epsilon$, where β is a scalar with $0 < \beta < 1$. The subsequent analysis applies, with minor modifications, to the corresponding version of the ϵ -relaxation method to be given shortly.

An arc (i, j) [or (j, i)] in the push list of i is said to be *unblocked* if there exists a $\delta > 0$ such that

$$p_i - p_j \geq f_{ij}^+(x_{ij} + \delta)$$

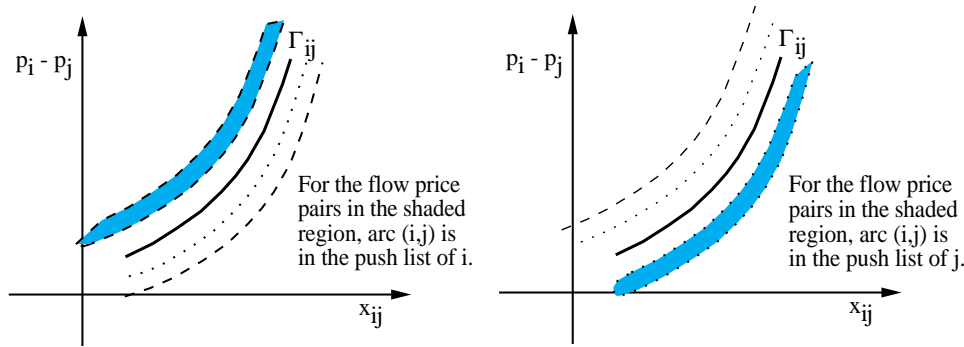


FIG. 4. A visualization of the conditions satisfied by a push list arc. The shaded area represents flow-price differential pairs corresponding to a push list arc.




FIG. 5. The flow margin of an unblocked push list arc.

(or $p_j - p_i \leq f_{ji}^-(x_{ji} - \delta)$, respectively). For an unblocked push list arc, the supremum of δ for which the above relation holds is called the *flow margin* of the arc. The flow margin of an arc (i, j) is illustrated in Fig. 5. An important property is noted in the following proposition.

PROPOSITION 1. *The arcs in the push list of a node are unblocked.*

Proof. Assume that for an arc $(i, j) \in \mathcal{A}$ we have

$$p_i - p_j < f_{ij}^+(x_{ij} + \delta) \quad \forall \delta > 0.$$

Since the function f_{ij}^+ is right continuous, this yields

$$p_i - p_j \leq \lim_{\delta \downarrow 0} f_{ij}^+(x_{ij} + \delta) = f_{ij}^+(x_{ij}),$$

and thus (i, j) cannot be in the push list of node i . A similar argument proves that an arc $(j, i) \in \mathcal{A}$ such that

$$p_j - p_i > f_{ji}^-(x_{ji} - \delta) \quad \forall \delta > 0$$

cannot be in the push list of node i . \square

For a given pair (x, p) satisfying ϵ -CS, consider an arc set \mathcal{A}^* that contains all push list arcs oriented in the direction of flow change. In particular, for each arc (i, j)

in the push list of a node i we introduce an arc (i, j) in \mathcal{A}^* , and for each arc (j, i) in the push list of node i we introduce an arc (i, j) in \mathcal{A}^* . The set of nodes \mathcal{N} and the set \mathcal{A}^* define the *admissible graph* $G^* = (\mathcal{N}, \mathcal{A}^*)$. Note that an arc can be in the push list of at most one node, so the admissible graph is well defined.

The ϵ -relaxation method starts with a flow-price vector pair (x, p) satisfying ϵ -CS and is such that the corresponding admissible graph G^* is acyclic. One possibility is to select an initial price vector p^0 and to set the initial arc flow for every arc $(i, j) \in \mathcal{A}$ to

$$(6) \quad x_{ij} = \sup\{\xi \mid f_{ij}^+(\xi) \leq p_i^0 - p_j^0 - \epsilon/2\}.$$

It can be seen with this choice that ϵ -CS is satisfied for every arc $(i, j) \in \mathcal{A}$ and that the initial admissible graph is empty and thus acyclic.

In the typical iteration of the method, we select a node i with positive surplus, and we perform one or more of the following two operations:

- (a) a *price rise* on node i , which increases the price p_i by the maximum amount that maintains ϵ -CS, while leaving all arc flows unchanged,
- (b) a *flow push* (also called a δ -*flow push*) along an arc (i, j) (or along an arc (j, i)), which increases (i, j) (or decreases (j, i)) by an amount $\delta \in (0, g_i]$, while leaving all node prices unchanged.

The iteration is as follows.

TYPICAL ITERATION OF THE ϵ -RELAXATION METHOD.

Step 1: Select a node i with positive surplus g_i ; if no such node exists, terminate the method.

Step 2: If the push list of i is empty, go to Step 3. Otherwise, choose an arc from the push list of i and perform a δ -flow push towards the opposite node j , where

$$\delta = \min\{g_i, \text{flow margin of arc}\}.$$

If the surplus of i becomes zero, go to the next iteration; otherwise go to Step 2.

Step 3: Increase the price p_i by the maximum amount that maintains ϵ -CS. Go to the next iteration.

We make the following observations about the ϵ -relaxation method:

1. The method preserves ϵ -CS and the prices are monotonically nondecreasing. This is evident from the initialization and Step 3 of the method.
2. Once the surplus of a node becomes nonnegative, it remains nonnegative for all subsequent iterations. The reason is that a flow push at a node i cannot make the surplus of i negative (cf. Step 2) and cannot decrease the surplus of neighboring nodes.
3. If at some iteration a node has negative surplus, then its price must be equal to its initial price. This is a consequence of observation 2 above and the fact that price changes occur only on nodes with positive surplus.

3. Termination of the ϵ -relaxation method. To prove the termination of the ϵ -relaxation method of section 2, we first prove that the total number of price rises that the method can perform is bounded.

PROPOSITION 2. *Each price rise increment in the ϵ -relaxation method is at least $\epsilon/2$.*

Proof. We first note that a price rise on a node i occurs only when its push list is empty. Thus, for every arc $(i, j) \in \mathcal{A}$ we have $p_i - p_j - f_{ij}^+(x_{ij}) \leq \epsilon/2$, and for every

arc $(j, i) \in \mathcal{A}$ we have $p_j - p_i - f_{ji}^-(x_{ji}) \geq -\epsilon/2$. This implies that all elements of the sets of positive numbers

$$S^+ = \{p_j - p_i + f_{ij}^+(x_{ij}) + \epsilon \mid (i, j) \in \mathcal{A}\},$$

$$S^- = \{p_j - p_i - f_{ji}^-(x_{ji}) + \epsilon \mid (j, i) \in \mathcal{A}\}$$

are greater than or equal to $\epsilon/2$. Since a price rise at i increases p_i by the increment $\gamma = \min\{S^+ \cup S^-\}$, the result follows. \square

The following proposition bounds the total number of price increases that the ϵ -relaxation method can perform on any node. The proof is patterned after that for the linear cost case [Ber86a], [BeE88].

PROPOSITION 3. *Assume that for some integer $K \geq 1$, the initial price vector p^0 for the ϵ -relaxation method satisfies $K\epsilon$ -CS together with some feasible flow vector x^0 . Then, the ϵ -relaxation method performs at most $2(K+1)(N-1)$ price rises per node.*

Proof. Consider the pair (x, p) at the beginning of an ϵ -relaxation iteration. Since the surplus vector $g = (g_1, \dots, g_N)$ is not zero and the flow vector x^0 is feasible, we conclude that for each node s with $g_s > 0$ there exists a node t with $g_t < 0$ and a path H from t to s that contains no cycles and is such that

$$(7) \quad x_{ij} > x_{ij}^0 \quad \forall (i, j) \in H^+,$$

$$(8) \quad x_{ij} < x_{ij}^0 \quad \forall (i, j) \in H^-,$$

where H^+ is the set of forward arcs of H and H^- is the set of backward arcs of H . (This can be seen from the conformal realization theorem ([Roc84] or [Ber91]) as follows. For the flow vector $x - x^0$, the net outflow from node t is $-g_t > 0$ and the net outflow from node s is $-g_s < 0$ (here we ignore the flow supplies), so, by the conformal realization theorem, there is a path H from t to s that contains no cycle and conforms to the flow $x - x^0$. That is, $x_{ij} - x_{ij}^0 > 0$ for all $(i, j) \in H^+$ and $x_{ij} - x_{ij}^0 < 0$ for all $(i, j) \in H^-$. Equations (7) and (8) then follow.)

From equations (7) and (8) and the convexity of the functions f_{ij} for all $(i, j) \in \mathcal{A}$, we have

$$(9) \quad f_{ij}^-(x_{ij}) \geq f_{ij}^+(x_{ij}^0) \quad \forall (i, j) \in H^+,$$

$$(10) \quad f_{ij}^+(x_{ij}) \leq f_{ij}^-(x_{ij}^0) \quad \forall (i, j) \in H^-.$$

Since the pair (x, p) satisfies ϵ -CS, we also have that

$$(11) \quad p_i - p_j \in [f_{ij}^-(x_{ij}) - \epsilon, f_{ij}^+(x_{ij}) + \epsilon] \quad \forall (i, j) \in \mathcal{A}.$$

Similarly, since the pair (x^0, p^0) satisfies $K\epsilon$ -CS, we have

$$(12) \quad p_i^0 - p_j^0 \in [f_{ij}^-(x_{ij}^0) - K\epsilon, f_{ij}^+(x_{ij}^0) + K\epsilon] \quad \forall (i, j) \in \mathcal{A}.$$

Combining equations (9)–(12), we obtain

$$p_i - p_j \geq p_i^0 - p_j^0 - (K+1)\epsilon \quad \forall (i, j) \in H^+,$$

$$p_i - p_j \leq p_i^0 - p_j^0 + (K + 1)\epsilon \quad \forall (i, j) \in H^-.$$

Applying the above inequalities for all arcs of the path H , we get

$$(13) \quad p_t - p_s \geq p_t^0 - p_s^0 - (K + 1)|H|\epsilon,$$

where $|H|$ denotes the number of arcs of the path H . We observed earlier that if a node has negative surplus at some time, then its price is unchanged from the beginning of the method until that time; thus, $p_t = p_t^0$. Since the path contains no cycles, we also have that $|H| \leq N - 1$. Therefore, equation (13) yields

$$(14) \quad p_s - p_s^0 \leq (K + 1)|H|\epsilon \leq (K + 1)(N - 1)\epsilon.$$

Since only nodes with positive surplus can increase their prices and, by Proposition 2, each price rise increment is at least $\epsilon/2$, we conclude from equation (14) that the total number of price rises that can be performed for node s is at most $2(K+1)(N-1)$. \square

The result of the preceding proposition is remarkable in that the bound on the number of price changes is independent of the cost functions but depends only on

$$K^0 = \min\{K \in \{0, 1, \dots\} \mid (x^0, p^0) \text{ satisfies } K\epsilon\text{-CS for some feasible flow vector } x^0\},$$

which is the minimum multiplicity of ϵ by which CS is violated by the starting price together with some feasible flow vector. This result will be used later to prove a particularly favorable complexity bound for the method. Note that K^0 is well defined for any p^0 because, for all K sufficiently large, $K\epsilon$ -CS is satisfied by p^0 and the feasible flow vector x satisfying equation (3).

In order to show that the number of flow pushes that can be performed between successive price increases is finite, we first prove that the method maintains the acyclicity of the admissible graph.

PROPOSITION 4. *The admissible graph remains acyclic throughout the ϵ -relaxation method.*

Proof. We use induction. Initially, the admissible graph G^* is empty, so it is trivially acyclic. Assume that G^* remains acyclic for all subsequent iterations up to the m th iteration for some m . We will prove that after the m th iteration G^* remains acyclic. Clearly, after a flow push the admissible graph remains acyclic, since it either remains unchanged or some arcs are deleted from it. Thus, we only have to prove that after a price rise at a node i , no cycle involving i is created. We note that after a price rise at node i , all incident arcs to i in the admissible graph at the start of the m th iteration are deleted and new arcs incident to i are added. We claim that i cannot have any incoming arcs which belong to the admissible graph. To see this, note that just before a price rise at node i , we have from (4) that

$$p_j - p_i - f_{ji}^-(x_{ji}) \leq \epsilon \quad \forall (j, i) \in \mathcal{A},$$

and since each price rise is at least $\epsilon/2$, we must have

$$p_j - p_i - f_{ji}^-(x_{ji}) \leq \frac{\epsilon}{2} \quad \forall (j, i) \in \mathcal{A}$$

after the price rise. Then, by equation (5), (j, i) cannot be in the push list of node j . By a similar argument, we have that (i, j) cannot be in the push list of j for all

$(i, j) \in \mathcal{A}$. Thus, after a price increase at i , node i cannot have any incoming incident arcs belonging to the admissible graph, so no cycle involving i can be created. \square

We say that a node i is a *predecessor* of a node j in the admissible graph G^* if a directed path from i to j exists in G^* . Node j is then called a *successor* of i . Observe that flow is pushed towards the successors of a node, and since G^* is acyclic, flow cannot be pushed from a node to any of its predecessors. A δ -flow push along an arc in G^* is said to be *saturating* if δ is equal to the flow margin of the arc. By our choice of δ (see Step 2 of the method), a nonsaturating flow push always exhausts (i.e., sets to zero) the surplus of the starting node of the arc. Thus we have the following proposition.

PROPOSITION 5. *The number of flow pushes between two successive price increases (not necessarily at the same node) performed by the ϵ -relaxation method is finite.*

Proof. We observe that a saturating flow push along an arc removes the arc from the admissible graph, while a nonsaturating flow push does not add a new arc to the admissible graph. Thus, the number of saturating flow pushes that can be performed between successive price increases is at most A . It will thus suffice to show that the number of nonsaturating flow pushes that can be performed between saturating flow pushes is finite. Assume the contrary; that is, there is an infinite sequence of successive nonsaturating flow pushes with no intervening saturating flow push. Then, the surplus of some node i^0 must be exhausted infinitely often during this sequence. This can happen only if the surplus of some predecessor i^1 of i^0 is exhausted infinitely often during the sequence. Continuing in this manner, we construct an infinite succession of predecessor nodes $\{i^k\}$. Thus, some node in this sequence must be repeated, which is a contradiction since the admissible graph is acyclic. \square

By refining the proof of Proposition 5, we can further show that the number of flow pushes between successive price increases is at most $(N + 1)A$, from which a complexity result for the ϵ -relaxation method may be derived. However, we will defer the analysis of complexity to section 4, where an implementation of the method with sharper complexity bound will be presented.

Propositions 3 and 5 prove that the ϵ -relaxation method terminates. Upon termination, we have that the flow-price vector pair satisfies ϵ -CS and that the flow vector is feasible since the surplus of all nodes will be zero. The following proposition, due to [BHT87], shows that the flow vector and the price vector obtained upon termination are primal optimal and dual optimal within a factor that is essentially proportional to ϵ .

PROPOSITION 6. *For each $\epsilon > 0$, let $x(\epsilon)$ and $p(\epsilon)$ denote any flow and price vector pair satisfying ϵ -CS with $x(\epsilon)$ feasible, and let $\xi(\epsilon)$ denote any flow vector satisfying CS together with $p(\epsilon)$ (note that $\xi(\epsilon)$ need not be feasible). Then*

$$0 \leq f(x(\epsilon)) + q(p(\epsilon)) \leq \epsilon \sum_{(i,j) \in \mathcal{A}} |x_{ij}(\epsilon) - \xi_{ij}(\epsilon)|.$$

Furthermore, $f(x(\epsilon)) + q(p(\epsilon)) \rightarrow 0$ as $\epsilon \rightarrow 0$.

Proposition 6 does not give an estimate of how small ϵ has to be in order to achieve a certain degree of optimality. However, in the common case where finiteness of the arc cost functions f_{ij} imply lower and upper bounds on the arc flows, i.e.,

$$-\infty < b_{ij} = \inf_{\xi} \{\xi \mid f_{ij}(\xi) < \infty\} \leq \sup_{\xi} \{\xi \mid f_{ij}(\xi) < \infty\} = c_{ij} < \infty,$$

Proposition 6 together with the fact that $q(p(\epsilon)) \geq -f^*$ yields the estimate

$$0 \leq f(x(\epsilon)) - f^* \leq \epsilon A \max_{(i,j) \in \mathcal{A}} |c_{ij} - b_{ij}|,$$

where f^* is the optimal cost of (P). Similarly, we obtain

$$0 \leq q(p(\epsilon)) - q^* \leq \epsilon A \max_{(i,j) \in \mathcal{A}} |c_{ij} - b_{ij}|,$$

where q^* is the optimal cost of (D).

4. Complexity analysis for the ϵ -relaxation method. We now derive a bound on the running time of the ϵ -relaxation method. Because the cost functions are convex, it is not possible to express the size of the problem in terms of the problem data. To deal with this difficulty, we introduce a set of simple operations performed by the method, and we estimate the number of these operations. In particular, in addition to the usual arithmetic operations with real numbers, we consider the following operations:

- (a) Given the flow x_{ij} of an arc (i, j) , calculate the cost $f_{ij}(x_{ij})$, the left derivative $f_{ij}^-(x_{ij})$, and the right derivative $f_{ij}^+(x_{ij})$.
- (b) Given the price differential t_{ij} of an arc (i, j) , calculate $\sup\{\xi \mid f_{ij}^+(\xi) \leq t_{ij}\}$ and $\inf\{\xi \mid f_{ij}^-(\xi) \geq t_{ij}\}$.

Operation (a) is needed to compute the push list of a node and a price increase increment; operation (b) is needed to compute the flow margin of an arc and the flow initialization of equation (6). We will thus estimate the total number of simple operations performed by the method (see Proposition 8).

To obtain a sharper complexity bound, we introduce an order in which the nodes are chosen in iterations. This rule is based on the *sweep implementation* of the ϵ -relaxation method, which was introduced in [Ber86a] and was analyzed in more detail in [BeE88], [BeT89], and [BeC93] for the linear cost network flow problem. All the nodes are kept in a linked list T , which is traversed from the first to the last element. The order of the nodes in the list is consistent with the successor order implied by the admissible graph; that is, if a node j is a successor of a node i , then j must appear after i in the list. If the initial admissible graph is empty, as is the case with the initialization of equation (6), the initial list is arbitrary. Otherwise, the initial list must be consistent with the successor order of the initial admissible graph. The list is updated in a way that maintains the consistency with the successor order. In particular, let i be a node on which we perform an ϵ -relaxation iteration, and let N_i be the subset of nodes of T that are after i in T . If the price of i changes, then node i is removed from its position in T and placed in the first position of T . The next node chosen for iteration, if N_i is nonempty, is the node $i' \in N_i$ with positive surplus, which ranks highest in T . Otherwise, the positive surplus node ranking highest in T is picked. It can be shown (see the references cited earlier) that with this rule of repositioning nodes following a price change, the list order is consistent with the successor order implied by the admissible graph throughout the method.

A *sweep cycle* is a set of iterations whereby all nodes are chosen once from the list T , and an ϵ -relaxation iteration is performed on those nodes that have positive surplus. The idea of the sweep implementation is that an ϵ -relaxation iteration at a node i that has predecessors with positive surplus may be wasteful, since the surplus of i will be set to zero and become positive again through a flow push at a predecessor node.

Our complexity analysis follows the line of the corresponding analysis for the linear cost problem. First we have a proposition that estimates the number of sweep cycles required for termination.

PROPOSITION 7. *Assume that for some integer $K \geq 1$, the initial price vector p^0 for the sweep implementation of the ϵ -relaxation method satisfies $K\epsilon$ -CS together with some feasible flow vector x^0 . Then, the number of sweep cycles up to termination is $O(KN^2)$.*

Proof. Consider the start of any sweep cycle. Let N^+ be the set of nodes with positive surplus that have no predecessor with positive surplus; let N^0 be the set of nodes with nonpositive surplus that have no predecessor with positive surplus. Then, as long as no price change takes place during the cycle, all nodes in N^0 remain in N^0 , and an iteration on a node $i \in N^+$ moves i from N^+ to N^0 . So if no node changed price during the cycle, then all nodes in N^+ will be moved to N^0 and the method terminates. Therefore, there is a price change in every cycle except possibly the last one. Since by Proposition 3 there are $O(KN^2)$ price changes, the result follows. \square

By using Proposition 7, we now bound the running time for the sweep implementation of the ϵ -relaxation method. The dominant computational requirements are as follows:

- (1) the computation required for price increases,
- (2) the computation required for saturating δ -flow pushes,
- (3) the computation required for nonsaturating δ -flow pushes.

PROPOSITION 8. *Assume that for some $K \geq 1$ the initial price vector p^0 for the sweep implementation of the ϵ -relaxation method satisfies $K\epsilon$ -CS together with some feasible flow vector x^0 . Then, the method requires $O(KN^3)$ operations up to termination.*

Proof. According to Proposition 3, there are $O(KN)$ price increases per node, so the requirements for (1) above are $O(KNA)$ operations. Furthermore, whenever a flow push is saturating, it takes at least one price increase at one of the end nodes before the flow on that arc can be changed again. Thus, the total requirement for (2) above is $O(KNA)$ operations also. Finally, for (3) above we note that for each sweep cycle there can be only one nonsaturating δ -flow push per node. Thus, a time bound for (3) is $O(N \cdot \text{total number of sweep cycles})$, which, by Proposition 7, is $O(KN^3)$ operations. Adding the computational requirements for (1), (2), and (3) and using the fact that $A \leq N^2$, the result follows. \square

It is well known that the theoretical and the practical performance of the ϵ -relaxation method can be improved by scaling. One possibility is *cost scaling* (see [BLJ92], [EdK72], [Roc80]). An analysis of cost scaling applied to ϵ -relaxation for the linear network flow problem is given in [BeE87] and also in [BeE88]. In the convex cost case, however, cost scaling may be difficult to implement since the arc cost functions may be unbounded. A second scaling approach in connection with the ϵ -relaxation method for linear cost problems is ϵ -*scaling*. This approach was originally introduced in [Ber79] as a means of improving the performance of the auction algorithm for the assignment problem. Its complexity analysis was given in [Gol87] and [GoT90].

The key idea of ϵ -scaling is to apply the ϵ -relaxation method several times, starting with a large value of ϵ , and to successively reduce ϵ up to a final value that will give the desirable degree of accuracy to our solution. Furthermore, the price and flow information from one application of the method is transferred to the next.

The procedure is as follows: first, we choose a scalar $\theta \in (0, 1)$, a price vector p^0 , and a desirable value $\bar{\epsilon}$ for ϵ on termination. Next, we choose a sufficiently large ϵ^0 so

that p^0 satisfies ϵ^0 -CS with some feasible flow vector x^0 . Then, for $k = 1, 2, \dots$, we set $\epsilon^k = \theta\epsilon^{k-1}$, and for $k = 1, 2, \dots, \bar{k}$, we apply the ϵ -relaxation method with $\epsilon = \epsilon^{k-1}$, where \bar{k} is the first positive integer k for which ϵ^{k-1} is below $\bar{\epsilon}$. Let (x^k, p^k) be the flow-price vector pair obtained at the k th application of the method for $k = 1, 2, \dots, \bar{k}$. Then, x^k is feasible and satisfies ϵ^{k-1} -CS with p^k . Furthermore, the admissible graph after the k th application of the method is acyclic. The initial price vector for the $(k + 1)$ st application is p^k , and the initial flow is x_{ij}^k for the arcs (i, j) that satisfy ϵ^k -CS with p^k ; otherwise,

$$\sup \{ \xi \mid f_{ij}^+(\xi) \leq p_i^k - p_j^k - \epsilon^k/2 \} .$$

This choice of initial flows does not introduce any new arcs to the admissible graph, so the initial admissible graph for the $(k + 1)$ st application of the method is acyclic. For the 1st application of the method, the initial price vector is p^0 and the initial flow vector is chosen so that the initial admissible graph is acyclic.

We observe that for the $(k + 1)$ st application of the method ($k = 0, 1, \dots, \bar{k} - 1$), the initial price vector p^k satisfies ϵ^k/θ -CS with the feasible flow vector x^k . Thus, based on Proposition 8, we conclude that the $(k + 1)$ st application of the method has a running time of $O(\lceil 1/\theta \rceil N^3)$, which is $O(N^3)$ since θ is a fixed scalar. The method will be applied at most $\bar{k} = \lceil \log_\theta(\epsilon^0/\bar{\epsilon}) \rceil$ times. We have thus obtained the following proposition.

PROPOSITION 9. *The running time of the ϵ -relaxation method using the sweep implementation and ϵ -scaling as described above is $O(N^3 \ln(\epsilon^0/\bar{\epsilon}))$ operations.*

We note that a complexity bound of $O(N A \ln(N) \ln(\epsilon^0/\bar{\epsilon}))$ operations was derived in [KaM93] for the tighten and cancel method. For relatively dense problems where $A = \Theta(N^2/\ln N)$, our complexity bound for the ϵ -relaxation method is more favorable, while for sparse problems, where $A = \Theta(N)$, the reverse is true.

5. The reverse and forward–reverse ϵ -relaxation methods. The ϵ -relaxation method we presented in the previous sections performed iterations only on nodes with positive surplus. We will refer to it as the *forward* method. We can also define a method (namely, the *reverse* method), which performs iterations on nodes of negative surplus. This involves a simple reformulation of the flow and price changing operations we introduced in previous sections for the forward method. The reverse ϵ -relaxation method is the “mirror image” of the forward method that we developed in the previous sections. Naturally, it has similar properties to the forward method and its validity follows from a similar analysis.

It is possible to combine the forward and the reverse methods so that the resulting method will operate on both positive and negative surplus nodes. Our intuition is that if we perform ϵ -relaxation iterations on both sources and sinks, we will be able to find the optimal solution faster for certain classes of problems. We refer to the resulting method as the *forward–reverse method*. We initialize the arc flows and node prices in the same way we initialized them for the forward and the reverse methods so that the initial admissible graph is acyclic. The forward–reverse method operates as follows.

TYPICAL ITERATION OF THE FORWARD–REVERSE ϵ -RELAXATION METHOD.

Pick a node i with nonzero surplus; if no such node exists then terminate. If i has positive surplus then perform an iteration of the forward ϵ -relaxation method. If i has negative surplus then perform an iteration of the reverse ϵ -relaxation method.

The idea of the forward–reverse method is recurrent in many relaxation-like methods. Termination of the method can be proved with an analysis similar to the one in section 3, provided that we also make the following assumption.

Assumption. The number of times the surplus of a node changes sign is finite.

The above assumption can be enforced by various mechanisms, some of which are discussed in [Tse86] for the relaxation method and in [Pol94] for the auction shortest path algorithm.

6. Computational results. We have developed and tested two experimental Fortran codes implementing the methods of this paper for convex cost problems. The first code, named NE-RELAX-F, implements the forward ϵ -relaxation method with the sweep implementation and ϵ -scaling as described in section 4. The second code, named NE-RELAX-FV, implements the forward–reverse version of NE-RELAX-F as described in section 5. These codes are based on the ϵ -relaxation code for linear cost problems described in Appendix 7 of [Ber91], which has been shown to be quite efficient. Several changes and enhancements were introduced in the codes for convex cost problems: all computations are done in real rather than integer arithmetic, and ϵ -scaling, rather than arc cost scaling, is used. Also, the updating of the push lists and prices are changed to improve efficiency. Otherwise, the sweep implementation and the general structure of the codes for linear and convex cost problems are identical. Initial testing on linear cost problems showed that the codes for convex cost problems perform as well as, and often better than, their counterparts for linear cost problems, which indicates that these codes are written efficiently. (The superior performance of the codes for convex cost problems may be due to the latter’s efficient management of the push lists and the speed of floating point computations of the machine on which the codes were run.)

The codes NE-RELAX-F and NE-RELAX-FV were compared to two existing Fortran codes, NRELAX and MNRELAX from [BHT87]. The latter implement the relaxation method for, respectively, strictly convex cost and convex cost problems, and they are believed to be quite efficient. All codes were compiled and run on a Sun Sparc-5 workstation with 24 megabytes of RAM under the Solaris operating system. We used the `-O` compiler option in order to take advantage of the floating point unit and the design characteristics of the Sparc-5 processor. Unless otherwise indicated, all codes terminated according to the same criterion; namely, the cost of the feasible flow vector and the cost of the price vector agree in their first 12 digits.

For our testing, we used convex linear/quadratic problems corresponding to the case of (P) where

$$f_{ij}(x_{ij}) = \begin{cases} a_{ij}x_{ij} + b_{ij}x_{ij}^2 & \text{if } 0 \leq x_{ij} \leq c_{ij}, \\ \infty & \text{otherwise} \end{cases}$$

for some a_{ij} , b_{ij} , and c_{ij} with $-\infty < a_{ij} < \infty$, $b_{ij} \geq 0$, and $c_{ij} \geq 0$. We call a_{ij} , b_{ij} , and c_{ij} the linear cost coefficient, the quadratic cost coefficient, and the capacity, respectively, of arc (i, j) . We created the test problems using two Fortran problem generators. The first is the public-domain generator NETGEN, written by Klingman, Napier, and Stutz [KNS74], which generates linear cost assignment/transportation/transshipment problems having a certain random structure. The second is the generator CHAINGEN, written by the second author, which generates transshipment problems having a chain structure as follows: starting with a chain through all the nodes, a user-specified number of forward arcs are added to each node (for example, if the user specifies 3 additional arcs per node then the arcs

TABLE 1

All problems are generated by NETGEN with linear cost coefficients in the range [1–100], total supply of 10000, one pure source and one pure sink, and arc capacities in the range [100–500] (except for problems 22–24 whose capacities are in the range [1000–2000]). For all problems, all arcs have quadratic cost coefficient in the range [5–10]. The run times for the codes (in seconds) were obtained on a Sun Sparc 5 with 24MB memory.

Problem	N	A	NRELAX	MNRELAX	NE-RELAX-F	NE-RELAX-FV
1	200	1300	7.9	6.0	1.9	1.7
2	200	1500	7.5	6.3	2.1	1.7
3	200	2000	2.8	5.6	2.1	1.7
4	200	2200	20.4	10.6	2.4	2.0
5	200	2900	2.3	24.8	1.4	1.2
6	300	3150	7.3	22.1	2.7	1.6
7	300	4500	7.5	21.1	3.9	2.6
8	300	5155	48.3	26.7	3.8	2.0
9	300	6075	7.2	22.7	3.2	2.4
10	300	6300	4.4	31.5	2.7	3.3
11	400	1500	69.2	15.0	8.7	7.8
12	400	2250	17.6	17.2	4.9	4.3
13	400	3000	22.0	20.4	7.3	6.0
14	400	3750	13.2	24.3	3.0	1.8
15	400	4500	10.0	35.9	7.4	6.2
16	400	1306	85.1	25.4	8.6	8.4
17	400	2443	31.6	21.5	7.3	6.2
18	400	1416	7.5	9.0	0.9	0.9
19	400	2836	45.4	26.7	8.6	7.8
20	400	1382	79.9	17.7	9.9	8.4
21	400	2676	33.4	23.9	6.8	5.8
22	1000	3000	64.4	50.9	8.4	4.1
23	1000	5000	26.7	49.0	4.0	3.5
24	1000	10000	26.3	323.2	5.5	5.5

($i, i + 2$), ($i, i + 3$), ($i, i + 4$) are added for each node i) and, for a user-specified percentage of nodes i , a reverse arc ($i, i - 1$) is also added. The graphs thus created have long diameters, and earlier tests on linear cost problems showed that the created problems are particularly difficult for all methods. As the above two generators create only linear cost problems, we modified the created problems as in [BHT87] so that a user-specified percent of the arcs generated a nonzero quadratic cost coefficient in a user-specified range.

Our tests were designed to study two key issues:

- (a) the performance of the ϵ -relaxation methods relative to the relaxation methods and the dependence of this performance on network topology and problem ill conditioning,
- (b) the sensitivity of the ϵ -relaxation methods to problem ill conditioning.

Ill-conditioned problems were created by assigning to some of the arcs smaller (but nonzero) quadratic cost coefficients compared to other arcs. When the arc cost functions have this structure, ill conditioning in the traditional sense of unconstrained nonlinear programming tends to occur.

We experimented with three sets of test problems: the first set comprises well-conditioned strictly convex quadratic cost problems generated using NETGEN (see Table 1); the second set comprises well-conditioned strictly convex quadratic cost problems generated using CHAINGEN (see Table 2); the third set comprises ill-conditioned strictly convex quadratic cost problems and mixed linear/quadratic cost

TABLE 2

All problems are generated by CHAINGEN with linear cost coefficients in the range [1–100], a supply of 1000 at node 1 and a demand of 1000 at node N , and arc capacities in the range [100–1000]. For all problems, all arcs have quadratic cost coefficient in the range [5–10] and half of the nodes have an additional reverse arc. The run times for the codes (in seconds) were obtained on a Sun Sparc 5 with 24MB memory. For these problems, the running times for NRELAX were excessively long even for 5 digits of accuracy and hence are not reported. For problem 10, MNRELAX did not terminate after the time shown, and this is indicated by the > in front of the time.

Problem	N	Added Arcs	A	MNRELAX	NE-RELAX-F	NE-RELAX-FV
1	50	4	269	1.1	0.1	0.2
2	100	4	544	14.9	0.6	0.8
3	150	4	819	15.6	1.2	1.0
4	200	4	1094	33.0	2.1	2.1
5	250	4	1369	41.0	2.4	2.7
6	300	6	2235	93.9	4.6	5.2
7	350	6	2610	266.9	5.9	6.3
8	400	8	3772	1102.6	10.4	10.3
9	450	8	4247	2152.5	10.8	11.3
10	500	10	5705	>1300	17.7	17.5

TABLE 3

All problems are generated by NETGEN with linear cost coefficients in the range [1–100], total supply of 1000, one pure source and one pure sink, and arc capacities in the range [100–300]. For all problems, half of the arcs have quadratic cost coefficient in the range [5–10] and the remaining half have the small quadratic coefficient shown. Note that problems 6 and 12 are mixed cost problems. The runs for the codes (in seconds) were obtained on a Sun Sparc 5 with 24MB memory. For NRELAX, the numbers in parentheses indicate the number of significant digits of solution accuracy obtained by NRELAX in the running time shown.

Problem	N	A	Small Quad Cost	NRELAX	MNRELAX	NE-RELAX-F	NE-RELAX-FV
1	200	1300	1	3.6	3.6	0.5	0.5
2	200	1300	0.1	20.9	4.3	0.6	0.9
3	200	1300	0.01	56.1	3.6	0.6	1.1
4	200	1300	0.001	(5)791.2	3.2	0.7	0.7
5	200	1300	0.0001	(5)1866.6	2.7	0.7	0.7
6	200	1300	0	-	-	0.6	0.8
7	400	4500	1	52.2	14.1	1.7	1.8
8	400	4500	0.1	53.4	11.2	1.8	2.0
9	400	4500	0.01	(5)80.5	13.7	2.3	2.6
10	400	4500	0.001	(5)710.7	15.0	2.6	2.6
11	400	4500	0.0001	(4)5753.4	13.5	3.6	3.1
12	400	4500	0	-	-	2.7	2.6

problems generated using NETGEN (see Table 3). The running time of the codes on these problems are shown in the last three to four columns of Tables 1–3. In all problems, the ϵ -relaxation codes were run to the point where they yielded higher or comparable solution accuracy than the relaxation codes. From the running times we can draw the following conclusions: first, the ϵ -relaxation codes NE-RELAX-F and NE-RELAX-FV have similar performance and both consistently outperform, by a factor of at least 3 and often much more, the relaxation codes NRELAX and MNRELAX on all test problems, independent of network topology and problem ill conditioning. In fact, on the CHAINGEN problems, the ϵ -relaxation codes outperform the relaxation codes by an order of magnitude or more. Other than the favorable complexity results that we obtained in this paper, we have no clear explanation of

this phenomenon.

REFERENCES

- [Ber79] D. P. BERTSEKAS, *A Distributed Algorithm for the Assignment Problems*, Laboratory for Information and Decision Systems Working Paper, M.I.T., Cambridge, MA, 1979.
- [Ber85] D. P. BERTSEKAS, *A unified framework for minimum cost network flow problems*, Math. Programming, 32 (1985), pp. 125–145.
- [Ber86a] D. P. BERTSEKAS, *Distributed Relaxation Methods for Linear Network Flow Problems*, in Proc. 25th IEEE Conference on Decision and Control, Athens, Greece, 1986, pp. 2101–2106.
- [Ber86b] D. P. BERTSEKAS, *Distributed Asynchronous Relaxation Methods for Linear Network Flow Problems*, Laboratory for Information and Decision Systems report P-1606, M.I.T., Cambridge, MA, 1986.
- [Ber91] D. P. BERTSEKAS, *Linear Network Optimization: Algorithms and Codes*, M.I.T. Press, Cambridge, MA, 1991.
- [Ber92] D. P. BERTSEKAS, *An Auction/Sequential Shortest Path Algorithm for the Min Cost Flow Problem*, Laboratory for Information and Decision Systems report P-2146, M.I.T., Cambridge, MA, 1992.
- [BeC93] D. P. BERTSEKAS AND D. A. CASTANON, *A generic auction algorithm for the minimum cost network flow problem*, Comput. Optim. Appl., 2 (1993), pp. 229–260.
- [BCE95] D. P. BERTSEKAS, D. CASTANON, J. ECKSTEIN, AND S. A. ZENIOS, *Parallel network optimization survey*, Handbooks Oper. Res. Management Sci., 7 (1995), pp. 331–339.
- [BeE87] D. P. BERTSEKAS AND J. ECKSTEIN, *Distributed Asynchronous Relaxation Methods for Linear Network Flow Problems*, in Proc. International Federation of Automatic Control '87, Munich, Germany, July 1987.
- [BeE88] D. P. BERTSEKAS AND J. ECKSTEIN, *Dual coordinate step methods for linear network flow problems*, Math. Programming, 42 (1988), pp. 203–243.
- [BeE87] D. P. BERTSEKAS AND D. EL BAZ, *Distributed asynchronous relaxation methods for convex network flow problems*, SIAM J. Control Optim., 25 (1987), pp. 74–85.
- [BHT87] D. P. BERTSEKAS, P. A. HOSEIN, AND P. TSENG, *Relaxation methods for network flow problems with convex arc costs*, SIAM J. Control Optim., 25 (1987), pp. 1219–1243.
- [BeT88] D. P. BERTSEKAS AND P. TSENG, *Relaxation methods for minimum cost ordinary and generalized network flow problems*, Oper. Res., 36 (1988), pp. 93–114.
- [BeT94] D. P. BERTSEKAS AND P. TSENG, *RELAX-IV: A Faster Version of the RELAX Code for Solving Minimum Cost Flow Problems*, Laboratory for Information and Decision Systems report P-2276, M.I.T., Cambridge, MA, 1994.
- [BeT89] D. P. BERTSEKAS AND J. N. TSITSIKLIS, *Parallel and Distributed Computation: Numerical Methods*, Prentice–Hall, Englewood Cliffs, NJ, 1989.
- [BlJ92] R. G. BLAND AND D. L. JENSEN, *On the computational behavior of a polynomial-time network flow algorithm*, Math. Programming, 54 (1992), pp. 1–39.
- [DMZ95] R. DE LEONE, R. R. MEYER, AND A. ZAKARIAN, *An ϵ -Relaxation Algorithm for Convex Network Flow Problems*, Computer Sciences Department Technical report, University of Wisconsin, Madison, WI, 1995.
- [EdK72] J. EDMONDS AND R. M. KARP, *Theoretical improvements in algorithmic efficiency for network flow problems*, J. Assoc. Comput. Mach., 19 (1972), pp. 248–264.
- [FoF62] L. R. FORD, JR., AND D. R. FULKERSON, *Flows in Networks*, Princeton University Press, Princeton, NJ, 1962.
- [GoT90] A. V. GOLDBERG AND R. E. TARJAN, *Solving minimum cost flow problems by successive approximation*, Math. Oper. Res., 15 (1990), pp. 430–466.
- [Gol87] A. V. GOLDBERG, *Efficient Graph Algorithms for Sequential and Parallel Computers*, Laboratory for Computer Science Technical report TR-374, M.I.T., Cambridge, MA, 1987.
- [Hag92] W. W. HAGER, *The dual active set algorithm*, in Advances in Optimization and Parallel Computing, P. M. Pardalos, ed., North–Holland, Amsterdam, the Netherlands, 1992, pp. 137–142.
- [HaH93] W. W. HAGER AND D. W. HEARN, *Application of the dual active set algorithm to quadratic network optimization*, Comput. Optim. Appl., 1 (1993), pp. 349–373.
- [KaM84] P. V. KAMESAM AND R. R. MEYER, *Multipoint methods for separable nonlinear networks*, Math. Programming Study, 22 (1984), pp. 185–205.

- [KaM93] A. V. KARZANOV AND S. T. MCCORMICK, *Polynomial methods for separable convex optimization in unimodular linear spaces with applications to circulations and cocirculations in network*, SIAM J. Comput., 26 (1997), pp. 1248–1278.
- [KNS74] D. KLINGMAN, A. NAPIER, AND J. STUTZ, *NETGEN - A program for generating large scale (un) capacitated assignment, transportation, and minimum cost flow network problems*, Management Sci., 20 (1974), pp. 814–822.
- [LiZ91] X. LI AND S. A. ZENIOS, *Data Parallel Solutions of Min-Cost Network Flow Problems Using ϵ -Relaxations*, Department of Decision Sciences report 1991-05-20, University of Pennsylvania, Philadelphia, PA, 1991.
- [Mey79] R. R. MEYER, *Two-segment separable programming*, Management Sci., 25 (1979), pp. 285–295.
- [NiZ93] S. S. NIELSEN AND S. A. ZENIOS, *On the massively parallel solution of linear network flow problems*, in Network Flow and Matching: First DIMACS Implementation Challenge, D. Johnson and C. McGeoch, eds., American Mathematical Society, Providence, RI, 1993, pp. 349–369.
- [Pol94] L. C. POLYMENAKOS, *Parallel shortest path auction algorithms*, Parallel Comput., 20 (1994), pp. 1221–1247.
- [Pol95] L. C. POLYMENAKOS, *ϵ -Relaxation and Auction Algorithms for the Convex Cost Network Flow Problem*, Electrical Engineering and Computer Science Department Ph.D. thesis, M.I.T., Cambridge, MA, 1995.
- [Roc80] H. RÖCK, *Scaling techniques for minimal cost network flows*, in Discrete Structures and Algorithms, U. Pape and Carl Hanser, eds., München, Germany, 1980, pp. 181–191.
- [Roc70] R. T. ROCKAFELLAR, *Convex Analysis*, Princeton University Press, Princeton, NJ, 1970.
- [Roc84] R. T. ROCKAFELLAR, *Network Flows and Monotropic Programming*, Wiley-Interscience, New York, NY, 1984.
- [Tse86] P. TSENG, *Relaxation Methods for Monotropic Programming Problems*, Operations Research Center Ph.D. thesis, M.I.T., Cambridge, MA, 1986.
- [TBT90] P. TSENG, D. P. BERTSEKAS, AND J. N. TSITSIKLIS, *Partially asynchronous, parallel algorithms for network flow and other problems*, SIAM J. Control Optim., 28 (1990), pp. 678–710.
- [Ven91] J. A. VENTURA, *Computational development of a lagrangian dual approach for quadratic networks*, Networks, 21 (1991), pp. 469–485.
- [Wei74] A. WEINTRAUB, *A primal algorithm to solve network flow problems with convex costs*, Management Sci., 21 (1974), pp. 87–97.