May 1995

APPROXIMATE MINIMUM-COST MULTICOMMODITY FLOWS IN $\tilde{O}(\varepsilon^{-2}KNM)$ TIME

Michael D. Grigoriadis and Leonid G. Khachiyan

LCSR-TR-245

Department of Computer Science
Hill Center for the Mathematical Sciences
Rutgers University
New Brunswick, NJ 08903.

APPROXIMATE MINIMUM-COST MULTICOMMODITY FLOWS IN $\tilde{O}(\varepsilon^{-2}KNM)$ TIME*

Michael D. Grigoriadis and Leonid G. Khachiyan

Department of Computer Science, Rutgers University, New Brunswick, NJ, USA

May 26, 1995

We show that an ε-approximate solution of the cost-constrained K-commodity flow problem on an N-node M-arc network G can be computed by sequentially solving $O\left(K(\varepsilon^{-2}+\log K)\log M\log(\varepsilon^{-1}K)\right)$ single-commodity minimum-cost flow problems on the same network. In particular, an approximate minimum-cost multicommodity flow can be computed in $\tilde{O}(\varepsilon^{-2}KNM)$ running time, where the notation $\tilde{O}(\cdot)$ means "up to logarithmic factors". This result improves the time bound mentioned in Grigoriadis and Khachiyan (1994) by a factor of M/N and that developed recently in Karger and Plotkin (1995) by a factor of ε^{-1}. We also provide a simple $\tilde{O}(NM)$-time algorithm for single-commodity budget-constrained minimum-cost flows which is $\tilde{O}\left(\varepsilon^{-3}\right)$ times faster than the algorithm of Karger and Plotkin (1995).

Key words: Approximation algorithm, block-angular program, minimum-cost multicommodity network flow, structured optimization

1. Introduction. In this note we consider the general block-angular feasibility problem

$$\sum_{k=1}^{K} f^k(x^k) \le \mathbf{1}, \quad x^k \in B^k, \quad k=1,\dots,K, \tag{1.1}$$

where the B^k are given nonempty convex compact sets, the $f^k : B^k \to \mathbb{R}_+^{M+1}$ are continuous nonnegative component-wise convex vector functions, and $\mathbf{1}$ is the vector of all ones. We wish to compute an *ε-approximate solution* $(\tilde{x}^1,\dots,\tilde{x}^K)$ of (1.1) for a given relative tolerance $\varepsilon \in (0,1)$:

$$\sum_{k=1}^{K} f^k(\tilde{x}^k) \le (1+\varepsilon)\mathbf{1}, \quad \tilde{x}^k \in B^k, \quad k=1,\dots,K. \tag{1.2}$$

The *cost-constrained K-commodity flow problem* for a given N-node M-arc directed graph $G=(V,\mathcal{A})$ is a special case of (1.1) as follows: 1) $x^k=\{x_e^k,\ e\in\mathcal{A}\}$ is a flow and B^k is the flow polytope for commodity k defined by conservation, supply-demand, nonnegativity and individual capacity constraints; 2) $f_0^k(x^k)=c^{kT}x^k/C$, where $c^k=\{c_e^k,\ e\in\mathcal{A}\}$ is a given non-negative cost vector for commodity k, and $C>0$ is a desired upper bound on the total cost of $(x^1,\dots,x^K)$; 3) $f_e^k(x^k)=x_e^k/u_e$, where $u_e>0$ is a given mutual capacity for each arc $e\in\mathcal{A}$. An ε-approximate *minimum-cost* multicommodity flow can be computed by approximately solving a logaritmic number of cost-constrained problems (1.1).

* This research was supported by NSF Grant CCR-9208539.

Returning to the general problem, we can equivalently write (1.1) as

$$\sum_{k=1}^{K} \mathcal{F}^k(x^k) \leq \mathbf{1}, \quad x^k \in B^k, \quad k = 1, \dots, K, \tag{1.3}$$

where $\mathcal{F}^k(x^k) = (1-\delta)f^k(x^k) + \delta\mathbf{1}/K$ and $\delta = \varepsilon/(1+\varepsilon)$. It is easy to see that any ε-approximate solution of (1.1) is a δ-approximate solution of (1.3) and vice versa. Now, the nonnegativity of each $f^k(x^k)$ implies that $\mathcal{F}^k(x^k) \geq \delta\mathbf{1}/K$ for any $x^k \in B^k$. We shall henceforth assume that each $f^k(x^k)$ has been replaced by $\mathcal{F}^k(x^k)$ in (1.1) and the original ε replaced by $\varepsilon/(1+\varepsilon)$, so that

$$f^k(x^k) \geq \varepsilon\mathbf{1}/K, \quad \text{for any } x^k \in B^k \ \text{ and } \ k = 1, \dots, K. \tag{1.4}$$

The reason we prefer for our problem to satisfy this condition will become evident in the next section.

Problem (1.1) can be approximately solved by the *exponential-potential reduction method* which is described in detail in Grigoriadis and Khachiyan [2], or by that in Plotkin, Tardos and Shmoys [7] for linear $f^k(x^k)$; also see Leighton et al. [5]. As shown in [2], an ε-approximate solution of a feasible problem (1.1) can be computed in expected $O\left(K \log M(\varepsilon^{-2} + \log K)\right)$ *block optimizations*, each of the form:

$$\min \ \left\{ \ p^T f^k(x^k) \mid f^k(x^k) \leq \mathbf{1}, \ x^k \in B^k \ \right\},$$

where $p \in \mathbb{R}^{M+1}$ is a positive vector and $\mathbf{1}^T p = 1$. We suppress the block index k to write the kth block optimization problem as

$$a^* \equiv a(x^*) = \min\{ \ a(x) \mid b(x) \leq 1, \quad x \in X\}, \tag{1.5}$$

where we use the notation:

$$\begin{aligned} &x \doteq x^k, \quad a(x) \doteq p^T f^k(x^k), \quad b \doteq f_0^k(x^k), \\ &X \doteq \{x^k \in B^k \mid f_m^k(x^k) \leq 1, \quad m = 1, \dots, M \ \}. \end{aligned} \tag{1.6}$$

These block minimizations need only be performed approximately (see [2], p. 99) to a relative accuracy of $\varepsilon' = \Theta(\varepsilon)$. In other words, rather than solving the block problems exactly it suffices to compute an *ε'-approximate block minimizer* $\hat{x}$ that satisfies

$$a(\hat{x}) \leq (1+\varepsilon')a^*, \quad b(\hat{x}) \leq 1 + \varepsilon', \quad \hat{x} \in X. \tag{1.7}$$

Finally, (1.4), (1.6), and the fact that $\mathbf{1}^T p = 1$, provide the range

$$\varepsilon/K \leq a^* \leq 1. \tag{1.8}$$

Hence, in order to achieve a relative accuracy of $\varepsilon' = \Theta(\varepsilon)$ for problem (1.5) it suffices to solve it to an absolute accuracy of $\Theta(\varepsilon^2/K)$. Since $f^k(x^k) \leq \mathbf{1}$ for any feasible block solution x^k, each component of p can thus be rounded with an absolute accuracy of $\Theta(\varepsilon^2/KM)$, so that p becomes a nonnegative rational vector with a common denominator of binary length $O(\log(\varepsilon^{-1}KM))$.

2. Results. In the next section we shall consider an arbitrary problem of the form (1.5) with a given positive range for a^*:

$$a_L \leq a^* \leq a_H. \tag{2.1}$$

We argue there that for $\varepsilon' \geq \mathrm{poly}(\varepsilon)$, the task of computing an ε'-approximate block minimizer $\hat{x}$ satisfying (1.7) can be (deterministically) reduced to $O\left(\log(\varepsilon^{-1}a_H/a_L)\right)$ optimizations over X with an objective function that is a convex combination of $a(x)$ and $b(x)$. This leads to the following observations for the multicommodity flow problem:

First, X is a capacitated flow polytope and each block optimization (1.5) is a single-commodity minimum-cost network flow problem on G subject to a budget constraint $b^T x \leq 1$.

Second, in view of (1.8), the computation of an ε'-approximate block solution (1.7) reduces to $O\left(\log(\varepsilon^{-1}K)\right)$ single commodity minimum-cost network flow computations without a budget constraint. Since p is a rational vector with a denominator of binary length $O(\log(\varepsilon^{-1}KM))$, Goldberg and Tarjan's algorithm [4] for the latter problems gives an overall $\tilde{O}(NM)$ time bound for the approximate solution of the budget-constrained minimum-cost flow problem (1.5). This is an $O(\varepsilon^{-3})$ improvement over Karger and Plotkin's bound derived in [6].

Third, the computation of an ε-approximate solution (1.2) of the entire cost-constrained K-commodity flow problem (1.1) reduces to an expected

$$O\left(K(\varepsilon^{-2} + \log K)\log M \log(\varepsilon^{-1}K)\right)$$

single commodity minimum-cost flow computations on G. Recently, Radzik [8] showed that the randomized exponential-potential reduction method that selects blocks randomly from $\{1, \ldots, K\}$ can be replaced by a deterministic selection with no increase in the order of its complexity. Hence,

Theorem 1. *An ε-approximate solution of the minimum cost K-commodity flow problem for an N-node M-arc network can be computed in $\tilde{O}(\varepsilon^{-2}KNM)$ time.*

This bound improves the $\tilde{O}(\varepsilon^{-2}KM^2)$ time bound obtained by invoking the *logarithmic potential reduction method* developed in [3]. It also improves by a factor of ε^{-1} the more recent bound in [6], obtained for the special case of multicommodity flows with "regular costs" $c^1 = \ldots = c^K$.

3. Optimization with a budget constraint. Let $X \subset \mathbb{R}^n$ be a non-empty convex compact set, and let $a(x)$, $b(x)$ be a pair of nonnegative continuous convex functions on X. Consider the optimization problem (1.5), and assume that we are given positive lower and upper bounds (2.1) on a^*. This problem is primarily motivated by the construction described in the previous section, but it has other independent applications in budget-constrained network flows, time-constrained shortest paths, and cyclic scheduling (see, e.g., the book by Ahuja, Magnanti and Orlin [1]).

An ε-approximate solution of (1.5) is a point $\hat{x} \in X$ such that $a(\hat{x}) \leq (1+\varepsilon)a^*$ and $b(x) \leq 1+\varepsilon$. We wish to show that an ε-approximate solution of (1.5) can be computed in a small number of optimizations over X without the budget constraint.

Let $\Delta = \varepsilon^{-1} a_H$, and consider the concave function $g(t)$ on $[0, \Delta]$ defined by

$$g(t) = \min\{a(x) + t\,(b(x) - 1) \mid x \in X\} \doteq a\,(x(t)) + t\,(b\,(x(t)) - 1)\,, \quad x(t) \in X. \quad (3.1)$$

We shall write $a(t)$ for $a\,(x(t))$, $b(t)$ for $b\,(x(t))$, and refer to the problem of computing an $x(t)$ for some fixed $t \in [0, \Delta]$ by X-*optimization*. Note that for any $t \in [0, \Delta]$ we have the bounds:

$$-\Delta \leq g(t) \leq a^*. \quad (3.2)$$

The lower bound follows from the nonnegativity of $a(x)$ and $b(x)$, because $g(t) = a(t) + t\,(b(t) - 1) \geq -t \geq -\Delta$. The upper bound is obtained from (3.1) for an optimal solution x^* of (1.5): $g(t) \leq a(x^*) + t\,(b(x^*) - 1) \leq a(x^*) = a^*$.

The definition (3.1) also implies that

$$g(t') \leq g(t) + (b(t) - 1)\,(t' - t) \;\text{ for all } t, t' \in [0, \Delta], \quad (3.3)$$

that is, $b(t) - 1$ is a "supergradient" of $g(t)$ at t.

Bisection search provides a fast way to approximately maximize $g(t)$ on $[0, \Delta]$. We show below that it can also be used to compute an ε-approximate solution of (1.5) in

$$\ell = 5 + \log\left(\varepsilon^{-2} a_H / a_L\right)$$

X-optimizations.

First of all, if $b(0) - 1 \leq 0$, then $x(0)$ is an optimal solution of (1.5). It is also easy to see that $x(\Delta)$ is an ε-optimal solution of (1.5) whenever $b(\Delta) - 1 \geq 0$. This is because $g(\Delta) = a(\Delta) + \Delta\,(b(\Delta) - 1) \leq a^*$ by (3.2); under the assumption $b(\Delta) - 1 \equiv b(x(\Delta)) - 1 \geq 0$, this provides $a(\Delta) \leq a^*$ and $b(\Delta) - 1 \leq a^*/\Delta = \varepsilon a^*/a_H \leq \varepsilon$.

At most two X-optimizations are needed to check for the conditions $b(0) - 1 \leq 0$ and $b(\Delta) - 1 \geq 0$. Now assume that $b(0) - 1 \geq 0$ and $b(\Delta) \leq 0$, and perform $\ell - 2$ additional X-optimizations to find an interval $0 \leq r < s \leq \Delta$ of length $s - r = 2^{-\ell-2}\Delta \leq \varepsilon a^*/8$ such that $b(r) - 1 \geq 0$ and $b(s) - 1 \leq 0$. By the continuity of $a(x)$, $b(x)$, there is a point

$$\hat{x} = \lambda x(r) + (1 - \lambda)x(s), \quad \lambda \in [0, 1], \quad (3.4)$$

at which $b(\hat{x}) = 1$. Such an $\hat{x}$ can be computed analytically for linear and quadratic $a(x)$, $b(x)$, or can be approximated to an arbitrary accuracy for general $a(x)$, $b(x)$. No additional X-optimizations are needed in either case.

In order to show that $\hat{x}$ is an ε-approximate solution of (1.5) we need only argue that $a(\hat{x}) \leq (1+\varepsilon)a^*$, since $\hat{x} \in [x(r), x(s)] \subseteq X$. Clearly, $\max\{r, \Delta - s\} \geq \Delta(1-2^{-l+2})/2 > \Delta/4$. Suppose that $r \geq \Delta/4$. (The case for $\Delta - s \geq \Delta/4$ is completely symmetric.) Apply the inequality (3.3) for $t = r$ and $t' = 0$ along with (3.2) to obtain

$$b(r) - 1 \leq \frac{g(r) - g(0)}{r} \leq \frac{a^* + \Delta}{\Delta/4} \leq \frac{1 + \varepsilon^{-1}}{\varepsilon^{-1}/4} \leq 8.$$

This gives

$$a(r) + s\,(b(r) - 1) = g(r) + (b(r) - 1)\,(s - r) \leq g(r) + 8(s - r) \leq (1 + \varepsilon)a^*.$$

Multiply this inequality by λ, where λ is defined by (3.4), and add it to $1 - \lambda$ times the inequality $a(s) + s\,(b(s) - 1) = g(s) \leq a^*$ to obtain:

$$\lambda a(r) + (1 - \lambda)a(s) + s\{\lambda b(r) + (1 - \lambda)b(s) - 1\} \leq (1 + \varepsilon)a^*.$$

By the convexity of $b(x)$ we have $1 = b(\hat{x}) \leq \lambda b(r) + (1 - \lambda)b(s)$, confirming the nonnegativity of the same expression in the curly brackets. Finally, $a(\hat{x}) \leq \lambda a(r) + (1 - \lambda)a(s) \leq (1 + \varepsilon)a^*$, by the convexity of $a(x)$.

Acknowledgments. We thank Neal Young for pointing out reference [6] to us and Jorge Villavicencio for drawing our attention to a flaw in an earlier draft of this note.

References

1 R.K. Ahuja, T.L. Magnanti, and J.B. Orlin, *Network Flows*, Prentice-Hall, NJ (1993).

2 M.D. Grigoriadis and L.G. Khachiyan, "Fast approximation schemes for convex programs with many blocks and coupling constraints", *SIAM J. Optimization* 4 (1994), pp. 86–107.

3 ——-and ——-, "Coordination complexity of parallel price-directive decomposition", Tech. Rep. 94-19, DIMACS, Rutgers University, New Brunswick, NJ, April 1994, rev. November 1994. To appear in *Mathematics of Operations Research.*

4 A. Goldberg and R.E. Tarjan, "Solving minimum cost flow problems by successive approximation", *Mathematics of Operations Research* 15 (1990), pp. 430–466.

5 T. Leighton, F. Makedon, S. Plotkin, C. Stein, E. Tardos, and S. Tragoudas, "Fast approximation algorithms for multicommodity flow problems", *Proceedings of the 23rd Annual ACM Symposium on Theory of Computing*, pp.101–111, May 1991.

6 D. Karger and S.A. Plotkin, "Adding multiple cost constraints to combinatorial optimization problems with applications to multicommodity flows", *Proc. 27th Annual ACM Symp. on Theory of Computing*, May 1995.

7 S.A. Plotkin, D.B. Shmoys, and E. Tardos, "Fast approximation algorithms for fractional packing and covering problems", in *Proc. 32nd Annual Symp. on Foundations of Computer Science*, pp. 495–504, October 1991.

8 T. Radzik, "Fast deterministic approximation for the multicommodity flow problem", *Proc. 6th ACM-SIAM Symp. on Discrete Algorithms*, pp. 486-492, January 1995.