MATHEMATICS OF OPERATIONS RESEARCH
Vol 16, No 1, February 1991
Printed in U S A

MAX-BALANCING WEIGHTED DIRECTED GRAPHS AND MATRIX SCALING*

HANS SCHNEIDER[†] AND MICHAEL H. SCHNEIDER[‡]

A weighted directed graph G is a triple (V, A, g) where (V, A) is a directed graph and g is an arbitrary real-valued function defined on the arc set A. Let G be a strongly-connected, simple weighted directed graph. We say that G is *max-balanced* if for every nontrivial subset of the vertices W, the maximum weight over arcs leaving W equals the maximum weight over arcs entering W. We show that there exists a (up to an additive constant) unique potential p_v for $v \in V$ such that (V, A, g^p) is max-balanced where $g_a^p = p_u + g_a - p_v$ for $a = (u, v) \in A$. We describe an $O(|V|^2|A|)$ algorithm for computing p using an algorithm for computing the *maximum cycle-mean* of G. Finally, we apply our principal result to the similarity scaling of nonnegative matrices.

1. Introduction. Let (V, A) be a directed graph (or simply a graph) with vertex set V and arc set A. We will use the notation $a = (u, v)$ to denote the arc a from vertex u to vertex v. The graph (V, A) may contain multiple arcs from u to v, although we will rule this out shortly. (Strictly speaking, we should write $a \in (u, v)$ where (u, v) is the set of all arcs from u to v.) A *weight function for* (V, A) is a real-valued function defined on the arcs A. A *potential for* (V, A) is a real-valued function defined on the vertices V. A *weighted graph* is a triple $G = (V, A, g)$ where (V, A) is a graph and g is weight function for (V, A). We will use g_a for $a \in A$ and p_v for $v \in V$ to denote the *weight of a* and the *potential of v*, respectively.

For a graph (V, A), a *cut for* (V, A) is a *nontrivial* subset W of V (i.e., $\varnothing \subset W \subset V$). (We will use the symbols $\subset$ and $\subseteq$ to denote strict and weak containment, respectively.) We define the set of arcs *leaving* W and the set of arcs *entering* W, written $\delta + (W; G)$ and $\delta^-(W; G)$, respectively, by

$$\delta^+(W; G) = \{a = (u, v) \in A | u \in W, \text{ and } v \in V \setminus W\}, \quad \text{and}$$

$$\delta^-(W; G) = \{a = (u, v) \in A | u \in V \setminus W, \text{ and } v \in W\}.$$

When there is no possibility of confusion, we will omit the dependence on G.

Let $G = (V, A, g)$ be a weighted graph, and let W be a cut for (V, A). Then G is *max-balanced at* W if

$$\max_{a \in \delta^+(W)} g_a = \max_{a \in \delta^-(W)} g_a,$$

that is, if the maximum weight over arcs leaving W equals the maximum weight over arcs entering W. (See Figure 1.) We define the maximum over the empty set to be $-\infty$. Further, G is *max-balanced* if G is max-balanced at every cut W. If G is max-balanced, we will also refer to the weight function g as max-balanced.

*Received December 21, 1987; revised March 21 and October 21, 1989.
AMS 1980 subject classification. Primary: 90C35. Secondary: 05C20
IAOR 1973 subject classification Main: Graphs
OR/MS Index 1978 subject classification Primary: 484 Networks/graphs/flows.
Key words. Directed graphs, weighted arcs, potentials.
[†]Research supported in part by NSF grants DMS 85-21521, ECSE 87-18971, and DMS 89-01445
[‡]Research supported in part by NSF grants ECS 85-04195 and ECSE 87-18971.

0364-765X/91/1601/0208/$01.25
Copyright © 1991, The Institute of Management Sciences/Operations Research Society of America

Copyright © 2001 All Rights Reserved

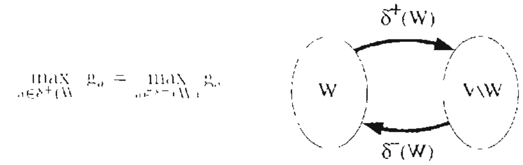

FIGURE 1. G is Max-balanced at Cut W.

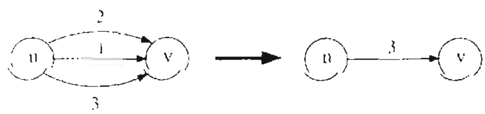

FIGURE 2. Identifying Parallel Arcs of G

Let p be a potential for (V, A). We defined the *reweighting of G with respect to p*, written G^p, to be the weighted graph (V, A, g^p) where

$$g_a^p = p_u + g_a - p_v \quad \text{for } a = (u, v) \in A. \tag{1}$$

A weight function f for (V, A) is a *reweighting of* g if $f = g^p$ for some potential p.

A graph (V, A) is *simple* if it contains no *loop* (i.e., an arc $a = (v, v)$) and no *parallel arcs* (i.e., multiple arcs from u to v). Let G be a weighted graph, and let H be the simple graph obtained from G be removing all loops and identifying parallel arcs. Thus, multiple arcs from u to v are replaced be a single arc whose weight equals the maximum weight over the arcs identified be forming (u, v). (See Figure 2.) It is easy to see that G is max-balanced if and only if H is max-balanced, and that a max-balanced graph must be the disjoint union of strongly-connected max-balanced graphs.

We are thus led to study the following problem:

PROBLEM 1 (The Max-Balancing Problem). *Given a strongly-connected, simple weighted graph (V, A, g), find a potential p, such that the reweighted graph $G^p = (V, A, g^p)$ is max-balanced.*

If p is a potential that solves Problem 1, we say that *p max-balances G*. We will show that the Max-Balancing Problem has a (up to additive constant) unique solution p and describe an $O(|V|^2|A|)$ algorithm for constructing p.

We now describe our paper in more detail. In §2 we recall analogous problems which have been studied and which motivate our investigations, and in §3 we present our notation and definitions. In §4 we define the operation of contraction which is used in our algorithm. In §5 we show that the max-balancing problem has at most one solution (up to an additive constant).

In §6 we show that there exists a reweighting of the weighted graph G such that each arc weight is less than or equal to the maximum cycle-mean of G (see [3]). We use a variant of Karp's algorithm [8] (see also [4, 5]) for finding the maximum cycle-mean. (Note, however, that any algorithm which computes the maximum cycle-mean could be used.) This algorithm will form the principal subroutine in our solution of the max-balancing problem.

In §7 we describe our algorithm for computing the potential p that max-balances G. Our algorithm constructs a sequence of weighted graphs

$$G = G_0 \to G_1 \to G_2 \to \cdots \to G_m,$$

where G^{i+1} is derived from G^i be reweighting and contracting a maximum-mean

Copyright © 2001 All Rights Reserved

cycle. The final term of this sequence is a singleton. At each iteration of the algorithm, we generate a potential σ' for G_t corresponding to a maximum-mean cycle of G_t. At the conclusion of the algorithm, the sum of the potentials σ' (suitable defined) computed at each iteration is a potential that max-balances G. The rest of §7 is devoted to a proof of this assertion. In §8 we apply our principal result to the similarity scaling of nonnegative matrices.

We remark that we consider strongly-connected, simple weighted graphs for the sake of simplicity of exposition. Our variant of Karp's algorithm finds the maximum cycle-mean of an arbitrary weighted graph G, and our algorithm (with slight modifications) can be shown to balance an arbitrary graph G at all cuts that are not the union of strong components of G. In particular our algorithm will max-balance all strongly-connected components of G (see [11] for a proof). A numerical example of the algorithm in the general case may be found in [11].

Max-balanced graph have also been studied in algebraic optimization under the name *algebraic flows*. See, for example, [1, 6, 7, 16]. Further results for max-balanced graphs can also be found in [10, 12].

2. Motivation. Let (V, A) be a graph. A nonnegative weight function g for (V, A) is a *circulation* if for every vertex v the sum of the weights over arcs entering v equals the sum of the weights over arcs leaving v. The following l_1 version of the max-balancing problem has been studied in [2, 9, 13, 14].

PROBLEM 2 (l_1-Balancing Problem). *Given a strongly-connected, simple weighted graph (V, A, g) with $g_a > 0$ for $a \in A$, find a potential $p > 0$ for (V, A) such that the weight function f defined by*

$$f_a = p_u g_a p_v^{-1} \quad \text{for } a = (u, v) \in A \tag{2}$$

is a circulation for (V, A).

The l_1-balancing problem occurs in economics, statistics, urban planning, and demography. For example, in development economics the weight function g represents an initial statistical estimate of the flow-of-funds between sectors of an economy. The circulation conditions are prescribed accounting identities requiring that after accounting for all transactions (including borrowing and saving) each sector's total receipts and expenditures must be equal. Since the data used to estimate the weights are incomplete, a numerical procedure must be used to modify the weights so that the initial estimates satisfy the accounting identities. The l_1-balancing problem is one approach for formulating this problem. See [15] for a discussion of the applications of l_1-balancing and related matrix balancing problems.

The l_1-balancing problem can be extended to l_ρ for $1 \leqslant \rho < \infty$ by requiring that for each vertex v the sum of the ρth powers of the weights over arcs entering and leaving v must be equal. It is not hard to see that for $1 \leqslant \rho < \infty$ the l_ρ-balancing problem can be reduced to the l_1-balancing problem for the weight function $g'_a = (g_a)^\rho$ for $a \in A$. The case of $\rho = \infty$, however, produces a significantly different problem which apparently cannot be reduced to the l_1 case. This is the problem we consider in this paper.

In the case of l_1-balancing, it is easy to see that if the weight function g is a circulation for (V, A), then for any cut W, the sum of the weights over arcs leaving W equals the sum of the weights over arcs entering W. Thus, if the circulation condition is satisfied with respect to single vertices, then the analogous circulation condition is also satisfied at every cut. This property is not satisfied in the extension to $\rho = \infty$. That is, a weight function that is max-balanced at singletons need not be max-balanced at larger cuts. (See Figure 3.)

Copyright © 2001 All Rights Reserved

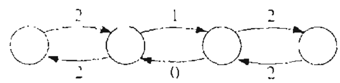

FIGURE 3. G is Max-balanced at Singletons, but Not Max-balanced.

Thus we are lead to study the following l_∞-balancing problem:

PROBLEM 3 (l_∞-Balancing Problem). *Given a strongly-connected, simple weighted graph (V, A, g) with $g_a > 0$ for $a \in A$, find a potential $p > 0$ for (V, A) such that for f defined by*

$$f_a = p_u g_a p_v^{-1} \quad \text{for } a = (u, v) \in A \tag{3}$$

the weighted graph (V, A, f) is max-balanced.

By taking logarithms in (3) it is straightforward to show that the l_∞-balancing problem is equivalent to the max-balancing problem. The original additive version (Problem 1) is more natural for presenting the algorithm described in §7. We will use the multiplicative version (Problem 3) in §8 when we apply our results to the similarity scaling of nonnegative matrices.

3. Notation and definitions. Let (V, A) be a graph, and let v_0 and v_k be vertices of (V, A). A *path from* v_0 to v_k of (V, A) is a sequence of the form $P = (v_0, a_1, v_1, \ldots, a_k, v_k)$ in which $a_i = (v_{i-1}, v_i)$ for $i = 1, 2, \ldots, k$. That is, a path is directed and may contain repeated arcs (or vertices). The path P is said to *start* and *end* at the vertices v_0 and v_k, respectively. We will identify a path with its underlying arc set. In particular, the *length* of a path P is the number of arcs of P and is denoted by $|P|$. A (*simple*) *cycle* is a path containing at least one arc that starts and ends at the same vertex and contains no repeated vertices. The set of all cycles of a graph (V, A) is denoted by cycles (G).

Let $G = (V, A, g)$ be a weighted graph. For a subset E of the arcs A we define the *weight of* E, written $g(E)$, by

$$g(E) = \sum_{a \in E} g_a.$$

In particular, for a cycle or path C, the weight of C is $g(C) = \sum_{a \in C} g_a$. For a cycle C in G, we define the *mean of* C, written $\bar{g}(C)$, by

$$\bar{g}(C) = \frac{1}{|C|} \sum_{a \in C} g_a.$$

We define the *maximum cycle-mean* for G, written $\text{mcm}(G)$, by

$$\text{mcm}(G) = \max\{\bar{g}(C) | C \in \text{cycles}(G)\}.$$

Note, $\text{mcm}(G) = -\infty$ if and only if G is acyclic (i.e., (V, A) contains no (directed) cycle). A cycle C of G is a *maximum-mean cycle* if

$$\bar{g}(C) = \text{mcm}(G).$$

Two vertices u and v are *connected* if there is a path from u to v and a path from v to u. Connectedness induces an equivalence relation on the set of vertices; the resulting equivalence classes are called the *strong components of* (V, A). We call a graph *strongly-connected* if it has exactly one strong component.

Copyright © 2001 All Rights Reserved

4. The operation of contraction. Let $G = (V, A, g)$ be a simple weighted graph. We need to define the graph derived from G by contraction. Our convention is that the operation of contraction is defined only for a graph whose vertex set is a partition of some underlying set. This is necessary for the consistency of the sequence of graphs generated by contraction in our max-balancing algorithm described in §7.

Let Π and Π' be partitions of a set V. Then we say that Π' is *coarser than* Π if every element of Π' can be expressed as the union of elements of Π. Let $G = (\Pi, A, g)$ be a weighted graph, and let Π' be coarser than Π. We define the *contraction of G with respect to* Π', written G/Π', to be the weighted graph (Π', A', g') where A' is the set of all $a' = (I', J')$ such that $I', J' \in \Pi', I' \neq J'$, $I \subseteq I'$ and $J \subset J'$ for some $(I, J) \in A$, and for $a' = (I', J')$

$$(4) \qquad g'_{a'} = \max\{g_a \mid a = (I, J) \in A, I \subseteq I', \text{and } J \subseteq J'\}.$$

The definition of A' ensures that this maximum is taken over a nonempty set. Intuitively, G' is derived from G by identifying all vertices of Π contained in the same element of Π'. Then all loops are removed and parallel arcs are identified. The weight function g' is derived by *max-projecting* g onto A'.

Restricting the operation of contraction to weighted graphs whose vertex sets are partitions is without loss of generality. For an arbitrary weighted graph (V, A, g), we define the *discrete partition of* V, written $\Pi(V)$, by

$$\Pi(V) = \{\{v\} | v \in V\}.$$

By identifying the element v of V and the element $\{v\}$ of $\Pi(V)$, there is an obvious graph isomorphism between (V, A, g) and $(\Pi(V), A, g)$.

In our algorithm, we shall consider the important case in which the partition Π' is induced by a cycle C to G. That is, one element of Π' is the set of vertices of C, and the others are the remaining elements of Π. In this case we denote the contracted graph by G/C, and refer to G/C as the *graph derived from G by contraction C to a point*.

Let $G = (V, A, g)$ and let $G_0 = (\Pi, A, g)$ be the isomorphic weighted graph in which Π is the discrete partition of V. Then for a sequence of weighted graphs

$$G \equiv G_0 \to G_1 \to G_2 \to \cdots \to G_m,$$

where G_{k+1} is constructed from G_k by contraction, the vertex set of each graph is a partition of V and is *coarser* than the preceding term.

5. Uniqueness. In this section we show that for a strongly-connected weighted graph G there is at most one (up to an additive constant) potential that max-balances G.

THEOREM 1. *Let $G = (V, A, g)$ be a strongly-connected weighted graph. If p and q are potentials for (V, A) that max-balance G, then, for some constant α, $p_v - q_v = \alpha$ for all $v \in V$. Therefore, the weight function g^p is the unique max-balanced reweighting of g.*

PROOF. Let p and q be potentials that max-balanced G, and let r be the potential defined by $r_v = p_v - q_v$. It is easy to see that $g^p = (g^q)^r$; that is,

$$(5) \qquad g_a^p = r_u + g_a^q - r_v \quad \text{for } a = (u, v) \in A.$$

Copyright © 2001 All Rights Reserved

We define $W \subseteq V$ by

$$W = \left\{ w \in V \mid r_w = \max_{v \in V} r_v \right\}.$$

It suffices to show that $W = V$. If not, then because G^p is max-balanced it follows from (5) that

$$\max_{\substack{a \in \delta^+(W) \\ a=(u,v)}} \{ r_u + g_a^q - r_v \} = \max_{\substack{a \in \delta^-(W) \\ a=(u,v)}} \{ r_u + g_a^q - r_v \}. \tag{6}$$

Note that both $\delta^+(W)$ and $\delta^-(W)$ are nonempty since (V, A) is strongly-connected. But since $r_u - r_v > 0$ for $a = (u, v) \in \delta^+(W)$, and $r_u - r_v < 0$ for $a = (u, v) \in \delta^-(W)$, line (6) contradicts the assumption that g^q is max-balanced. □

6. Computing maximum-mean cycles. The principal subroutine used by our max-balancing algorithm computes the maximum cycle-mean (mcm(G)) of a weighted graph $G = (V, A, g)$. Given mcm(G), we can find a potential p for (V, A) with the property that in the reweighted graph G^p, every arc has weight no larger than mcm(G). The following lemma shows that such a potential exists (see also, [3, Theorem 7.5]).

THEOREM 2. *Let $G = (V, A, g)$ be a weighted graph containing a cycle, and let $H = (V, A, g - \text{mcm}(G))$ be the weighted graph in which the arc weights are shifted down by mcm(G). For each $v \in V$, let p_v be the maximum weight over all paths of H ending at v. (Note, the length and starting point are arbitrary.) Then*

$$p_u + g_a - p_v \leqslant \text{mcm}(G) \quad \textit{for every } a = (u, v) \in A. \tag{7}$$

Further, if r is any potential satisfying (7) *and a is an arc contained in some maximum-mean cycle for G, then*

$$g_a^r = r_u + g_a - r_v = \text{mcm}(G). \tag{8}$$

PROOF. Since H has no positive cycles, p_v is finite for each $v \in V$. Since $a = (u, v)$ extends any path ending at u to a path ending at v, it follows directly from the definition of p_v that

$$p_u + g_a - \text{mcm}(G) \leqslant p_v \quad \text{for each } a = (u, v) \in A,$$

and (7) follows. (See Figure 4.)

Let C be a maximum-mean cycle for G, and let r be a potential satisfying (7). Since $g(C) = g^r(C)$, we have

$$\text{mcm}(G) = \bar{g}(C) = \bar{g}^r(C).$$

Now (8) follows directly from (7). □

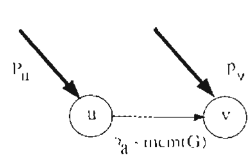

FIGURE 4. $p_u + g_a - p_v \leqslant \text{mcm}(G)$.

Copyright © 2001 All Rights Reserved

We call a potential satisfying (7) an *optimal potential for G*. The optimal potential p defined in Theorem 2 can be computed in time $O(|V||A|)$ as follows: Let $n = |V|$; for $k = 0, 1, \dots, n$ and $v \in V$, let $F_k(v)$ be the maximum weight over the set of paths of length k ending at v (in the original graph (V, A, g)). The $F_k(v)$'s can be computed using the recurrence

$$F_0(v) = 0 \quad \text{for } v \in V, \quad \text{and}$$

$$F_{k+1}(v) = \max_{\substack{a \in \delta^-(v) \\ a = (u,v)}} \{F_k(u) + g_a\} \quad \text{for } k = 0, 1, 2, \dots, n-1.$$

Then $\text{mcm}(G)$ is given by

$$\text{mcm}(G) = \max_{v \in V} \left\{ \min_{0 \le k \le n-1} \left\{ \frac{F_n(v) - F_k(v)}{n-k} \right\} \right\}. \tag{9}$$

Now the optimal potential p in Theorem 2 can be computed by

$$p_v = \max_{0 \le k \le n-1} \{F_k(v) - k \cdot \text{mcm}(G)\} \quad \text{for } v \in V. \tag{10}$$

The method described here is a modification of the maximum cycle-mean algorithm described in Karp [8]. Specifically, Karp assumes that the graph G is strongly-connected and defines the $F_k(v)$'s as the maximum weight over paths ending at v from some fixed vertex. In our modification we define the $F_k(v)$'s as the maximum weights over all paths ending at v, thereby extending Karp's algorithm to arbitrary weighted graph. Karp's proof that the $\text{mcm}(G)$ is given by (9) in the strongly-connected case extends to this more general setting (see [11] for the details).

7. The balancing algorithm. In each iteration of our max-balancing algorithm we compute the maximum cycle-mean $\text{mcm}(G)$ and a maximum-mean cycle C for a graph $G = (V, A, g)$. Using these we can compute an optimal potential for G, that is, a potential with the property that in the reweighted graph G^p the weights on the arcs of C are equal to $\text{mcm}(G)$ and the weights on the remaining arcs are no larger that $\text{mcm}(G)$ (Theorem 2). We then contract the cycle C to a point (in the reweighted graph) and repeat the iteration. Since each contraction operation decreases the number of vertices, the algorithm terminates after at most $|V|$ iterations.

At each iteration, the vertex set is a partition of V. Thus, for $v \in V$ we consider the element of each partition containing v. We define a potential for the original graph G be adding up the optimal potentials computed at each iteration evaluated at the element of the partition containing v. We show in Theorem 6 that the resulting potential max-balances G.

The max-balancing algorithm

Input: A strongly-connected, simple weighted graph $G = (V, A, g)$.

Output: A potential p for (V, A) such that the reweighted graph G^p is max-balanced, and an integer m equal to the number of iterations of the algorithm. (Recall, the weight of $a = (u, v)$ in G^p is $p_u + g_a - p_v$.)

0: (Initialization) Let $G_0 = (\Pi^0, A^0, g^0)$ where Π^0 is the discrete partition of V, $A^0 = A$, and $g^0 = g$. Set $k = 0$.

1: (Termination) If G_k is a singleton, set $m = k$, and go to Step 5.

2: (Compute Maximum-Mean Cycle) Find a maximum-mean cycle C^k and corresponding optimal potential σ^k for G_k. (See Theorem 2.) Set Π^{k+1} equal to the partition induced by C^k. (See §4.)

Copyright © 2001 All Rights Reserved

3: (Reweight and Contract G_k) Let G_{k+1} be the weighted graph formed by reweighting G_k with respect to the potential σ^k and contracting the cycle C^k to a point. That is,

$$G_{k+1} = (G_k)^{\sigma^k}/C^k.$$

Thus, arc $a' = (I', J') \in A^{k+1}$ has weight

$$g_{a'}^{k+1} = \max\{\sigma_I^k + g_a^k - \sigma_J^k | a = (I, J) \in A^k, I \subseteq I', \text{and } J \subseteq J'\}.$$

4: (Increment) Set $k = k + 1$, and return to Step 1.

5: (Compute Max-Balancing Potential) For $v \in V$ and $k = 0, 1, 2, \ldots, m - 1$, let $I(v, k)$ be the element of Π^k containing v. Define the potential p for (V, A) by

$$(11) \qquad p_v = \sum_{k=0}^{m-1} \sigma_{I(v,k)}^k \quad \text{for } v \in V.$$

Return p and m; STOP.

It is easy to see that if a weighted graph G is strongly-connected and simple then so is any contraction of G. Therefore, the weighted graphs generated by the max-balancing algorithm are strongly-connected and simple. Further, at any iteration if the algorithm does not terminate in Step 1, the graph G_k must contain a cycle. Note that we use the notation g^k to denote the weight function of the graph G_k. Strictly speaking, we should write $g^{(k)}$ to distinguish this from g^p which denotes the reweighting of g with respect to the potential p. No confusion should result, however, since k will always be an index.

THEOREM 3. *Let $G = (V, A, g)$ be a strongly-connected, simple weighted graph. Then the max-balancing algorithm terminates after at most $|V|$ contraction-reweighting operations.*

PROOF. After each contraction operation $|\Pi^{k+1}| > |\Pi^k|$ since a cycle has at least two arcs. □

LEMMA 4. *Let $G = (V, A, g)$ be a strongly-connected, simple weighted graph, and let $G_k = (\Pi^k, A^k, g^k)$ for $k = 0, 1, \ldots, m$ be the sequence of weighted graphs produced by the max-balancing algorithm. Then*

(i) $g_a^{k+1} \leqslant \text{mcm}(G_k)$ *for* $a \in A^{k+1}$ *and* $k = 0, 1, \ldots, m - 1$, *and*

(ii) $\text{mcm}(G_0) \geqslant \text{mcm}(G_1) \geqslant \cdots \geqslant \text{mcm}(G_{m-1})$.

PROOF. Part (i) follows directly from Theorem 2 since reweighting G_k by σ^k decreases all arcs' weights below $\text{mcm}(G_k)$. Part (ii) follows directly from part (i) since $\text{mcm}(G_{k-1})$ is an average of arc weights from G_{k+1}. □

For a given strongly-connected, simple weighted graph $G = (V, A, g)$, let Π^k for $k = 0, 1, \ldots, m$ be the sequence of partitions of V produced by the max-balancing algorithm, and let $I(v, k)$ be the element of Π^k containing v (as in Step 5 of the max-balancing algorithm). For the graph (V, A) we define the potentials $p^k, k = 0, 1, 2, \ldots, m$, by $p^0 = 0$, and

$$(12) \qquad p_v^k = \sum_{l=0}^{k-1} \sigma_{I(v,l)}^l \quad \text{for } v \in V, \text{ for } k = 1, 2, \ldots m.$$

Note that the p^k's are the partial sums in (11).

Copyright © 2001 All Rights Reserved

The next lemma contains technical results that are needed to prove correctness of the max-balancing algorithm.

LEMMA 5. *Let $G = (V, A, g)$ be a strongly-connected, simple weighted graph. Let $G_k = (\Pi^k, A^k, g^k)$ and p^k for $k = 0, 1, \ldots, m$ be, respectively, the weighted graphs produced by the max-balancing algorithm and the potentials defined in* (12). *Then the following are true*:

(i) *If $u, v \in V$ are in the same element of Π^j, $j = 0, 1, \ldots, m$, then*

$$p_u^k - p_v^k = p_u^j - p_v^j \quad \text{for } k = j, j+1, \ldots, m.$$

(ii) *If I and J are distinct elements of Π^{j-1}, $j = 1, 2, \ldots, m$, then for $a' = (I, J)$*

$$g_{a'}^{j-1} = \max\{p_u^{j-1} + g_a - p_v^{j-1} | u \in I, v \in J, \text{ and } a = (u, v) \in A\}, \tag{13}$$

and

$$\sigma_I^{j-1} + g_{a'}^{j-1} - \sigma_J^{j-1} = \max\{p_u^j + g_a - p_v^j | u \in I, v \in J, \text{ and } a = (u, v) \in A\}. \tag{14}$$

PROOF. (i) If u and v are in the same element of Π^j then they are in the same element of Π^k for $k \geqslant j$ since Π^k is coarser than Π^j. Therefore $I(u, k) = I(v, k)$ for $k \geqslant j$, and part (i) follows directly from (12).

(ii) Line (13) follows by induction and the definition of contraction. Line (14) follows from (13) and the definition of p^j in (12). □

Next, we provide additional explanation of the results in Lemma 5. Consider the weight functions g^{p^k} for (V, A) derived from g by reweighting with respect to the potential p^k defined in (12). For a given arc $a = (u, v) \in A$, consider the sequence of weights $g_a^{p^k}$ for $k = 0, 1, \ldots, m$. It follows from part (i) that as soon as the endpoints u and v are in the same vertex of G_j, then the weight of a in this sequence is constant. We will use this observation in the proof of Theorem 6. Since the maximum cycle-mean calculations are performed on the contracted graphs G^k, to prove that $p = p^m$ max-balanced (V, A), it is necessary to relate the weight function g^{p^k} for (V, A) with the weight function g^k of the contracted graph G_k. It follows from part (ii) that g^k can be derived by (13) from g^{p^k}. Line (14) is the relation that we will actually use in Theorem 6.

We need the following definition in the proof of Theorem 6. For a graph (V, A), let Π and W be, respectively, a partition and a cut. We say that W is *compatible with* Π if for every element I or Π, either $I \subseteq W$ or $I \subseteq V \setminus W$. Equivalently, W is *compatible with* Π if W can be written as the union of elements of Π. We can now state and prove the main theorem of our paper.

THEOREM 6. *Let $G = (V, A, g)$ be a strongly-connected, simple weighted graph, and let p be the potential for (V, A) produced by the max-balancing algorithm. Then G^p is max-balanced.*

PROOF. Let W be a cut for (V, A). We must show that

$$\max_{\substack{a \in \delta^+(W) \\ a = (u,v)}} \{p_u + g_a - p_v\} = \max_{\substack{a \in \delta^-(W) \\ a = (u,v)}} \{p_u + g_a - p_v\}.$$

Let p^k for $k = 0, 1, \ldots, m$ be the potentials defined in (12), and consider G^{p^k}, the reweighting of G with respect to p^k.

Copyright © 2001 All Rights Reserved

Intuitively, our proof technique is as follows: We run the max-balancing algorithm until the first time, say j, that the computed maximum-mean cycle C^{j-1} contains vertices from both W and $V \setminus W$. Then we show that the reweighting operation forces the weights on arcs of G^{p^j} leaving and entering W to lie below $\mathrm{mcm}(G_{j-1})$. Since C^{j-1} must enter and leave W, it follows that G^{p^j} is max-balanced at W. Since the contraction operation freezes the weight of at least one arc from $\delta^+(W)$ and $\delta^-(W)$ and since the computed maximum cycle-means are decreasing, it follows that G^{p^k} must remain max-balanced at W at subsequent iteration. Since $p = p^m$, it follows that G^p is max-balanced. We now give the formal argument.

Let $G_k = (\Pi^k, A^k, g^k)$ for $k = 0, 1, \dots, m$ be the sequence of graphs generated by the max-balancing algorithm with input G. Define j, $0 > j \leqslant m$, to be the smallest integer such that W is not compatible with Π^j. There must be one since Π^0 and Π^m are, respectively, the discrete and indiscrete partitions of V.

Claim 1. G^{p^j} is max-balanced at W.

We define $W' \subseteq \Pi^{j-1}$ by

$$W' = \{I \in \Pi^{j-1} | I \subseteq W\}.$$

Note that W' is a cut for G^{j-1} since W is a compatible with Π^{j-1}. Since W in not compatible with Π^j, however, the cycle C^{j-1} computed by the max-balancing algorithm must contain arcs from both $\delta^+(W'; G_{j-1})$ and $\delta^-(W'; G_{j-1})$. Since σ^{j-1} is an optimal potential for G_{j-1} it follows from Theorem 2 that G_{j-1} reweighted by σ^{j-1} is max-balanced at W'. That is

$$(15) \qquad \mathrm{mcm}(G_{j-1}) = \max_{\substack{a' \in \delta^+(W'; G_{j-1}) \\ a' = (I, J)}} \{\sigma_I^{j-1} + g_{a'}^{j-1} - \sigma_J^{j-1}\}$$

$$= \max_{\substack{a' \in \delta^-(W', G_k) \\ a' = (I, J)}} \{\sigma_I^{j-1} + g_{a'}^{j-1} - \sigma_J^{j-1}\}.$$

Furthermore, both maxima in (15) are attained at some arc of C^{j-1}.

Let $a = (u, v)$ be an arc of $\delta^+(W; G) \cup \delta^-(W; G)$. Since W is compatible with Π^{j-1} it follows that u and v are in distinct elements of Π^{j-1}. Therefore, combining Lemma 5, part (ii) and (15) we have

$$\mathrm{mcm}(G_{j-1}) = \max_{\substack{a \in \delta^+(W, G) \\ a = (u, v)}} \{p_u^j + g_a - p_v^j\} = \max_{\substack{a \in \delta^-(W, G) \\ a = (u, v)}} \{p_u^j + g_a - p_v^j\},$$

which shows that G^{p^j} is max-balanced at W.

Claim 2. G^{p^k} is max-balanced at W for $k = j + 1, j + 2, \dots, m$.

Let $a = (u, v)$ be any arc of $\delta^-(W; G) \cup \delta^+(W; G)$. If u and v are in distinct elements I and J of Π^{k-1}, then for $a' = (I, J)$

$$(16) \qquad p_u^k + g_a - p_v^k \leqslant \sigma_I^{k-1} + g_{a'}^{k-1} - \sigma_J^{k-1} \leqslant \mathrm{mcm}(G_{k-1}) \leqslant \mathrm{mcm}(G_{j-1}).$$

The first inequality follows from (14), the second from Theorem 2, and the third from Lemma 4, part (ii).

If u and v are in the same element of Π^{k-1}, then let i, $j \leqslant i < k$, be the smallest index such that u and v are in the same element of Π^i. Then

$$(17) \qquad p_u^k + g_a - p_v^k = p_u^i + g_a - p_v^i \leqslant \mathrm{mcm}(G_{i-1}) \leqslant \mathrm{mcm}(G_{j-1}).$$

Copyright © 2001 All Rights Reserved

The equality follows from Lemma 5, part (i); since u and v are in distinct elements of Π^{j-1}, the first inequality follows from (13) and Theorem 2; the second inequality follows from Lemma 4, part (ii).

Combining (16) and (17), we have

$$(18) \qquad p_u^k + g_a - p_v^k \leqslant \mathrm{mcm}(G_{j-1}) \quad \text{for } a \in \delta^+(W;G) \cup \delta^-(W;G).$$

Since the cycle C^{j-1} is contracted to a point when Π^j is formed in Step 3 of the max-balancing algorithm, for $k \geqslant j$ each set $\delta^+(W;G)$ and $\delta^-(W;G)$ contains an arc $a = (u,v)$ such that u and v are in the same element of Π^k and $p_u^j + g_a - p_v^j = \mathrm{mcm}(G_{j-1})$. It follows from Lemma 5, part (i), that for such arcs $p_u^k + g_a - p_v^k = \mathrm{mcm}(G_{j-1})$. Therefore, it follows from (18) that

$$\max_{\substack{a \in \delta^+(W;G) \\ a=(u,v)}} \{p_u^k + g_a - p_v^k\} = \max_{\substack{a \in \delta^-(W,G) \\ a=(u,v)}} \{p_u^k + g_a - p_v^k\} = \mathrm{mcm}(G_{j-1}).$$

This proves the claim.

The theorem follows from Claim 2, since $p = p^m$. □

A referee has proposed a different method for defining the potential σ^k in Step 2 of the max-balancing algorithm. Namely, let $G_k = (\Pi^k, A^k, g^k)$ be the graph in Step 2, and let C^k be any maximum-mean cycle of mean $\lambda^k = \mathrm{mcm}(G_k)$. Now let σ^k be any potential for G_k such that

$$(19) \qquad \sigma_I^k + g_a^k - \sigma_J^k = \lambda^k \quad \text{for } a = (I,J) \in C^k.$$

Such a σ^k can be computed as follows: Choose any $I \in C^k$ and set $\sigma_I^k = 0$. Then travel around C^k and set σ^k so that (19) is satisfied. The value of σ^k on the remaining vertices of G_k can be set to zero.

It is not hard to show that if σ^k is defined as above, then the computed cycle-means are decreasing (Lemma 4, part (i)). Moreover, we still have the result that for $a = (u,v) \in A$, if j is the first integer such that u and v are in the same element of the partition Π_j, then

$$p_u^j + g_a - p_v^j \leqslant \mathrm{mcm}(G_{j-1}).$$

Using these results, the proof of Theorem 6 can be adapted to show that the max-balancing algorithm with this modification is correct.

In the referee's modified approach one first finds a maximum mean cycle C and the corresponding cycle mean. One then computes a potential for reweighting. This approach is more general in the sense that it does not depend on a specific way of identifying C, and because it shows that the value of the potential at vertices which do not lie on C are irrelevant. Our approach, on the other hand, repeats a modified version of Karp's algorithm for identifying C. We thus clearly reveal the connections between our algorithm, the max balancing problem and the min-max characterization of the maximum-cycle mean in Theorem 2 and [3, Theorem 7.5], namely

$$\mathrm{mcm}(G) = \min_p \Big\{ \max_{a=(u,v) \in A} \{p_u + g_a - p_v\} \Big\}.$$

The two approaches are computationally equivalent.

8. Similarity scaling of nonnegative matrices. Let B be an $n \times n$ nonnegative matrix, and let $V = \{1, 2, \ldots, n\}$. For $I \subset V$ (recall, $\subset$ denotes strict containment) we

Copyright © 2001 All Rights Reserved

use I' to denote $V \setminus I$, the complement of I in V. The matrix B is called *max-balanced* if for every nontrivial $I \subset V$ we have

$$\max\{b_{ij} \mid i \in I, j \in I'\} = \max\{b_{ij} \mid i \in I', j \in I\}.$$

In the definition of max-balanced matrices, we define the maximum over the empty set to be 0, whereas in the definition of max-balanced graphs we define the maximum over the empty set to be $-\infty$. This is consistent with the log transformation used to convert the multiplicative matrix problem into the additive graph problem. (See §2.)

We define the *graph associated with* B, written Graph (B), to be the weighted graph (V, A, g) where

$$V = \{1, 2, \ldots, n\},$$

$$A = \{a = (i, j) \mid b_{ij} > 0\}, \quad \text{and}$$

$$g_a = \ln b_{ij} \quad \text{for } a = (i, j) \in A.$$

This correspondence induces a bijection between nonnegative matrices and weighted graphs with no parallel arcs.

An $n \times n$ matrix $D = (d_{ij})_{i,j=1}^n$ is called *diagonal* if $d_{ij} = 0$ whenever $i \neq j$. We denote a diagonal matrix by $\mathrm{diag}(d_1, d_2, \ldots, d_n)$ where d_i for $i = 1, 2, \ldots, n$ is the ith diagonal entry. A diagonal matrix is called *positive* if all of its diagonal entries are positive. Given $n \times n$ nonnegative matrices B and C, the matrix C is a *(diagonal) similarity scaling of* B if $C = DBD^{-1}$ for some positive diagonal matrix D.

There is a close connection between the operations of scaling nonnegative matrices and reweighting graphs. We state the following lemma without proof.

LEMMA 7. *Let B be an $n \times n$ nonnegative matrix, and let* $D = \mathrm{diag}(d_1, d_2, \ldots, d_n)$ *be an $n \times n$ positive diagonal matrix.*

(i) *The matrix B is max-balanced if and only if* Graph(B) *is max-balanced.*

(ii) *Let p be the potential for* Graph(B) *defined by $p_i = \ln d_i$ for $i = 1, 2, \ldots, n$. Then*

$$\mathrm{Graph}(DBD^{-1}) = \mathrm{Graph}(B)^p.$$

Let (V, A, g) be the graph associated with the nonnegative matrix B. Then B is called *irreducible* if (V, A) is strongly-connected. Also, B is called *completely reducible* if every arc of A is contained in a strong component of (V, A). We state the following lemma without proof.

LEMMA 8. *Let B be an $n \times n$ nonnegative matrix, and let (V, A, g) be the graph of B. Then the following are equivalent:*

(i) *The matrix B is completely reducible;*

(ii) *There exists a permutation matrix P such that $P^T BP$ is the direct sum of irreducible matrices;*

(iii) *For each cut W of (V, A), $\delta^+(W)$ is nonempty if and only if $\delta^-(W)$ is nonempty.*

Let G be the weighted graph obtained from Graph (B) by removing all loops. It is easy to see that a potential p max-balances G if and only p max-balances Graph (B). Thus we have the following corollary of Theorems 1 and 6.

COROLLARY 9. *Let B be an $n \times n$ nonnegative, irreducible matrix. Then there exists a unique (up to multiplicative constant) positive diagonal matrix D such that the similarity scaling DBD^{-1} of B is max-balanced.*

Copyright © 2001 All Rights Reserved

We have the following corollary of Corollary 9.

COROLLARY 10. *Let B be an $n \times n$ nonnegative matrix. Then there exists a similarity scaling of B that is max-balanced if and only if B is completely reducible.*

PROOF. ($\Leftarrow$) Suppose that B is completely reducible. Let $G = (V, A, g)$ be the graph associated with B, and let $\Pi = \{I_1, I_2, \ldots, I_\omega\}$ be the partition of V determined by the strong components of (V, A). Let G_i be the (strongly-connected) weighted graph induced by I_i, and let p^i be a potential that max-balances G_i. We define the potential q for (V, A) by

$$q_v = p_v^i \quad \text{for } v \in I_i \in \Pi. \tag{20}$$

It is easy to see that q max-balances G, since it follows from the definition of completely reducible that G is the union of the induced subgraphs G_i for $i = 1, 2, \ldots, \omega$.

($\Rightarrow$) Conversely, suppose that B is not completely reducible. Then by Lemma 8 there exists a cut W of Graph (B) such that $\delta^+(W) \neq \emptyset$ and $\delta^-(W) = \emptyset$. Clearly, no similarity scaling of B can be max-balanced at W. □

Next we show that Corollary 9 can be generalized to arbitrary nonnegative matrices. First, we need the following elementary lemma:

LEMMA 11. *Let $G = (V, A, g)$ be an acyclic weighted graph. Then for any scalar $M > 0$ there exists a potential σ for (V, A) such that*

$$g_a^\sigma \leqslant -M \quad \text{for all } a \in A. \tag{21}$$

PROOF. For $v \in V$, let r_v be the maximum length over all paths in (V, A) ending at v. For $\alpha > 0$, define the potential σ^α by $\sigma_v^\alpha = \alpha r_v$. Since $r_v \geqslant r_u + 1$ for $a = (u, v) \in A$, we have

$$g_a^{\sigma^\alpha} = \alpha r_u + g_a - \alpha r_v \leqslant g_a - \alpha.$$

Now we can choose α large enough so that (21) is satisfied. □

Let $G = (V, A, g)$ be a weighted graph, and let Π be the partition of V determined by the strong components of (V, A). We define the *condensed graph of* G, written condense (G), to be the weighted graph G/Π. It is easy to see that condense (G) is acyclic.

We can now state the main theorem of this section.

THEOREM 12. *Let B be a $n \times n$ nonnegative matrix, and let $G = (V, A, g)$ be the graph associated with B. Then for each $\epsilon > 0$ there exists an $n \times n$ positive diagonal matrix D such that the matrix $C = DBD^{-1}$ satisfies the following properties:*

(i) *For each nontrivial $I \subset V$ that is not the union of strong components of (V, A)*

$$\max\{c_{ij} \mid i \in I, j \in I'\} = \max\{c_{ij} \mid i \in I', j \in I\}.$$

(ii) *For each nontrivial $I \subset V$*

$$|\max\{c_{ij} \mid i \in I, j \in I'\} - \max\{c_{ij} \mid i \in I', j \in I\}| \leqslant \epsilon.$$

PROOF. Let $\Pi = \{I_1, I_2, \ldots, I_\omega\}$ be the partition of V determined by the strong components of (V, A), and for $i = 1, 2, \ldots, \omega$ let p^i be a potential that max-balances the subgraph G_i induced by the strong component I_i. Define the potential q for

Copyright © 2001 All Rights Reserved

(V, A) by

$$q_v = p_v^I \quad \text{for } v \in I, \in \Pi. \tag{22}$$

Let β be the minimum over the set of all arc weights of G^q, and define $\alpha = \min\{\beta, \ln \epsilon/2\}$. Let $H = \text{condense}(G^q)$, and let σ be a potential for H such that every arc of H^σ has weight less than or equal to β. (See Lemma 11.) Define the potential p for (V, A) by

$$p_v = q_v + \sigma_I \quad \text{for } v \in I \in \Pi, \tag{23}$$

and let $D = \text{diag}(d_1, d_2, \dots, d_n)$ where $d_i = e^{p_i}$.

We claim that $C = DBD^{-1}$ satisfies (i) and (ii). Let $a = (u, v) \in A$, and suppose that the vertices u and v are contained in the strong components J and K, respectively. If $J = K$, then if follows from (23) that

$$p_u + g_a - p_v = q_u + g_a - q_v \geqslant \beta. \tag{24}$$

If $J \neq K$, then (J, K) is an arc of H, and therefore

$$p_u + g_a - p_v = \sigma_J + q_u + g_a - q_v - \sigma_K \leqslant \alpha \leqslant \beta, \tag{25}$$

since $q_u + g_a - q_v$ is less than or equal to the weight of the arc (J, K) in H^σ.

To prove part (i), let I be a cut for (V, A) that is not the union of strong components of (V, A). We will show that

$$\max_{a \in \delta^+(I)} g_a^p \leqslant \max_{a \in \delta^-(I)} g_a^p. \tag{26}$$

Since I is not compatible with the partition Π, both $\delta^+(I)$ and $\delta^-(I)$ contain an arc whose endpoints are in the same strong component. Let $a = (u, v)$ be any arc of $\delta^+(I)$. If u and v are in distinct strong components, then it follows from (24) and (25) that there exists an arc $b \in \delta^-(I)$ such that $g_a^p \leqslant g_b^p$. If u and v are in the same strong component I_t, then since the subgraph of G^p induced by I_t is max-balanced, it follows that there exists $b \in \delta^-(I)$ such that $g_a^p \leqslant g_b^p$. This proves the inequality (26). A similar argument show that the opposite inequality holds as well; this proves part (i).

If I is the union of strong components of (V, A), then (25) holds for all arcs leaving and entering I. Therefore, part (ii) follows from (25) and part (i). □

Of course, Corollary 10 and Theorem 12 may be restated in an additive form for weighted graphs.

Acknowledgement. We wish to thank an anonymous referee for suggestions that significantly improved the presentation and exposition of our paper.

References

[1] Burkard, Rainer E., and Zimmerman, Uwe (1982). Combinatorial Optimization in Linearly Ordered Semimodules: A Survey. In Bernhard Korte, Ed., *Modern Applied Mathematics—Optimization and Operations Research*, North-Holland, Amsterdam.

[2] Eaves, B. Curtis. Hoffman, Alan J., Rothblum, Uriel G and Schneider, Hans (1985). Line-sum-symmetric Scalings of Square Nonnegative Matrices. *Math. Programming Studies* **25** 124–141.

[3] Engel, Gernot M. and Schneider, Hans. (1975). Diagonal Similarity and Equivalence for Matrices over Groups with 0. *Czechoslovak Math. J.* **25** 389–403.

[4] Golitschek. M. V (1982). Optimal Cycles in Doubly Weighted Graphs and Approximation of Bivariate Functions by Univariate Ones *Numer. Math.* **39** 65–84.

Copyright © 2001 All Rights Reserved

[5] Golitschek, M. V. and Schneider, Hans (1984). Applications of Shortest Path Algorithms to Matrix Scalings. *Numer. Math.* **44** 111–126.

[6] Hamacher, Horst (1980). Algebraic Flows in Regular Matroids. *Discrete Appl. Math.* **2** 27–38.

[7] ______. 1980. Maximal Algebraic Flows: Algorithms and Examples. In Uwe Pape, Ed., *Discrete Structures and Algorithms*, Carl Hauser Verlag, Munich, 1980. 153–166.

[8] Karp, Richard M (1978). A Characterization of the Minimum Cycle Mean in a Digraph. *Discrete Math.* **23** 309–311.

[9] Osborne, E. E. (1960). On Pre-conditioning of Matrices. *J. Assoc. Comput. Mach.* **7** 338–345

[10] Rothblum, Uriel G, Schneider, Hans and Schneider, Michael H. (1989). *Characterizations of Max-Balanced Flows.* OR Group Report Series #90-1, The Johns Hopkins University, Baltimore, MD 21218. Submitted to *Discrete Appl Math.*

[11] Schneider, Hans and Schneider, Michael H. (1987). An l_∞ balancing of a weighted directed graph. OR Group Report Series #88–1. Johns Hopkins University, Baltimore, Maryland.

[12] ______ and ______. (1990). Towers and Cycle Covers for Max-Balanced Graphs. *Congress Numer.* **73** (January) 159–170.

[13] Schneider, Michael H. (1989). Matrix Scaling Entropy Minimization, and Conjugate Duality (I): Positivity Conditions. *Linear Algebra Appls.* **114 / 115** 785–813, 1989.

[14] ______. (1990). Matrix Scaling, Entropy Minimization, and Conjugate Duality (II): The Dual Problem *Math. Programming Ser B* **46**.

[15] ______ and Zenios Stavros (1990) A Comparative Study of Algorithms for Matrix Balancing. *Oper. Res.* **383** (May/June) 439–455

[16] Zimmerman, Uwe (1981). *Linear and Combinatorial Optimization in Ordered Algebraic Structures.* Vol. 10 *Ann. Discrete Math.*. North-Holland Publishing Company, Amsterdam

HANS SCHNEIDER: DEPARTMENT OF MATHEMATICS, UNIVERSITY OF WISCONSIN, MADISON, WISCONSIN 53706

MICHAEL H. SCHNEIDER: DEPARTMENT OF MATHEMATICAL SCIENCES, JOHNS HOPKINS UNIVERSITY, BALTIMORE, MARYLAND 21218

Copyright © 2001 All Rights Reserved