

SCALING METHODS FOR FINDING A MAXIMUM FREE MULTIFLOW OF MINIMUM COST

ANDREW V. GOLDBERG AND ALEXANDER V. KARZANOV

Suppose we are given an undirected graph with nonnegative integer-valued edge capacities and costs in which a subset of nodes is specified. We consider the problem of finding a collection of flows between arbitrary pairs of specified nodes such that the capacity constraints are satisfied and the sum of costs of flows is minimum, provided that the sum of values of flows is maximum. It is known that this problem has a *half-integer* optimal solution and such a solution can be found in strongly polynomial time using the ellipsoid method.

In this paper we give two "purely combinatorial" polynomial algorithms for finding a half-integer optimal solution. These are based on capacity and cost scaling techniques and use the double covering method earlier worked out for the problem.

1. Introduction. By a *graph* we mean a finite undirected graph with possible multiple edges; VG and EG are the sets of nodes and edges of a graph G , respectively. A *network* is a quadruple $N = (G, T, c, a)$ consisting of a graph G , a subset $T \subseteq VG$ of its nodes, called *terminals*, and nonnegative integer-valued functions c (of *capacities*) and a (of *costs*) on the edges of G .

A simple path in G between two distinct terminals is called a *T-path*; let $\mathcal{P} = \mathcal{P}(G, T)$ denote the set of *T-paths*. A (*c-admissible*) *multicommodity flow*, or, briefly, a *multiflow*, in N is a nonnegative rational-valued function $f: \mathcal{P} \rightarrow \mathbb{Q}_+$ satisfying the capacity constraint

$$\xi^f(e) := \sum (f(P): e \in P \in \mathcal{P}) \leq c(e) \quad \text{for all } e \in EG$$

(when writing $e \in P$, we consider a path as an edge-set). The *total value*, v_f , of f is $\sum (f(P): P \in \mathcal{P})$, and the *total cost*, a_f , of f is $\sum (a(e)\xi^f(e): e \in EG)$. A multiflow of the maximum total value is called *maximum*. We deal with the following *minimum cost maximum multiflow* problem:

(1.1) Find a maximum multiflow f in N whose total cost a_f is as small as possible.

The simplest case arises when $|T| = 2$. Then (1.1) turns into the well-known (undirected) minimum cost maximum *flow* problem. A fundamental fact is that the latter problem has an optimal solution which is integer-valued; see, e.g., Ford and Fulkerson (1962). This is, in general, not true for $|T| \geq 3$. Nevertheless, a weaker property of *half-integrality* has been proved, as follows.

THEOREM 1 (KARZANOV 1979). *There exists an optimal solution f to (1.1) such that $2f$ is integer-valued.*

REMARK 1. Problem (1.1) is a special case of the general undirected minimum cost maximum multiflow problem in which pairs $\{s_1, t_1\}, \dots, \{s_r, t_r\}$ of nodes are

specified, and the goal is to connect these pairs by flows so that the resulting multiflow is c -admissible and has the minimum cost provided that its total value is maximum. The special case we study allows flows between arbitrary pairs of terminals; in other words, the graph H induced by the pairs $\{s_1, t_1\}, \dots, \{s_r, t_r\}$ as edges is a complete graph. Following Lomonosov (1985), we call a multiflow for a complete H a *free* multiflow. The relationship between the denominators of multiflow entries and the structure of H is studied in Karzanov (1987). It turned out that if H is not a complete multi-partite graph, then for each $k \in \mathbb{Z}_+$, there exists an instance of G, c, a such that kf is not integral for any optimal solution f for G, H, c, a . On the other hand, when H is a complete multi-partite graph, the problem is easily reduced to the free multiflow case, whence the problem has a half-integer optimal solution by Theorem 1.

REMARK 2. The existence of a half-integer maximum free multiflow was established, independently, by Lovász (1976) and Cherkassky (1977). Theorem 1 extends that result to the cost case.

Theorem 1 appeared as a consequence of a similar result concerning a more general *parametric problem*, namely:

(1.2) Given $p \in \mathbb{Z}_+$, find a multiflow f in N which maximizes $pv_f - a_f$.

THEOREM 2 (KARZANOV 1979). For any $p \geq 0$, (1.2) has a half-integer optimal solution.

Obviously, (1.2) becomes equivalent to (1.1) when p is large enough (from the existence of a half-integer solution to (1.2) it follows that $p = 2a(EG)c(EG) + 1$ is sufficient). So Theorem 1 immediately follows from Theorem 2. Here and later on for a function g on a set S and a subset $S' \subseteq S$, $g(S')$ denotes $\sum(g(e) : e \in S')$.

Theorem 2 was originally proved by constructive means, relying on a pseudo-polynomial algorithm which solves (1.2) in time $O(\min\{(c(EG) + 1)q, (a(EG) + 1)q', 2^{q''}\})$, where q, q', q'' are polynomials in $|VG|, |EG|$. Further results were obtained in Karzanov (1994) where a relatively simple, nonalgorithmic, proof of Theorem 2 was given and it was shown how to find a half-integer optimal solution to (1.2) (and, therefore, (1.1)) in strongly polynomial time by use of the ellipsoid method. However, no polynomial algorithm which uses solely "combinatorial means" has been known until the present.

Like the classic algorithm for the min-cost max-flow problem due to Ford and Fulkerson (1962), the algorithm in Karzanov (1979) works in frameworks of the primal-dual method of linear programming and iteratively reduces (1.2) to simpler, noncost, subproblems on multiflows in "feasible subgraphs" Γ of G , which are solved by use of combinatorial techniques. The key idea is that the subproblem allows a reduction to a certain single-commodity flow problem in a skew-symmetric digraph D such that each edge of Γ corresponds to a pair of "skew-symmetric arcs" in D . Such a D is called the *double covering* over Γ . This explains the existence of a half-integer optimal solution to (1.2): The flow problem in D has an integer solution g and the back reduction increases the solution "fractionality" twice, transforming g into a half-integer multiflow in Γ .

In this paper we design two "purely combinatorial" polynomial algorithms for finding a half-integer optimal solution to (1.2). The first algorithm works with an arbitrary $p \in \mathbb{Z}_+$, while the second algorithm works when p is large; in particular, both algorithms solve (1.1). At the high level, instead of the primal-dual approach, they apply variants of the scaling method: the first algorithm scales capacities (cf. Edmonds and Karp 1972) and runs in $O(\log(c(EG) + 2))$ iterations, while the second one scales costs (cf. Röck 1980; Bland and Jensen 1992) and runs in $O(\log(a(EG) +$

$2) + \log(c(Z) + 2))$ iterations. Here Z is the set of edges e with $a(e) = 0$. An iteration of each algorithm deals with a subproblem which is somewhat different from, but closely related to, that of Karzanov (1979). We show that each subproblem is reduced to a sequence of directed flow problems and that the total time required to perform an iteration is polynomial in $|VG|, |EG|$.

It should be noted that the original version of this paper (Goldberg and Karzanov 1993) presented polynomial time capacity and cost scaling algorithms based on similar ideas, but was technically more complicated, using the concept of so-called transitive fork environments—a formalism which allows us to extend standard augmenting and alternating path techniques to rather general situations. This formalism was, in fact, introduced in Karzanov (1993) for solving the integer strengthening of (1.1).

For technical reasons, throughout the paper we assume that the cost function a is *positive*, i.e., $a(e) > 0$ for all $e \in EG$. This assumption does not lead to loss of generality. Indeed, if the set $Z = \{e \in EG: a(e) = 0\}$ is nonempty, we can replace p by $p' = (2c(Z) + 1)p$, and a by a' defined by

$$\begin{aligned} a'(e) &= (2c(Z) + 1)a(e) \quad \text{for } e \in EG - Z, \\ &= 1 \quad \text{for } e \in Z. \end{aligned}$$

Then for any two *half-integer* maximum multiflows f and f' with $\Delta = (pv_f - a_f) - (pv_{f'} - a_{f'}) > 0$, we have $\Delta \geq 1/2$ and

$$\begin{aligned} &(p'v_f - a'_f) - (p'v_{f'} - a'_{f'}) \\ &= (2c(Z) + 1)((pv_f - a_f) - (pv_{f'} - a_{f'})) - \zeta^f(Z) + \zeta^{f'}(Z) \\ &\geq (2c(Z) + 1)\Delta - c(Z) \geq c(Z) + \frac{1}{2} - c(Z) > 0. \end{aligned}$$

Therefore, if f is an optimal solution for p', a' , then f is also an optimal solution for p, a . Note that a factor of $c(Z) + 1$ in the definition of a' causes the rising of the term involving c in the complexity of the cost scaling algorithm declared above.

Although we allow multiple edges in G , when it is not confusing, we denote an edge with end nodes u and v by uv .

This paper is organized as follows. Section 2 contains background for the algorithms that we develop. It considers the dual of (1.2), describes the construction of double covering digraphs D , and briefly reviews the relationship between multiflows in G and flows in D . Sections 3–6 are devoted to the capacity scaling algorithm. Section 7 contains the cost scaling algorithm; its description is shorter because both algorithms utilize many common tools.

2. Double covering. The linear program dual of (1.2) is

$$(2.1) \quad \text{Minimize } c \cdot \gamma \text{ over all } \gamma: EG \rightarrow \mathbb{Q}_+ \text{ such that } \gamma(P) \geq p - a(P) \text{ for each } P \in \mathcal{P},$$

viewing P as an edge-set and denoting by $g \cdot h$ the inner product $\sum(g(e)h(e): e \in S)$ of functions g and h on S . For $l \in \mathbb{Q}_+^{EG}$, let $\text{dist}_l(x, y)$ denote the l -distance between nodes $x, y \in VG$, i.e., the minimum l -length $l(P)$ of a path P from x to y . The constraints in (2.1) can be expressed in a more compact form as

$$(2.2) \quad \text{dist}_{a+\gamma}(s, t) \geq p \quad \text{for any } s, t \in T, s \neq t.$$

4

The linear programming duality theorem applied to (1.2) and (2.1) implies that a multiflow f and a vector $\gamma \in \mathbb{Q}_+^{EG}$ satisfying (2.2) are optimal solutions to these problems if and only if the following "complementary slackness" conditions hold:

$$(2.3) \quad \text{For } P \in \mathcal{P}, \text{ if } f(P) > 0 \text{ then } a(P) + \gamma(P) = p;$$

$$(2.4) \quad \text{For } e \in EG, \text{ if } \gamma(e) > 0 \text{ then } e \text{ is saturated by } f, \text{ i.e., } \zeta^f(e) = c(e).$$

In the rest of this section we consider a *positive* function l on EG which satisfies $\text{dist}_l(s, t) \geq p$ for all distinct $s, t \in T$ (usually l is equal to $a + \gamma$). For brevity, we denote $\text{dist}_l(\cdot, \cdot)$ by $\text{dist}(\cdot, \cdot)$ and omit the prefix l in the terms l -length, l -distance and l -shortest. A T -path of the length exactly p is called a (T, l) -line, or, briefly, a T -line. A part of a T -line is called a *line*. By (2.2), every line P is a shortest path; so P is simple (as l is positive). The structure of lines has important properties described below; for detailed proofs, see Karzanov (1979, 1994).

We need some terminology and notations. The subgraph of G formed by T and the nodes and edges contained in lines is called the l -graph and denoted by $\Gamma = \Gamma^l$. For $v \in V\Gamma$, the *potential*, $\pi(v) = \pi_l(v)$, of v is the distance from v to T (i.e., $\min\{\text{dist}(v, s) : s \in T\}$), and $T(v) = T_l(v)$ denotes the set of terminals $s \in T$ closest to v (i.e., $\text{dist}(v, s) = \pi(v)$).

Obviously, $\pi(v) \leq p/2$ for any $v \in V\Gamma$; moreover, $\pi(v) = p/2$ implies $|T(v)| \geq 2$, and $\pi(v) < p/2$ implies $|T(v)| = 1$. Therefore, the nodes of Γ are naturally partitioned into the sets $V_s = \{v \in V\Gamma : T(v) = \{s\}\}$ for $s \in T$ and $V^\bullet = \{v \in V\Gamma : |T(v)| \geq 2\}$; the nodes in V^\bullet are called *central*. One can obtain the following characterization of the lines (in particular, the edges of Γ) in terms of potentials:

(2.5) A path $P = (v_0, e_1, v_1, \dots, e_k, v_k)$ in G with both ends in Γ is a line if and only if either (i) $v_0, v_k \in V_s \cup V^\bullet$ and $|\pi(v_k) - \pi(v_0)| = l(P)$ for some $s \in T$, or (ii) $v_0 \in V_s, v_k \in V_t$ and $\pi(v_0) + \pi(v_k) + l(P) = p$ for some distinct $s, t \in T$.

In particular, in view of the positivity of l , $\pi(x) \neq \pi(y)$ for any edge $xy \in E\Gamma$ with $x, y \in V_s$, and no edge of Γ connects two central nodes. Property (2.5) easily implies that

(2.6) A T -path $P = (v_0 = s, e_1, v_1, \dots, e_k, v_k = t)$ in Γ is a T -line if and only if there is $0 \leq i < k$ such that $v_0, \dots, v_i \in V_s; v_{i+2}, \dots, v_k \in V_t; \pi(v_0) < \dots < \pi(v_i); \pi(v_{i+2}) > \dots > \pi(v_k)$; and either $v_{i+1} \in V^\bullet$, or $v_{i+1} \in V_t$ and $\pi(v_{i+1}) > \pi(v_{i+2})$.

Using (2.5) for one-edge paths (edges), we now construct the digraph $D^l = D = (VD, AD)$, the *double covering* over Γ , mentioned in the Introduction, as follows. Split each $v \in V\Gamma$ into $2|T(v)|$ nodes v_s^1 and v_s^2 ($s \in T(v)$). If v is noncentral, i.e., $T(v)$ consists of a single terminal s , then v_s^1 and v_s^2 are called the *first* and *second* copies of v in D , respectively. The arcs of D are defined as follows:

- (2.7) (i) Each edge $e = uv \in E\Gamma$ with $u \in V_s, v \in V_s \cup V^\bullet$ and $\pi(u) < \pi(v)$ generates two arcs (u_s^1, v_s^1) and (v_s^2, u_s^2) ;
 (ii) Each edge $e = uv \in E\Gamma$ with $u \in V_s$ and $v \in V_t$ ($s \neq t$) generates two arcs (u_s^1, v_t^2) and (v_t^1, u_s^2) ;
 (iii) Each central node v generates arcs (v_s^1, v_t^2) for all distinct $s, t \in T(v)$.

An illustration is given in Figure 1 where $T = \{s, t, q\}$, $p = 4$, the numbers on edges indicate values of l , and the arcs of D are directed upward. An arc as in (2.7)(iii) is

called *central*, or a *v-arc*, and has infinite capacity. An arc as in (2.7)(i) or (ii) has the same capacity as that of the edge which generates this arc. We keep the same notation c for the capacities in D and think of $T^1 = \{s^1: s \in T\}$ and $T^2 = \{s^2: s \in T\}$ as the sets of *sources* and *sinks* of D , respectively.

Since for each $b = (u_s^i, v_t^j) \in AD$, the pair $b' = (v_t^{3-j}, u_s^{3-i})$ is also an arc of D , the mapping $\sigma: v_s^i \rightarrow v_t^{3-i}$ gives a (skew) *symmetry* of D ; we extend σ to the arcs denoting b' by $\sigma(b)$. Next, the construction of D yields a natural mapping ω of $VD \cup AD$ to $V\Gamma \cup E\Gamma$ which brings a node v_s^i to v , a noncentral arc (u_s^i, v_t^j) to the edge uv , and a central arc (v_s^1, v_t^2) to the node v . We extend σ and ω in a natural way to the dipaths and other objects in D . For example, the dipath $\sigma(P)$ symmetric to a dipath $P = (x_0, b_1, x_1, \dots, b_k, x_k)$ is $(\sigma(x_k), \sigma(b_k), \sigma(x_{k-1}), \dots, \sigma(b_1), \sigma(x_0))$; the path $\omega(P)$ for this P is the sequence $(\omega(x_0), \omega(b_1), \omega(x_1), \dots, \omega(b_k), \omega(x_k))$ with the repeated central nodes (if any) deleted. A function h on AD is *symmetric* if $h(b) = h(\sigma(b))$ for each $b \in AD$.

Important properties of D , easily derived from (2.6) and (2.7), are as follows:

(2.8) Dipaths P and $\sigma(P)$ are disjoint, and $\omega(\sigma(P))$ is reverse to $\omega(P)$;

(2.9) ω yields a *one-to-one correspondence* between the set of T -lines and the set of T^1 to T^2 dipaths in D .

Such a correspondence is further extended to a relationship between certain multiflows in G and flows in D , as follows. A multiflow f is called *going along T-lines* if each T -path P with $f(P) > 0$ is a T -line. A (c -admissible T^1 to T^2) *flow* in D is a function $g: AD \rightarrow \mathbb{Q}_+$ satisfying the conservation condition

$$\text{div}_g(x) := \sum_{y: (x,y) \in AD} g(x,y) - \sum_{y: (y,x) \in AD} g(y,x) = 0$$

for all $x \in VD - (T^1 \cup T^2)$,

and the capacity constraint $g(b) \leq c(b)$ for all $b \in AD$. Note that D is acyclic, therefore, a flow g can be represented as the sum of elementary flows along dipaths. That is, there are T^1 to T^2 dipaths P_1, \dots, P_m and rationals $\alpha_1, \dots, \alpha_m \geq 0$ such that $\sum(\alpha_i; b \in P_i) = g(b)$ for each $b \in AD$; we call $\mathcal{D} = \{(P_i, \alpha_i)\}$ a *decomposition* of




FIGURE 1.

g . Such a \mathcal{D} induces the function $f = f^{\mathcal{D}}$ on \mathcal{P} by setting $f(\omega(P_i)) = \alpha_i/2$ for $i = 1, \dots, m$ and $f(P) = 0$ for the remaining T -paths P . For any noncentral arc $b \in AD$, we have

$$(2.10) \quad \zeta^f(\omega(b)) = \frac{1}{2}(g(b) + g(\sigma(b))) \leq \frac{1}{2}(c(b) + c(\sigma(b))) = c(\omega(b)).$$

Hence, f is c -admissible. Also f goes along T -lines (by (2.9)). Furthermore, if g is integer-valued, then it has an *integer* decomposition \mathcal{D} (i.e., with all α_i 's integral), whence $f^{\mathcal{D}}$ is *half-integral* (this is the crucial observation in the proof of Theorem 2 in Karzanov 1979, 1994).

Conversely, let $f: \mathcal{P} \rightarrow \mathbb{Q}_+$ be a multiflow going along T -lines. Define the function $g = g^f$ on AD so that for $b \in AD$, $g(b)$ is the sum of numbers $f(\omega(P))$ over all T^1 to T^2 dipaths P in D that contain b or $\sigma(b)$. Then g is a symmetric flow, and for each noncentral arc $b \in AD$, we have

$$(2.11) \quad g(b) = g(\sigma(b)) = \zeta^f(\omega(b)).$$

Now return to consideration of a function γ satisfying (2.2) and let $l = a + \gamma$. It is convenient to extend γ to the arcs b of D^l by setting $\gamma(b) = \gamma(\omega(b))$ if b is noncentral, and $\gamma(b) = 0$ otherwise. We say that an arc b is *feasible* if $\gamma(b) = 0$, and a flow g in D is *feasible* if it saturates each infeasible arc b , $g(b) = c(b)$. Summing up arguments above and using (2.10) and (2.11), we derive the following.

STATEMENT 2.1 (KARZANOV 1979). The flow g^f induced by an optimal solution f to (1.2) is feasible. Moreover, if, in addition, f is half-integral, so is g . Conversely, the multiflow $f^{\mathcal{D}}$ induced by a feasible flow g and a decomposition \mathcal{D} of g is an optimal solution to (1.2). Moreover, if both g and \mathcal{D} are integral, then f is half-integral.

Our algorithms will also use a result on the fractionality of dual problem (2.1).

THEOREM 2.2 (KARZANOV 1994). (2.1) has a half-integer optimal solution.

3. Capacity scaling algorithm. As before, we assume that a is positive. Also one may assume that $p > 0$. At the high level, the algorithm for (1.2) that we now develop applies a standard capacity scaling approach and consists of *big* (or *scaling*) *iterations*. The numbers of these iterations is equal to the number of ones in the binary notation of the capacities.

More precisely, let I be the maximum of numbers $\lceil \log_2(c(e) + 1) \rceil$ among all $e \in EG$. For $i = 0, \dots, I$ and $e \in EG$, define the truncated capacity $c^i(e) = \lfloor c(e)/2^{I-i} \rfloor$. Then $c^0 = 0$ and $c^I = c$. For $i = 1, \dots, I$ define U^i to be the set of edges e with $c^i(e) > 2c^{i-1}(e)$. Assuming that the elements of U^i are arbitrarily ordered as $z_1, \dots, z_{k(i)}$ ($k(i) = |U^i|$), define the capacities $c_j^i(e) = c^i(e) - 1$ ($= 2c^{i-1}(e)$) for $e = z_{j+1}, \dots, z_k$, and $c_j^i(e) = c^i(e)$ for the other edges e in G , $j = 0, \dots, k(i)$. In particular, $c_0^i = 2c^{i-1}$.

The algorithm starts with obvious optimal primal and dual solutions f^0 and γ^0 to (1.2) for G, T, c^0, a, p . More precisely, we put $f^0 = 0$, and choose γ^0 so as to provide $\text{dist}_{a+\gamma^0}(s, t) \geq p$ for all distinct $s, t \in T$ (e.g., for $e = xy \in EG$, one can take $\gamma^0(e)$ to be $\max(0, p - a(e))$ if $x, y \in T$, 0 if $x, y \in VG - T$, and $\max(0, p/2 - a(e))$ otherwise).

The input of a current big iteration (with $1 \leq i \leq I$ and $1 \leq j \leq k(i)$) consists of optimal primal and dual solutions f_{j-1}^i and γ_{j-1}^i for G, T, c_{j-1}^i, a, p , and the goal is to transform them into optimal f_j^i and γ_j^i for c_j^i (letting $f_0^i = 2f^{i-1}$ and $\gamma_0^i = \gamma^{i-1}$, where f^{i-1} and γ^{i-1} are optimal solutions found for c^{i-1}). So the last big iteration finds optimal solutions for the original c .

Thus, the following auxiliary problem has to be solved at most $|EG|$ times; for convenience we keep the same notation as for the original problem.

(3.1) Given optimal half-integral f and γ for G, T, \hat{c}, a, p and an edge $z \in EG$, find optimal half-integral f' and γ' for G, T, c, a, p , where $c(z) = \hat{c}(z) + 1$ and $c(e) = \hat{c}(e)$ for all $e \in EG - \{z\}$.

We describe an algorithm to solve (3.1) in strongly polynomial time, thus providing the desired result for (1.2). For reasons which will be clear later, we assume that each pair of distinct terminals is connected by an edge of a large capacity and zero cost; the set of these edges is denoted by W . The latter assumption does not lead to loss of generality because, obviously, each element of W must be saturated in any optimal solution for \hat{c} or c (in view of $p > 0$). Also we may assume that c is positive; for if $c(e) = 0$ for some e , we can simply delete e from G .

If $\gamma(z) = 0$, then f and γ remain optimal for c (since (2.4) holds for z). So assume that $\gamma(z) > 0$. Then $\zeta^f(z) = \hat{c}(z) = c(z) - 1$, by (2.4). Our aim is either to make z saturated by increasing the total multiflow through z by one, or to reduce $\gamma(z)$ to zero, in order to restore the violated complementary slackness condition at z . This task will require at most two *primal updates* (which increase $\zeta(z)$) and $O(|VG|)$ *dual updates* (which decrease $\gamma(z)$).

We use the construction from the previous section. Let $l = a + \gamma$ and construct the l -graph $\Gamma = \Gamma^l$, digraph $D = D^l$ and flow $g = g^f$. As before, γ is extended to AD . We may assume that z is in Γ because if z belongs to no T -line, we can decrease $\gamma(z)$ by the maximum amount allowed by (2.2). Then γ takes zero values on all edges in $EG - E\Gamma$ (in view of (2.3), (2.4) (for \hat{c}) and the positivity of c). Let $\bar{z} = (q, r)$ and $\bar{z}' = (r', q')$ be the arcs forming $\omega^{-1}(z)$. We need to slightly modify D and g to $H = H^l$ and h , the *extended double covering* and a *circulation* in it, as follows.

(i) H is formed by adding to D arcs of infinite capacity and zero cost from each sink s^2 to each source t^1 . The set of these arcs is denoted by \bar{W} ; we formally extend γ by zero to \bar{W} . Extend g to each arc in \bar{W} by a *rather large* positive half-integer in such a way that the resulting function g' is a circulation, i.e., $\text{div}_{g'}(v) = 0$ for all nodes v . This can be done since the edges from W generate arcs b from sources s^1 to sinks t^2 , for all distinct $s, t \in T$, and $g(b) = c(b)$ is large.

(ii) Transform g' into an *integer* circulation h feasible for \hat{c}, γ by increasing or decreasing $g'(b)$ by $1/2$ at each arc b for which $g'(b)$ is not integral. This can be done since these b 's are feasible (as, obviously, $0 < g'(b) < \hat{c}(b)$, yielding $\gamma(b) = 0$), and the fact that each node v satisfies $\text{div}_{g'}(v) = 0$ implies that v is incident to an even number of arcs b with $g'(b)$ not integral.

Thus, h is an integer circulation in H , and \bar{z}, \bar{z}' are the only arcs which make h infeasible for c .

4. Primal update. We attempt to increase h at the arcs \bar{z}, \bar{z}' using techniques motivated by usual augmenting paths in flow problems. Let F be the set of feasible arcs in H , and H_F the digraph $(VH, F \cup \{\bar{z}, \bar{z}'\})$. For a digraph D' and a subset $X \subseteq VD'$, let $\delta_{D'}^+(X)$ ($\delta_{D'}^-(X)$) denote the set of arcs $b = (x, y) \in AD'$ *leaving* X , i.e., with $x \in X \not\equiv y$ (respectively, *entering* X , i.e., with $x \notin X \ni y$), and let $\delta_{D'}(X) = \delta_{D'}^+(X) \cup \delta_{D'}^-(X)$ (the *cut* induced by X).

A path $P = (v_0, b_1, v_1, \dots, b_k, v_k)$ in H_F is called *active* if $h(b_i) < c(b_i)$ for all *forward* arcs $b_i = (v_{i-1}, v_i)$, $h(b_i) > 0$ for all *backward* arcs $b_i = (v_i, v_{i-1})$ in P , and P uses neither \bar{z} nor \bar{z}' as a backward arc. If, in addition, P is a simple circuit and contains at least one of \bar{z} and \bar{z}' , P is called an *augmenting circuit*. One can

efficiently find such a circuit or certify that it does not exist (cf. Ford and Fulkerson 1962). Three cases are possible.

Case 1. An augmenting circuit P containing both \bar{z}, \bar{z}' is found. In this case we push a unit of flow along P , i.e., increase h by one at each forward arc of P , and decrease it by one at each backward arc. Then the resulting function h' is a (c -admissible) integer circulation with $h'(\bar{z}) = h'(\bar{z}') = c(z)$. Therefore, the half-integer multiflow $f^{\mathcal{D}}$ induced by an integer decomposition \mathcal{D} of $h'_{|AD}$, along with γ , solves (3.1).

Case 2. No augmenting circuit exists. Then the set $X \subseteq VH$ reachable by active paths from r does not contain q , and the set Y reachable by active paths from q' does not contain r' . Note that if, say, Y contains q , then, obviously, X does not contain r' . So, without loss of generality, one may assume that $q, r' \notin X$.

Case 3. An augmenting circuit P containing exactly one of \bar{z} and \bar{z}' , say $\bar{z}' \in P$, is found. Push a unit of flow along P , obtaining a circulation h' with $h'(\bar{z}') = c(z)$. Then search for an augmenting circuit P' for h' and \bar{z} . (i) If such a P' is found, push unit flow along P' . The resulting circulation saturates (for c) both \bar{z}, \bar{z}' and induces an optimal solution to (3.1). (ii) Now suppose that P' does not exist. Then the set $X \subseteq VH$ reachable from r by active paths for h' does not contain q . Furthermore, the following is true (this will also be used in §§6 and 7).

STATEMENT 4.1. Let h' be a circulation in H with $h'(\bar{z}) < h'(\bar{z}')$, and let X be the set of nodes of H reachable from r by active paths for h' . Let $q \notin X$. Then $r' \notin X$.

PROOF. Let F_1 and F_2 (respectively, Q_1 and Q_2) be the sets of feasible (respectively, infeasible) arcs in $\delta_H^+(X)$ and $\delta_H^-(X)$, respectively. Then $\bar{z} \in Q_2, h'(b) = c(b)$ for each $b \in (F_1 \cup Q_1 \cup Q_2) - \{\bar{z}, \bar{z}'\}$, and $h'(b) = 0$ for each $b \in F_2$ (in view of the definition of X).

Suppose that $r' \in X$. Let $\alpha = c(\bar{z}) - h'(\bar{z})$ and $\beta = c(\bar{z}') - h'(\bar{z}')$; then $\alpha > \beta$. If $q' \in X$ (whence $\bar{z}' \notin \delta_H(X)$), put $\Delta = \alpha$, and if $q' \notin X$ (whence $\bar{z}' \in Q_1$), put $\Delta = \alpha - \beta$. We have

$$(4.1) \quad \begin{aligned} 0 = \operatorname{div}_{h'}(X) &= h'(F_1) + h'(Q_1) - h'(F_2) - h'(Q_2) \\ &= c(F_1) + c(Q_1) - c(Q_2) + \Delta. \end{aligned}$$

Consider the symmetric set $X' = \sigma(X)$. By symmetry, $r, r' \in X' \not\equiv q'$. Also $q \in X'$ if and only if $q' \in X$. Define the subsets F'_1, F'_2, Q'_1, Q'_2 of $\delta_H(X')$ similarly to those for $\delta_H(X)$. Then $\bar{z}' \in Q'_1$. If $q \in X'$ (whence $\bar{z} \notin \delta_H(X')$), put $\Delta' = \beta$, and if $q \notin X'$ (whence $\bar{z} \in Q'_2$), put $\Delta' = \beta - \alpha$. We have

$$(4.2) \quad \begin{aligned} 0 = -\operatorname{div}_{h'}(X') &= h'(F'_2) + h'(Q'_2) - h'(F'_1) - h'(Q'_1) \\ &\leq c(F'_2) + c(Q'_2) - c(Q'_1) + \Delta'. \end{aligned}$$

Note that F'_2, Q'_1, Q'_2 are symmetric to F_1, Q_2, Q_1 , respectively. Therefore, $c(F'_2) + c(Q'_2) - c(Q'_1) = c(F_1) + c(Q_1) - c(Q_2)$. Now (4.1) and (4.2) yield $\Delta \leq \Delta'$, which is impossible (in both cases) since $\alpha > \beta$. \square

In Case 1 and in (i) of Case 3, the primal updates applied (that update h) give an optimal solution to (3.1). In what follows we assume that no augmenting circuit exists, i.e., we are in Case 2 or in (ii) of Case 3. Then the algorithm continues by performing the dual update described in the next section. For convenience we keep the notation h for the current circulation in H , and the notation f for a corresponding multiflow in G (defined up to a decomposition of $h_{|AD}$).

For $v \in VH$, let X_v denote the set of nodes reachable from v by active paths in H . The above analysis shows that, up to symmetry, we deal with one of the following situations:

(4.3)

- (i) $h(\bar{z}) < h(\bar{z}') (= c(z))$, and $q, r' \notin X_r$;
- (ii) $h(\bar{z}) = h(\bar{z}') < c(z)$, $q, r' \notin X_r$, and $q \in X_{q'} \not\equiv r'$;
- (iii) $h(\bar{z}) = h(\bar{z}') < c(z)$, $q, r' \notin X_r \cup X_{q'}$, and $X_r \cap X_{q'} \neq \emptyset$;
- (iv) $h(\bar{z}) = h(\bar{z}') < c(z)$, $q, r' \notin X_r \cup X_{q'}$, and $X_r \cap X_{q'} = \emptyset$.

We call (i), (ii), (iii), (iv) *Situations (of type) 1, 2, 3, 4*, respectively. In this hierarchy, type i is considered as more preferable than type j if $i < j$. We shall see in §6 that each dual update either solves (3.1), or leads to a primal update, or is followed by another dual update of the same or a more preferable type. Define R to be $\{r, q'\}$ in Situation 3, and $\{r\}$ in the other situations (R is the set of *roots* of active paths we are going to deal with). In what follows by the *reachability* we always mean the reachability by use of active paths (from or to a specified node or subset of nodes). The set of nodes reachable from R is denoted by X , and the symmetric set $\sigma(X)$ by X' .

REMARK. At first glance, it may look artificial to distinguish between Situations 3 and 4 and define R (and, therefore, X) in these situations differently. However, we shall see later that if R were defined in the same way (to be either $\{r\}$ or $\{r, q'\}$) in these situations, we could not guarantee that both the dual half-integrality and the property of z to be an edge of the current l -graph are simultaneously maintained.

In Situation 2, we also introduce the set

$$(4.4) \quad M = \{x \in VH: x \text{ is reachable from } q' \text{ and reachable to } q\}.$$

A subset $Y \subseteq VH$ is called *tight* if $h(b) = c(b)$ for each $b \in \delta_H^+(Y) \cap F$, and $h(b) = 0$ for each $b \in \delta_H^-(Y) \cap F$; e.g., X is tight. The sets X and X' have important properties, which will often be used later on.

STATEMENT 4.2. (i) If $Y \subseteq VH$ is tight, then $VH - \sigma(Y)$ is also tight; in particular, $VH - X'$ is tight. (ii) $X \cap X' = \emptyset$. (iii) X' coincides with the set X'' of nodes reachable to $R' = \sigma(R)$.

PROOF. (i) is proved similarly to Statement 4.1. To see (ii), observe that $q, r' \notin X$ implies $R \cap X' = \emptyset$. Therefore, $R \subseteq X - X'$. Next, it is easy to see that if Y and Z are tight, then both $Y \cap Z$ and $Y \cup Z$ are tight too. This implies that $X - X'$ is tight (since both X and $VH - X'$ are tight). From the fact that each element of X is reachable from R it easily follows that no proper subset of X which includes R can be tight. Hence, $X - X' = X$, yielding (ii). Finally, to see (iii), observe that $VH - (X' \cap X'')$ is tight (since both $VH - X'$ and $VH - X''$ are tight). Also $R' \subseteq X' \cap X''$. From the definition of X'' it follows that $VH - X''$ is the maximum tight set disjoint from R' . Therefore, $X' \cap X'' = X''$, whence $X'' \subseteq X'$. The reverse inclusion $X' \subseteq X''$ follows from the inclusion $X \subseteq \sigma(X'')$ (since $\sigma(X'')$ is tight, by (i)). \square

The fact that H contains the set \bar{W} of arcs, each with nonzero flow, leads to the following important statement.

STATEMENT 4.3. $X \cap (T^1 \cup T^2) = \emptyset$ (and, therefore, $X' \cap (T^1 \cup T^2) = \emptyset$).

PROOF. Suppose that X contains some $s^i \in T^1 \cup T^2$. The existence of an arc $b = (s^2, s^1)$ (in \bar{W}) with $0 < h(b) < c(b)$ implies that s^{3-i} is also reachable. Hence, both s^1, s^2 are in X . By symmetry, $s^1, s^2 \in X'$, contrary to the fact that X and X' are disjoint (by Statement 4.2). \square

Finally, in Situation 2, we observe the following:

- (4.5)
- (i) M is self-symmetric, i.e., $M = \sigma(M)$;
 - (ii) M and $X \cup X'$ are disjoint.

PROOF. M is the intersection of X_q and the set Y of nodes reachable to q . Arguing as in the proof of (iii) in Statement 4.2, one shows that $Y = \sigma(X_q)$. This implies (i). Next, if (ii) is false, then M and X have a common node (as $M = \sigma(M)$ and $X' = \sigma(X)$). Hence, there is an active path from r to q ; a contradiction. \square

5. Dual update. In the classic algorithm by Ford and Fulkerson (1962) for the min-cost max-flow problem, the current dual vector is transformed within a single cut. In contrast, the dual update described below deals simultaneously with two cuts, namely, $\delta_H(X)$ and $\delta_H(X')$, where X and X' were defined in the previous section. Another feature, which makes the procedure more involved, is that the update of γ may change the l -graph Γ and, therefore, the double covering D over Γ . In fact, our method uses some basic ideas of the dual update from the algorithm in Karzanov (1979). However, compared with the latter, we mainly work with the l -graph Γ rather than the double covering D , which seems to be simpler and more enlightening. When it is not confusing, we omit H in $\delta_H^+(\cdot)$, $\delta_H^-(\cdot)$ and $\delta_H(\cdot)$.

First of all we have to explain how X can intersect the subgraph $H(v) = \omega^{-1}(v)$ generated by a central vertex v ($H(v)$ is induced by the v -arcs and called a *central subgraph*, or the *v -subgraph*). For $i = 1, 2$, let $V^i(v)$ denote the set $\{v_s^i: s \in T(v)\}$ (so $\{V^1(v), V^2(v)\}$ is a partition of $VH(v)$).

STATEMENT 5.1. Let $Y = X \cap VH(v)$ be nonempty. Then: (i) $Y = \{v_s^1\} \cup (V^2(v) - \{v_s^2\})$ for some $s \in T(v)$; and (ii) $h(b) = 0$ for each arc from $V^1(v) - \{v_s^1\}$ to $V^2(v) - \{v_s^2\}$.

PROOF. Suppose that Y contains v_s^1 for some $s \in T(v)$. Since X is tight and each arc (v_s^1, v_t^2) is not saturated (as it has infinite capacity), X contains v_t^2 for all $t \in T(v) - \{s\}$. Then X' contains the symmetric nodes v_s^2 and v_t^1 ($t \neq s$), and now the fact that X and X' are disjoint (by Statement 4.2) gives (i). To see (ii), observe that each arc b from $V^1(v) - \{v_s^1\}$ to $V^2(v) - \{v_s^2\}$ is feasible and enters X ; therefore, $h(b) = 0$.

Next suppose that X contains a node v_s^2 and none of the nodes in $V^1(v)$. Then $h(b) = 0$ for each arc b from $V^1(v)$ to v_s^2 . By the construction of D , no other arcs enter v_s^2 . Therefore, the total flow entering v_s^2 is zero, whence $\text{div}_h(v_s^2) = 0$ implies that $h(b) = 0$ for each arc b leaving v_s^2 . But then v_s^2 cannot be reachable from R ; a contradiction. \square

In order to describe the update of γ , we associate with X certain subsets and functions in Γ . For $s \in T$, define the sets

$$(5.1) \quad A_s = \{v \in V_s: v_s^1 \in X\}; \quad A_s^\bullet = \{v \in V^\bullet: s \in T(v), v_s^1 \in X\};$$

$$B_s = \{v \in V_s: v_s^2 \in X\}; \quad \tilde{V}_s = V_s \cup A_s^\bullet.$$

From Statements 4.2(ii) and 5.1 it follows that the sets A_s, A_s^\bullet, B_s for all $s \in T$ are pairwise disjoint, and their preimages in D give a partition of $X \cup X'$. Let $\tilde{V}^\bullet = V^\bullet - \cup(A_s^\bullet: s \in T)$. Define the function ρ on $V\Gamma$ by

$$(5.2) \quad \begin{aligned} \rho(v) &= -1 && \text{if } v \in A_s \cup A_s^\bullet, s \in T, \\ &= 1 && \text{if } v \in B_s, s \in T, \\ &= 0 && \text{otherwise;} \end{aligned}$$

and the function λ on $E\Gamma$ by

$$(5.3) \quad \begin{aligned} \lambda(e) &= \rho(v) - \rho(u) && \text{if } e = uv, u, v \in \tilde{V}_s \cup \tilde{V}^\bullet, \pi(u) < \pi(v), \\ &= -\rho(u) - \rho(v) && \text{if } e = uv, u \in \tilde{V}_s, v \in \tilde{V}_t, s \neq t. \end{aligned}$$

The algorithm chooses some rational number $\epsilon \geq 0$ and makes the ϵ -transformation $\gamma \rightarrow \gamma^\epsilon$ which replaces $\gamma(e)$, $e \in E\Gamma$, by

$$(5.4) \quad \gamma^\epsilon(e) = \max\{0, \gamma(e) + \epsilon\lambda(e)\},$$

and maintains $\gamma^\epsilon(e) = 0$ for $e \in EG - E\Gamma$. Accordingly, define the length function l^ϵ to be $a + \gamma^\epsilon$. Since γ^ϵ is nonnegative, l^ϵ is positive.

First we show that such an update of γ is correct for a sufficiently small $\epsilon > 0$. The proof falls into several parts. For $j = -2, -1, 0, 1, 2$, define $E^j = \{e \in E\Gamma: \lambda(e) = j\}$. For $U, W \subseteq VH$, let (U, W) denote the set of arcs from U to W in H . Let C denote the set of central arcs in H .

STATEMENT 5.2. Let $Z = VH - (X \cup X')$. The following are true:

$$(5.5) \quad \begin{aligned} E^{-2} &= \{\omega(b): b \in (X', X) - C\}, \\ E^{-1} &= \{\omega(b): b \in (Z, X)\} \quad (= \{\omega(b): b \in (X', Z)\}), \\ E^1 &= \{\omega(b): b \in (X, Z)\} \quad (= \{\omega(b): b \in (Z, X')\}), \\ E^2 &= \{\omega(b): b \in (X, X')\}. \end{aligned}$$

PROOF. Consider an arc $b = (u_s^i, v_t^j) \in \delta(X)$. Let b belong to (X', X) . Since $i \leq j$ (in view of (2.7)) and $\sigma(b) \in (X', X)$, we may assume that $i = 1$. (2.7) shows that either (i) $s = t$, $j = 1$, $u \in V_s$, $v \in V_s \cup V^\bullet$ and $\pi(u) < \pi(v)$, or (ii) $s \neq t$, $j = 2$, $u \in V_s$ and $v \in V_t$. In case (i), X contains u_s^2 and v_s^1 , and we have $u \in B_s$ and $v \in A_s \cup A_s^\bullet$ (by (5.1)), whence $\rho(u) = 1$ and $\rho(v) = -1$ (by (5.2)). In case (ii), X contains u_s^2 and v_t^2 , whence $\rho(u) = \rho(v) = 1$. In both cases, $\lambda(\omega(b)) = -2$, by (5.3).

The cases with b in (Z, X) , (X, Z) or (X, X') are examined in a similar way (note that in these cases b is noncentral, by Statement 5.1). Also a similar analysis shows that if b is noncentral and not in $\delta(X) \cup \delta(X')$, then $\lambda(\omega(b)) = 0$. A careful examination of these remaining cases is left to the reader. \square

For $i = -1, -2$, define $J^i = \{e \in E\Gamma: \gamma(e) = 0, \lambda(e) = i\}$. Expression (5.5) and the tightness of X imply that

$$(5.6) \quad \begin{aligned} (i) \quad & z \in E^{-1} \cup E^{-2}; \\ (ii) \quad & \zeta^f(e) = 0 \text{ for each } e \in J^{-1} \cup J^{-2}; \\ (iii) \quad & \zeta^f(e) = c(e) \text{ for each edge } e \neq z \text{ in } (E^{-2} - J^{-2}) \cup (E^{-1} - J^{-1}) \cup E^1 \cup E^2. \end{aligned}$$

Therefore, increasing ϵ decreases $\gamma(z)$, and each edge $e \neq z$ with $\gamma^\epsilon(e) > 0$ remains saturated by f . The choice of ϵ for the update of γ is restricted by several factors. The first upper bound on ϵ is

$$(5.7) \quad \epsilon_1 = \min\{\min\{\gamma(e): e \in E^{-1} - J^{-1}\}, \min\{\gamma(e)/2: e \in E^{-2} - J^{-2}\}\}.$$

In particular, $\epsilon_1 \leq \gamma(z)$. Obviously, $\epsilon_1 > 0$. In other words, ϵ_1 is the minimum ϵ for which the ϵ -transformation reduces some positive $\gamma(e)$ to zero. In what follows we

always assume that $0 < \epsilon \leq \epsilon_1$. Another important property provided by the choice of ϵ_1 is as follows.

STATEMENT 5.3. Let P be a (T, l) -line. Then $l^\epsilon(P) \geq p$. Moreover, $l^\epsilon(P) = p$ if $f(P) > 0$.

PROOF. Consider the T^1 to T^2 dipath $Q = \omega^{-1}(P)$. By Statement 4.3, the ends of Q are outside both X and X' . Hence, Q enters X as many times as it leaves X , and similarly for X' . Also if $\delta(X) \cup \delta(X')$ contains a central arc b , then $b \in (X', X)$. These arguments together with (5.5) show that $\lambda(P) = 2\alpha$, where α is the number of times Q meets central arcs in $\delta(X)$ (clearly α is 0 or 1). Next, since $\epsilon \leq \epsilon_1$, we derive from (5.4) and (5.7) that $\gamma^\epsilon(P) = \gamma(P) + \epsilon\lambda(P) + \epsilon\beta_1 + 2\epsilon\beta_2$, where β_i is the number of edges of P in J^{-i} . Hence,

$$\begin{aligned} l^\epsilon(P) &= a(P) + \gamma^\epsilon(P) = a(P) + \gamma(P) + 2\epsilon\alpha + \epsilon\beta_1 + 2\epsilon\beta_2 \\ &\geq a(P) + \gamma(P) = l(P) = p. \end{aligned}$$

The second part of the statement is valid because if $f(P) > 0$ then $\beta_1 = \beta_2 = 0$ (by (5.6)(ii)) and $\alpha = 0$ (by Statement 5.1 and the fact that $h(b) > 0$ for all arcs b of Q). \square

A (T, l) -line P with $l^\epsilon(P) = p$ is called a *nonbroken line*. By Statement 5.3, each edge $e \in E\Gamma$ with $\zeta^J(e) > 0$ belongs to a nonbroken line. Let us call an arc b of D *forbidden* if it is feasible (i.e., $\gamma(b) = 0$) and either enters X or leaves X' or both. From the above proof one can see that

(5.8) for a (T, l) -line P , the following are equivalent:

- (i) P is nonbroken;
- (ii) $\omega^{-1}(P)$ has no forbidden arcs;
- (iii) P neither meets $J^{-1} \cup J^{-2}$ nor contains three consecutive nodes $u \in V_t$, $v \in A_s^\bullet$, $w \in V_{t'}$ for $s \neq t, t'$.

Statement 5.3 implies that for a sufficiently small ϵ , the l^ϵ -distance between distinct terminals is at least p , and f goes along (T, l^ϵ) -lines (taking into account that l^ϵ is continuous in ϵ , that $l(P) > p$ for each T -path P in G which is not a (T, l) -line, and that the number of these (simple) paths is finite). Our aim is to find a reasonable bound on ϵ to guarantee these properties.

Let Γ' be the subgraph of Γ formed by T and the nodes and edges of nonbroken lines. By Statement 5.3, f goes along lines in Γ' . This implies that γ^ϵ takes nonzero values only within $E\Gamma' \cup \{z\}$ (since the edges in $E^1 \cup E^2$ are saturated, by (5.6)(iii)). Using (5.8), one can give an efficient procedure to extract Γ' from Γ . There is a simple explicit description of Γ' , as follows.

Let J be the set of edges e of Γ such that any (T, l) -line passing through e meets $J^{-1} \cup J^{-2}$; clearly J can be found efficiently. By (5.8), the elements of J cannot belong to nonbroken lines. The converse property also holds, as shown by the following statement which will be proved in the next section.

STATEMENT 5.4. Each of the following belongs to a nonbroken line: (i) the node $\omega(x)$ for each $x \in X$; (ii) each edge in $E\Gamma - J$; (iii) the edge z .

Thus, assuming that this statement is true, we observe that $E\Gamma' = E\Gamma - J$, that Γ' includes A_s , A_s^\bullet and B_s for all $s \in T$, and that $\gamma^\epsilon(e) = 0$ for all $e \in E\Gamma - E\Gamma'$. For $s \in T$, let $V'_s = \tilde{V}_s \cap V\Gamma'$, and let $V'^\bullet = \tilde{V}^\bullet \cap V\Gamma'$ (cf. (5.1)). Then $A_s, A_s^\bullet, B_s \subseteq V'_s$. For $v \in V\Gamma'$, define

$$(5.9) \quad \pi^\epsilon(v) = \pi(v) + \epsilon\rho(v).$$

Comparing (2.5) with (5.3) and (5.4), we observe that for $e = uv \in E\Gamma$,

$$(5.10) \quad \begin{aligned} l^\epsilon(e) &= |\pi^\epsilon(v) - \pi^\epsilon(u)| && \text{if } u, v \in V'_s \cup V'^\bullet, s \in T, \\ &= p - \pi^\epsilon(u) - \pi^\epsilon(v) && \text{if } u \in V'_s, v \in V'_t, s \neq t. \end{aligned}$$

When ϵ is growing from zero, π^ϵ decreases by ϵ on the elements of each $A_s \cup A_s^\bullet$, increases by ϵ on the elements of each B_s , and remains unchanged on the other nodes of Γ . This motivates the next upper bound on ϵ to be

$$(5.11) \quad \epsilon_2 = \min\{p/2 - \pi(v) : v \in B_s, s \in T\},$$

which says that the growth of ϵ should be stopped as soon as some increasing number $\pi^\epsilon(v)$ achieves $p/2$. From (5.10) it follows that π^ϵ is the potential function for l^ϵ in Γ' whenever $\epsilon \leq \epsilon_2$ (i.e., for each $v \in V\Gamma'$, $\pi^\epsilon(v)$ is the minimum l^ϵ -length of a path from T to v in Γ'). Since each set B_s has no central nodes for l , $\epsilon_2 > 0$.

Consider two nodes u and v in Γ' . Let $d(u, v)$ be the minimum cost $a(P)$ of a path P from u to v in G with all edges and inner nodes (if any) not in Γ' ; if such a path does not exist, $d(u, v) = \infty$. When ϵ is growing, it may happen at some moment that such a minimum cost path P becomes a line for the current length function (and further growing ϵ would make the distance between some terminals less than p). To account for such a moment, we associate with the *ordered* pair (u, v) two numbers $\bar{l}(u, v)$ and $\bar{\lambda}(u, v)$, where

$$(5.12) \quad \begin{aligned} \bar{l}(u, v) &= \pi(v) - \pi(u) && \text{and } \bar{\lambda}(u, v) = \rho(v) - \rho(u) \\ &&& \text{if } u, v \in V'_s \cup V'^\bullet, s \in T, \\ \bar{l}(u, v) &= p - \pi(u) - \pi(v) && \text{and } \bar{\lambda}(u, v) = -\rho(u) - \rho(v) \\ &&& \text{if } u \in V'_s, v \in V'_t, s \neq t. \end{aligned}$$

One can see that if u and v are connected by a line (e.g., an edge) P in Γ' , then $l^\epsilon(P)$ equals $\max\{\bar{l}(u, v) + \epsilon\bar{\lambda}(u, v), \bar{l}(v, u) + \epsilon\bar{\lambda}(v, u)\}$. Our third upper bound on ϵ is

$$(5.13) \quad \epsilon_3 = \min\{\epsilon(u, v) : u, v \in V\Gamma', \epsilon(u, v) \geq 0\},$$

where

$$(5.14) \quad \epsilon(u, v) = (d(u, v) - \bar{l}(u, v)) / \bar{\lambda}(u, v).$$

We observe that $d(u, v) > \bar{l}(u, v)$; therefore, $\epsilon_3 > 0$.

STATEMENT 5.5. Let $0 \leq \epsilon \leq \min\{\epsilon_1, \epsilon_2, \epsilon_3\}$. Then $\text{dist}_T^\epsilon(s, t) \geq p$ for all distinct $s, t \in T$. In particular, Γ' is a subgraph of Γ^{l^ϵ} .

PROOF. We know that if $\epsilon' \leq \epsilon_1, \epsilon_2$, then $\pi^{\epsilon'}$ is the potential function for $l^{\epsilon'}$ within Γ' . Suppose that $l^{\epsilon'}(P) < p$ for some T -path P . Using the potentials $\pi^{\epsilon'}$ in Γ' and the fact that each node of Γ' belongs to a $(T, l^{\epsilon'})$ -line in it, one can see that there is a part P' of P , say, from u to v , such that both u, v are in Γ' while the other nodes of P' are not, and either (i) $u, v \in V'_s \cup V'^\bullet$ and $l^{\epsilon'}(P') < |\pi^{\epsilon'}(v) - \pi^{\epsilon'}(u)|$, or (ii) $u \in V'_s, v \in V'_t$ ($s \neq t$) and $l^{\epsilon'}(P') < p - \pi^{\epsilon'}(u) - \pi^{\epsilon'}(v)$.

In case (i), assume for definiteness $\pi^\epsilon(v) \geq \pi^\epsilon(u)$. We have $l^\epsilon(P') = a(P') \geq d(u, v)$ (by the definition of d). Also $\pi^\epsilon(v) - \pi^\epsilon(u) = \pi(v) + \epsilon\rho(v) - \pi(u) - \epsilon\rho(u) = \bar{l}(u, v) + \epsilon\bar{\lambda}(u, v)$. Therefore, $l^\epsilon(P') < \pi^\epsilon(v) - \pi^\epsilon(u)$ implies $d(u, v) < \bar{l}(u, v) + \epsilon\bar{\lambda}(u, v)$. This is impossible when $\bar{\lambda}(u, v) \leq 0$ (since $d(u, v) > \bar{l}(u, v)$). And if $\bar{\lambda}(u, v) \geq 0$, then $\epsilon(u, v) > 0$, and now (5.14) and $\epsilon \leq \epsilon_3 \leq \epsilon(u, v)$ imply $d(u, v) \geq \bar{l}(u, v) + \epsilon\bar{\lambda}(u, v)$; a contradiction.

Case (ii) leads to a contradiction in a similar way. \square

Thus, the choice of ϵ to be $\min\{\epsilon_1, \epsilon_2, \epsilon_3\}$ gives γ^ϵ which, together with f , satisfies (2.2)–(2.4) for \hat{c} . In the next section we show that such a choice ensures fast convergence to optimal solutions for (3.1).

6. Complexity of the algorithm and dual half-integrality. Our goal is to show that $O(|VG|)$ consecutive dual updates either reduces $\gamma(z)$ to zero or make it possible to apply the primal update increasing the total multiflow at z . First of all we establish three facts (Statements 6.1–6.3).

STATEMENT 6.1. The edge z belongs to a nonbroken line, i.e., (iii) in Statement 5.4 is true.

PROOF. If $\zeta^f(z) > 0$, the result immediately follows from Statement 5.3. So assume that $\zeta^f(z) = 0$. Then $h(\bar{z}) = h(\bar{z}') = 0$, therefore, we are in Situation 2, 3 or 4 (see (4.3)). Choose a dipath Q_1 from r to T^2 in D , and let b be the first arc of Q_1 leaving X . Since $h(b) > 0$ (as X is tight), we may assume that Q_1 is chosen so that all its arcs following b have nonzero flow. Then Q_1 has no forbidden arcs.

Suppose that a situation of type 3 takes place. Then $q \in X'$ (as $R = \{r, q'\}$). The argument “symmetric” to that above shows that there is a T^1 to q dipath Q_2 without forbidden arcs. Now the concatenation of the path Q_2 , the arc \bar{z} and the path Q_1 gives a T^1 to T^2 dipath which contains \bar{z} and has no forbidden arcs.

In Situation 4, we argue in a similar way using the facts that q belongs to the set $\sigma(X_{q'})$ which is disjoint from X and that $VH - \sigma(X_{q'})$ is tight (as $X_r \cap X_{q'} = \emptyset$, $\sigma(X_r) \cap X_{q'} = \emptyset$ and $X_{q'}$ is tight).

Now consider a situation of type 2. Let $P = (x_0, b_1, x_1, \dots, b_k, x_k)$ be an active path from q' to q . Note that P cannot contain only forward arcs; for otherwise there would exist a T^1 to T^2 dipath L in D which passes through both \bar{z} and \bar{z}' , implying that the path $\omega(L)$ is not a T -line, contrary to (2.9). Thus, P has a backward arc; let b_i be the last of such arcs. Since $h(b_i) > 0$, there is a dipath Q_2 from T^1 to x_i with nonzero flow at all its arcs. Concatenating the path Q_2 , the part of P from x_i to q , the arc \bar{z} and the path Q_1 , we obtain a T^1 to T^2 dipath which contains \bar{z} and avoids forbidden arcs. \square

STATEMENT 6.2. Let $b = (x, y)$ be an arc of D with $x, y \in X$. Then the element $\omega(b)$ belongs to a nonbroken line.

PROOF. Choose a dipath Q_1 from y to T^2 and an active path P from R to x . By the argument as in the proof of Statement 6.1, we may assume that $h(b) = 0$, and Q_1 has no forbidden arcs. Let w be the last node of P such that either $w \in R$ or w is the tail of an arc with nonzero flow. Then there is a T^1 to w dipath Q_2 without forbidden arcs (if $w \in R$, Q_2 exists by Statement 6.1). Now the concatenation of the path Q_2 , the part of P from w to x , the arc b and the path Q_1 gives a T^1 to T^2 dipath which contains b and has no forbidden arcs. \square

STATEMENT 6.3. Assuming that situation of type 2 occurs, let M be the set defined in (4.4). Let $b = (x, y)$ be an arc of D with $x, y \in M$. Then the element $\omega(b)$ belongs to a nonbroken line.

PROOF. Choose an active path P from y to q and a dipath Q from q to T^2 passing through \bar{z} . Arguing as in the proofs of Statements 6.1 and 6.2, one can transform P and Q into a y to T^2 dipath Q_1 without forbidden arcs. Similarly, there exists a T^1 to x dipath Q_2 without forbidden arcs. Then the concatenation of Q_2 , b and Q_1 is as required. \square

As a corollary from Statements 6.1 and 6.2, we immediately obtain (i) in Statement 5.4 from the previous section. To see (ii) in that statement, consider $e \in E\Gamma - J$ and take a T -line P which contains e but none of the elements of $J^{-1} \cup J^{-2}$. Let $Q = \omega^{-1}(P) = (x_0, b_1, x_1, \dots, b_k, x_k)$ and $e = \omega(b_i)$. Suppose that Q has a forbidden arc b_j , and let for definiteness $j < i$. Since P does not meet $J^{-1} \cup J^{-2}$, the only possible case is when b_j is a central arc going from X' to X . Choose a T^1 to T^2 dipath $L = (y_0, d_1, y_1, \dots, d_m, y_m)$ which contains x_j ($x_j = y_i$, say) and has no forbidden arcs (L exists by (i)). Then the concatenation of the part of L from y_0 to y_i and the part of Q from x_j to x_k contains b_i and none of forbidden arcs, whence (ii) follows.

Return to consideration of the dual update described in the previous section and suppose that it finishes with $\gamma^\epsilon(z) > 0$, where $\epsilon = \min\{\epsilon_1, \epsilon_2, \epsilon_3\}$. Let $\hat{\Gamma}$, \hat{D} and \hat{H} stand for the l^ϵ -graph Γ^{l^ϵ} , double covering D^{l^ϵ} and extended double covering H^{l^ϵ} , respectively, and let $\hat{\sigma}$ and $\hat{\omega}$ be the corresponding mappings for \hat{D} . We know that the graph Γ' generated by the nonbroken lines (see §5) is a subgraph of $\hat{\Gamma}$. However, the subgraph $D' = \omega^{-1}(\Gamma')$ of D^l can differ from the subgraph $\hat{D}' = \hat{\omega}^{-1}(\Gamma')$ of \hat{D} . This happens if the sets of central nodes in Γ' for l and l^ϵ are different.

More precisely, for $s \in T$, let B_s^\bullet be the set of nodes $v \in B_s$ with $\pi^\epsilon(v) = p/2$, and let $A^\bullet = \cup(A_s^\bullet: s \in T)$ and $B^\bullet = \cup(B_s^\bullet: s \in T)$ (where A_s^\bullet and B_s^\bullet are defined in (5.1)). We observe that the set of central nodes for (Γ', l^ϵ) is $(V'^\bullet - A^\bullet) \cup B^\bullet$, where V'^\bullet is the set of central nodes for (Γ', l) . This implies that \hat{D}' is obtained from D' by the following operations which destroy the v -subgraphs generated by the elements $v \in A^\bullet$ and create the v -subgraphs for the elements $v \in B^\bullet$:

(6.1)

- (i) for $v \in A_s^\bullet$, $s \in T$, delete the v -arcs, identify the nodes v_s^1 and v_t^2 for all $t \in T(v) - \{s\}$ with a new node \hat{v}_s^1 (the *first* copy of v in \hat{D}'), and identify the nodes v_s^2 and v_t^1 , $t \in T(v) - \{s\}$, with a new node \hat{v}_s^2 (the *second* copy of v in \hat{D}'); then
- (ii) for each $v \in B_s^\bullet$, $s \in T$, split v_s^1 into nodes \hat{v}_t^1 , $t \in \hat{T}'(v)$, split v_s^2 into nodes \hat{v}_t^2 , $t \in \hat{T}'(v)$, and add central arcs $(\hat{v}_t^1, \hat{v}_t^2)$ for all distinct $t, t' \in \hat{T}'(v)$, where $\hat{T}'(v)$ is the set of terminals of Γ' closest to v with respect to l^ϵ .

Let \hat{x}_t^i stand for the node in \hat{D}' induced by $x \in V\Gamma'$, $t \in \hat{T}'(x)$ and $i \in \{1, 2\}$. One can see that, under transformation (6.1), for each $v \in A_s^\bullet$, an arc of the form (u_s^1, v_s^1) (respectively, (v_t^2, w_t^2) , $t \in T(v) - \{s\}$) in D' becomes an arc of the form $(\hat{u}_s^1, \hat{v}_s^1)$ (respectively, $(\hat{v}_t^2, \hat{w}_t^2)$) in \hat{D}' ; and for each $v \in B_s^\bullet$, an arc of the form (u_t^1, v_s^2) (respectively, (v_s^2, w_s^2)) in D' becomes an arc of the form $(\hat{u}_t^1, \hat{v}_t^1)$ (respectively, $(\hat{v}_s^2, \hat{w}_s^2)$) in \hat{D}' ; and accordingly for the symmetric arcs. (Note that for $v \in B^\bullet$, no arc of the form (u_s^2, v_s^2) can exist.)

Based on these observations, we introduce the mapping τ from VD' to nodes or subsets of nodes of \hat{D}' that brings v_s^i to \hat{v}_s^i if $v \in V\Gamma - (A^\bullet \cup B^\bullet)$, brings each of v_s^1 and v_t^2 ($t \neq s$) to \hat{v}_s^1 if $v \in A_s^\bullet$, and brings v_s^2 to $\{\hat{v}_s^2\} \cup \{\hat{v}_t^2: t \neq s\}$ if $v \in B_s^\bullet$; and accordingly for the symmetric nodes. This induces a one-to-one mapping, also denoted by τ , of the set of arcs of D' different from v -arcs for $v \in A^\bullet$ to the set of arcs of \hat{D}' different from v -arcs for $v \in B^\bullet$ (so that for $b = (x, y)$, $\tau(b)$ goes from $\tau(x)$ to $\tau(y)$). Moreover, obviously, $\hat{h}(\tau(b)) = h(b)$ holds for all arcs b to which τ is applicable, where \hat{h} is a circulation in \hat{H} corresponding to f (\hat{h} is determined

uniquely within \hat{D}). Also each arc $b \in AD'$ with both ends in X satisfies $\gamma^\epsilon(\tau(b)) = \gamma(b)$ (in view of (5.5)); in particular, if b is feasible for γ , then $\tau(b)$ is feasible for γ^ϵ .

Let $\hat{z} = (\hat{q}, \hat{r})$ and $\hat{z}' = (\hat{r}', \hat{q}')$ be the arcs $\tau(\bar{z})$ and $\tau(\bar{z}')$, respectively, and let \hat{X} be the set of nodes of \hat{H} reachable by active (for γ^ϵ and \hat{h}) paths from the corresponding set \hat{R} (that is $\{\hat{r}\}$ or $\{\hat{r}, \hat{q}'\}$). We distinguish between two situations:

(6.2)

- (i) \hat{X} contains some of \hat{q}, \hat{r}' ;
- (ii) none of \hat{q}, \hat{r}' is in \hat{X} .

Case (i) (giving "breakthrough") will be studied later, and now we consider case (ii). By Statement 4.2, $\hat{X} \cap \hat{X}' = \emptyset$, where $\hat{X}' = \hat{\sigma}(\hat{X})$. Define \tilde{X} to be the set of nodes of \hat{D} occurring in $\tau(x)$ for $x \in X$. Summing up the above observations and using Statement 6.2, one can see that each active path from R in D' determines in a natural way an active path from \hat{R} in \hat{H} . Moreover, the following important property holds (in fact, this is valid in case (i) of (6.2) as well):

$$(6.3) \quad \tilde{X} \subseteq \hat{X} \quad (\text{and therefore, } \omega(X) \subseteq \hat{\omega}(\hat{X})).$$

One can see that the first inclusion in (6.3) is strict if and only if the second inclusion is strict.

Note also that, in Situation 2, the set $\omega(M)$, as well as the subgraph of Γ induced by this set, is contained in Γ' , by Statement 6.3. Moreover, since M and $X \cup X'$ are disjoint (by (4.5)), we have $\omega(M) \cap \omega(X) = \emptyset$ (taking into account that if X meets a central subgraph, then $X \cup X'$ contains all nodes of this subgraph). Therefore, the subgraph of D' induced by M remains in \hat{D} . This implies that every node of M remains reachable from \hat{q}' and reachable to \hat{q} in \hat{H} . So the set as in (4.4) does not decrease under the dual update in question. Moreover, if the resulting set \hat{X} meets M , then an active path from \hat{r} to \hat{q} is created, and we obtain "breakthrough" (6.2)(i). Similar arguments show that, in Situation 4, the transformation preserves the subgraph induced by $X_{q'}$, and the set reachable from \hat{q}' does not decrease.

Next, case (6.2)(ii) is further divided into three subcases depending on which of the bounds $\epsilon_1, \epsilon_2, \epsilon_3$ is achieved by ϵ . We show that each of them results in at least one of the following:

(6.4)

- (i) $\hat{\omega}(\hat{X})$ strictly includes $\omega(X)$; or
- (ii) \hat{A} strictly includes A .

Here A denotes $\cup(A_s \cup A_s^\bullet: s \in T)$, and \hat{A} denotes the analogous set for $\hat{\Gamma}, \hat{X}$. Note that $A \subseteq \hat{A}$ always holds because for each $s \in T$ and $v \in A_s \cup A_s^\bullet$, we have $\pi^\epsilon(v) < \pi(v) \leq p/2$ (so v is noncentral for l^ϵ , and $v_s^1 \in X$ implies $\hat{v}_s^1 \in \hat{X}$, whence $v \in \hat{A}_s$).

(A) Let $\epsilon = \epsilon_1$. This means that $\gamma^\epsilon(e) = 0$ for some edge $e \neq z$ with $\gamma(e) > 0$ and $\lambda(e) < 0$. At least one arc $b = (x, y)$ in $\omega^{-1}(e)$ enters X , and $h(b) = c(b) > 0$. The corresponding arc $\tau(b) = (\hat{x}, \hat{y})$ of \hat{D} enters \hat{X} . Therefore (in view of (6.3)), \hat{x} is reachable from \hat{R} by an active path that uses $\tau(b)$ as a backward arc. So \hat{x} is in \hat{X} . Since $\hat{x} \notin \tilde{X}$, we obtain (i) in (6.4).

(B) Let $\epsilon = \epsilon_3$. This means that there appears a new line P which connects some $u, v \in V\Gamma'$ and has no inner node in Γ' . Suppose that $u, v \in V'_s \cup V'^\bullet$ and $\pi(u) < \pi(v)$ for some s . Then $\pi^\epsilon(u) < \pi^\epsilon(v)$, and the dipath $Q = \hat{\omega}^{-1}(P)$ goes from $x = \hat{u}_s^1$ to $y = \hat{v}_s^1$. Furthermore, $d(u, v) - \tilde{l}(u, v) = \epsilon(\rho(v) - \rho(u)) > 0$ shows that either (a) $u \in A_s \not\ni v$, or (b) $u \notin B_s \ni v$, or both.

In case (a), we have $\hat{u}_s^1 \in \hat{X}$. Since all arcs of Q are feasible and nonsaturated, y is reachable, i.e., $y \in \hat{X}$. On the other hand, y cannot be in \tilde{X} . Indeed, $y \in \tilde{X}$ would imply $v \in B_s$ (as $v \notin A_s$). Then X contains v_s^2 , whence $\hat{v}_s^2 \in \tilde{X}$. But \tilde{X} cannot contain both \hat{v}_s^1 and \hat{v}_s^2 . Thus, we have (i) in (6.4).

In case (b), applying similar arguments to the symmetric dipath $\hat{\sigma}(Q)$, we deduce that the end \hat{u}_s^2 of $\hat{\sigma}(Q)$ is in \hat{X} but not in \tilde{X} .

Finally, let $u \in V'_s$ and $v \in V'_t$ for $s \neq t$. Then $\rho(u) + \rho(v) < 0$ (cf. (5.12)). This is possible only if some of u and v , u say, is in the A -set, while the other, v , is not in the B -set. Arguing as above, one shows that \hat{v}_t^2 is in \hat{X} but not in \tilde{X} , again obtaining (i) in (6.4).

(C) Let $\epsilon = \epsilon_3$. Then B^\bullet is nonempty. Consider some $v \in B_s^\bullet$. We know that \tilde{X} contains \hat{v}_s^2 and \hat{v}_t^1 for all $t \in \hat{T}'(v) - \{s\}$. On the other hand, Statement 5.1 shows that if $\hat{X} \cap \hat{X}' = \emptyset$ then the set $\hat{X} \cap \hat{\omega}^{-1}(v)$ is of the form $\{v_s^1\} \cup \{v_t^2: t' \in \hat{T}(v) - \{s'\}\}$. Hence, $|\hat{T}(v)| \geq 3$ is impossible (in this case "breakthrough" happens). Now assuming that $\hat{T}(v)$ consists of two elements s and s' , we conclude that v becomes a member of \hat{A}_s . This gives (ii) in (6.4).

Thus, (6.4) is always valid. In case (6.2)(ii), the next iteration of the algorithm makes the dual update for γ^ϵ , and so on until (i) in (6.2) happens. (6.4) gives an obvious upper bound on the number of such iterations.

STATEMENT 6.3. The number of consecutive dual updates does not exceed $|VG|$. \square

We now consider case (i) in (6.2). If \hat{q} is reachable from \hat{r} , the algorithm makes the corresponding primal update increasing the total multiframe at z ; and similarly if \hat{r}' is reachable from \hat{q}' when \hat{z}' is not saturated (recall that $\hat{h}(\hat{z}) = c(z) - 1 \leq \hat{h}(\hat{z}')$). So we may assume that

(6.5) $\hat{r}, \hat{r}' \in \hat{X} \not\cong \hat{q}$ and either (a) $\hat{h}(\hat{z}') = c(z)$, or (b) $\hat{h}(\hat{z}') = c(z) - 1$ and \hat{r}' is unreachable from \hat{q}' .

Statement 4.1 applied to \hat{h} and \hat{X} shows that (a) is impossible. Consider (b). If before the dual update we were in Situation 2, then, as it was shown above, the current $\hat{H}, \hat{h}, \gamma^\epsilon$ admit an active path P within M which goes from \hat{q}' to \hat{q} . Also \hat{r}' is reachable from \hat{r} by an active path Q . Since $\hat{X} \cap M = \emptyset$, these paths are disjoint. Therefore, P and Q along with \hat{z} and \hat{z}' form an augmenting circuit.

Thus, if we cannot augment the total multiframe at z acting as above, then before the dual update we were in Situation 3 or 4. On the other hand, after the update we fall in Situation 2 (with \hat{z} and \hat{z}' permuted). The algorithm continues by making the dual update with respect to the set reachable from \hat{q}' . Summing up the above arguments, we conclude with the following.

STATEMENT 6.4. The above algorithm solves (3.1) in $O(|VG|)$ iterations, and each iteration is performed in time polynomial in $|VG|, |EG|$. \square

Thus, the whole capacity scaling algorithm for (1.2) is polynomial and has the complexity as indicated in the Introduction.

It remains to show that the optimal dual solution γ' found by the algorithm is half-integral. Since the function γ in the input of (3.1) is assumed to be half-integral, it suffices to prove the following.

STATEMENT 6.5. The numbers ϵ_1, ϵ_2 and ϵ_3 found at the dual iteration are half-integers.

PROOF. By (5.7), ϵ_1 is a half-integer if $\gamma(e)$ is an integer for each edge e with $\gamma(e) > 0$ and $\lambda(e) = -2$. To see the latter, consider an arc $b = (x, y) \in \omega^{-1}(e)$;

then $x \in X'$ and $y \in X$ (cf. (5.5)). Choose in H a path P_1 from r to y and a path P_2 from x to r' such that all nodes of P_1 are in X , all nodes of P_2 are in X' , and all arcs of these paths are feasible. Such a P_1 (and similarly, P_2) exists because, in situation of each type, the subgraph induced by feasible arcs with both ends within X is connected and spans X (this is just the point where we essentially use distinction in the definition of X for types 3 and 4).

The images of the arcs of P_1 and P_2 together with e form a subgraph of Γ which contains a circuit $C = (v_0, e_1, v_1, \dots, e_k, v_k = v_0)$ involving e ; let $e = e_1$. Since the arcs of P_1 and P_2 are feasible, $l(e_i) = a(e_i)$ for $i = 2, \dots, k$. Also $l(e_i)$ is either $|\pi(v_{i-1}) - \pi(v_i)|$ or $p - \pi(v_{i-1}) - \pi(v_i)$, $i = 1, \dots, k$ (cf. (2.5)). Therefore, the integrality of p and a and the half-integrality of π (as γ is half-integral) imply that the numbers $2\pi(v_i)$ are integers and have the same parity, i.e., $2\pi(v_1) \equiv \dots \equiv 2\pi(v_{k-1}) \equiv 2\pi(v_0) \pmod{2}$. Now $\gamma(e) = l(e) - a(e)$ and $l(e) \in \{|\pi(v_0) - \pi(v_1)|, p - \pi(v_0) - \pi(v_1)\}$ show that $\gamma(e)$ is an integer, whence $\epsilon_1 \in \frac{1}{2}\mathbb{Z}$.

The half-integrality of ϵ_3 follows from the integrality of $d(u, v) - \bar{l}(u, v)$ for each pair (u, v) with $\lambda(u, v) = 2$ (cf. (5.13)–(5.14)). The latter is shown by arguments similar to those for ϵ_1 . Finally, ϵ_2 is a half-integer since p is an integer and π is half-integral (see (5.11)). \square

7. Cost scaling algorithm. It is based on ideas and tools similar to those of the capacity scaling algorithm. As before, a is assumed to be positive. At the high level, the algorithm consists of at most $I|EG|$ big iterations, where $I = \max\{\lceil \log_2 a(e) \rceil : e \in EG\}$. For $i = 0, 1, \dots, I$ and $e \in EG$, define $a^i(e) = \lceil a(e)/2^i \rceil$. Then $a^0(e) = a(e)$, $a^I(e) = 1$, $a^i(e) > 0$ and $2a^{i+1}(e) - a^i(e) \in \{0, 1\}$.

For $i = 0, \dots, I-1$, let U^i be the set of edges e with $a^i < 2a^{i+1}$, and let these edges be $z_1, \dots, z_{k(i)}$ ($k(i) = |U^i|$). For $j = 0, \dots, k(i)$, define $a_j^i(e) = a^i(e)$ ($= 2a^{i+1}(e) - 1$) for $e = z_1, \dots, z_j$, and $a_j^i(e) = 2a^{i+1}(e)$ for the other edges e of G . In particular, $a_0^i = 2a^{i+1}$.

Let p' be a large enough positive integer (for our purposes, it suffices to assume that $p' \geq 2|EG|c(EG)$), and let $p^i = 2^{I-i}p'$. In the beginning we find optimal primal and dual solutions f^I and γ^I for G, T, c, a^I, p' , using the pseudo-polynomial combinatorial algorithm from Karzanov (1979). Since $a^I(EG) = |EG|$, this takes strongly polynomial time. Note that f^I is half-integral, while γ^I may not be so (the above-mentioned algorithm does not guarantee the dual half-integrality). Nevertheless, Karzanov (1994) suggests a strongly polynomial combinatorial procedure that transforms an arbitrary optimal dual vector to a half-integer one. So we may assume that γ^I is half-integral as well.

The input of a current big iteration consists of optimal primal and dual solutions f_{j-1}^i and γ_{j-1}^i for G, T, c, a_{j-1}^i, p^i , and the goal is to transform them into optimal f_j^i and γ_j^i for G, T, c, a_j^i, p^i . Here we put $f_0^i = f^{i+1}$ and $\gamma_0^i = 2\gamma^{i+1}$, where f^{i+1} and γ^{i+1} are optimal solutions found for a^{i+1} and p^{i+1} . Hence, the last big iteration finds optimal solutions for a and $p^0 = 2^I p'$ (thus solving (1.1)).

So we come to the following auxiliary problem, to be solved at most $I|EG|$ times; for convenience we keep the original notation.

(7.1) Given optimal half-integral f and γ for G, T, c, \hat{a}, p , and an edge z , find optimal half-integral f' and γ' for G, T, c, a, p , where a coincides with \hat{a} everywhere except for z for which $a(z) = \hat{a}(z) - 1 > 0$.

First of all we examine complementary slackness conditions (2.3) and (2.4) for a, γ, f . If (2.3) is violated, we restore it by updating $\gamma(z) := \gamma(z) + 1$ (thus compensating the decrease of the cost of z by one). Therefore, without loss of generality we

may assume that (2.3) is valid; also we may assume that z belongs to a T -line for a, γ, p (otherwise $\gamma(z)$ should be duly decreased). If (2.4) is violated at z , we have a situation similar to that in §3. The only difference is that now $c(z) - \zeta^f(z)$ may be an arbitrary positive half-integer, while $\gamma(z)$ is exactly one (as (2.4) holds for the initial γ). Again our aim is either to saturate z or to reduce $\gamma(z)$ to zero. We act in a way similar to that in §§4 and 5 for solving (3.1). In what follows we outline only the differences between two algorithms, leaving it to the reader to check the necessary details in full.

We construct Γ, D, H, h as in §§2 and 3 and attempt to increase h at $\bar{z} = (q, r)$ and/or $\bar{z}' = (r', q')$ (where $\{\bar{z}, \bar{z}'\} = \omega^{-1}(z)$). More precisely, first we try to change the circulation h on feasible arcs plus \bar{z} and \bar{z}' so as to get an integer circulation h' with $h'(\bar{z})$ as large as possible provided that $h'(\bar{z}') \geq h(\bar{z}')$ (therefore, our primal update corresponds to several primal updates from §4). This is done in strongly polynomial time by use of one of numerous strongly polynomial maximum flow algorithms (for a survey, see, e.g., Ahuja et al. (1993), Goldberg et al. (1990)). Then we change h' in a similar way to get a circulation h'' with $h''(\bar{z}')$ maximum provided that $h''(\bar{z}) = h'(\bar{z})$.

Thus, we may assume that h is subject to one of the situations in (4.3) (with $h(\bar{z}')$ not necessarily equal to $c(z)$ in (4.3)(i)). Accordingly, define R, X and M as in §4. Obviously, (4.5) and Statements 4.2 and 4.3 remain valid for our case. Next the algorithm makes the dual update with respect to X as described in §5. It results in a new function $\gamma' = \gamma^\epsilon$ for which $\gamma'(z) < \gamma(z)$. By Statement 6.5 (which remains valid), γ' is half-integral; therefore, $\gamma'(z)$ is 0 or $1/2$ (as $\gamma(z) = 1$). If $\gamma'(z) = 0$, we are done, and if $\gamma'(z) = 1/2$, the algorithm continues with either the next dual update (solving (7.1)) or the primal update as above. In the latter case the total multiflow at z increases, and if z still remains nonsaturated, one more dual update is applied, which eventually solves (7.1). Thus, we obtain the following.

STATEMENT 7.1. The algorithm for (7.1) consists of at most two primal updates and at most two dual updates. Each update is executed in strongly polynomial time. \square

Moreover, the final multiflow f' and dual vector γ' are half-integral, as required. Summing up the above arguments we conclude that the whole cost scaling algorithm solves (1.2) for a rather large p and runs in polynomial time as indicated in the Introduction.

Acknowledgement. We are thankful to the referees for meticulously reading the original version of this paper, correcting errors and numerous suggestions intended to improve and simplify the text.

Part of this work was done while the first author was at the Computer Science Department, Stanford University, and supported by ONR Young Investigator Award N00014-91-J-1855, NSF Presidential Young Investigator Grant CCR-8858097 with matching funds from AT & T, DEC, and 3M, and a grant from Powell Foundation.

The original version was done while the second author was visiting Stanford University in 1993 and supported by the above mentioned NSF and Powell Foundation Grants. The revised version was supported by European Union Grant INTAS-93-2530-ext.

References

- Ahuja, R. K., T. L. Magnanti, J. B. Orlin (1993). *Network Flows: Theory, Algorithms and Applications*, Prentice Hall, New York.
- Bland, R. G., D. L. Jensen (1992). On the computational behavior of a polynomial-time network flow algorithm. *Math. Programming* **54** 1–41.

- Cherkassky, B. V. (1977). A solution of a problem on multicommodity flows in a network. *Ekonomika i Matematicheskie Metody* **13** (1) 143–151 (Russian).
- Edmonds, J., R. M. Karp (1972). Theoretical improvements in algorithmic efficiency for network flow problems. *J. ASM* **19** 248–264.
- Ford, L. R., D. R. Fulkerson (1962). *Flows in Networks*, Princeton Univ. Press, Princeton, New Jersey.
- Goldberg, A. V., A. V. Karzanov (1993). *Transitive fork environments and minimum cost multiflows*, Technical Report STAN-CS-93-1476, Stanford University, Stanford, California.
- , E. Tardos, R. E. Tarjan (1990). Network flow algorithms, B. Korte, L. Lovász, H. J. Prömel, A. Schrijver, eds., *Paths, Flows, and VLSI-Layout*, Springer, Berlin, 101–164.
- Karzanov, A. V. (1979). A minimum cost maximum multiflow problem. *Combinatorial Methods for Flow Problems*, Inst. for System Studies, Moscow, iss. 3, 138–156 (Russian).
- (1987). Unbounded fractionality in the maximum and minimum-cost multicommodity flow problems. *Problems of Discrete Optimization and Methods for Their Solution*, Central Economic and Mathematical Inst., Moscow, 123–135 (Russian); English translation in *Amer. Math. Soc. Translations, Ser. 2*, **158** (1994) 71–80.
- (1993). *Edge-disjoint T-paths of minimum total cost*, Technical Report STAN-CS-92-1465, Stanford University, Stanford, California.
- (1994). Minimum cost multiflow in undirected networks. *Math. Programming* **66** (3) 313–325.
- Lomonosov, M. V. (1985). Combinatorial approaches to multiflow problems. *Discrete Appl. Math.* **11**(1) 1–94.
- Lovász, L. (1976). On some connectivity properties of Eulerian graphs. *Acta Math. Akad. Sci. Hung.* **28** 129–138.
- Röck, H. (1980). Scaling techniques for minimal cost network flows. U. Pape, ed., *Discrete Structures and Algorithms*, Carl Hansen, Munich, 181–191.

A. V. Goldberg: NEC Research Institute, Inc., 4 Independence Way, Princeton, New Jersey 08540; e-mail: avg@research.nj.nec.com

A. V. Karzanov: Institute for System Analysis of the Russian Academy of Sciences, 9, Prospect 60 Let Oktyabrya, 117312 Moscow, Russia