CORE DISCUSSION PAPER
9947

A FASTER CAPACITY SCALING ALGORITHM FOR MINIMUM COST SUBMODULAR FLOW

Lisa Fleischer[1], Satoru Iwata[2], and S. Thomas McCormick[3]

July 1999

Abstract

We describe an $O(n^4 h \min\{\log U, n^2 \log n\})$ capacity scaling algorithm for the minimum cost submodular flow problem. Our algorithm modifies and extends the Edmonds–Karp capacity scaling algorithm for minimum cost flow to solve the minimum cost submodular flow problem. The modification entails scaling a relaxation parameter δ. Capacities are relaxed by attaching a complete directed graph with uniform arc capacity δ in each scaling phase. We then modify a feasible submodular flow by relaxing the submodular constraints, so that complementary slackness is satisfied. This creates discrepancies between the boundary of the flow and the base polytope of a relaxed submodular function. To reduce these discrepancies, we use a variant of the successive shortest path algorithm that augments flow along min cost paths of residual capacity at least δ. The shortest augmenting path subroutine we use is a variant of Dijkstra's algorithm modified to handle exchange capacity arcs efficiently. The result is a weakly polynomial time algorithm whose running time is better than any existing submodular flow algorithm when U is small and C is big. We also show how to use max mean cuts to make the algorithm strongly polynomial. The resulting algorithm is the first capacity scaling algorithm to match the current best strongly polynomial bound for submodular flow.

Keywords: submodular flow, shortest paths, strongly polynomial time algorithm

[1]Department of Industrial Engineering and Operations Research, Columbia University, New York, NY 10027, USA. Email: lisa@ieor.columbia.edu. This research was done in part while visiting RIMS, Kyoto University, with support by Grant-in-Aid for Scientific Research No. 10750053 from Ministry of Education, Science, Sports, and Culture of Japan, and while on leave at CORE, Université catholique de Louvain, Belgium.

[2]Division of Systems Science, Graduate School of Engineering Sciences, Osaka University, Toyonaka, Osaka 560-8531, Japan. Email: iwata@sys.es.osaka-u.ac.jp. Partially supported by Grant-in-Aid for Scientific Research No. 10750053 from Ministry of Education, Science, Sports, and Culture of Japan.

[3]Faculty of Commerce and Business Administration, University of British Columbia, Vancouver, BC V6T 1Z2 Canada. Email: stmv@adk.commerce.ubc.ca. Research supported by an NSERC Operating Grant; part of this research was done during a visit to SORIE at Cornell University.
This text presents research results of the Belgian Program on Interuniversity Poles of Attraction initiated by the Belgian State, Prime Minister's Office, Science Policy Programming. The scientific responsibility is assumed by the authors.

1 Introduction

The submodular flow problem, introduced by Edmonds and Giles [3], is one of the most important frameworks of efficiently solvable combinatorial optimization problems. It includes minimum cost flow, graph orientation, polymatroid intersection, and directed cut covering problems as special cases. To talk about running time we use n for the number of nodes, h for the time to compute a single exchange capacity (defined below), C for the maximum absolute value of the arc costs, and U for the maximum absolute value of the capacities. When we use C or U we assume that those data are integral.

A number of minimum cost flow algorithms have been extended to solve the submodular flow problem. The first such polynomial algorithm is due to Cunningham and Frank [1]. It generalizes a cost-scaling primal-dual method. A variant of this algorithm, combined with the push/relabel algorithm of Fujishige and Zhang [14], runs in $\mathrm{O}(n^4 h \log C)$ time [19]. This is the previous best weakly polynomial bound for submodular flow. A different cost-scaling approach is based on cycle canceling [17], and achieves a running time of $\mathrm{O}(n^4 h \min\{\log(nC), n^2 \log n\})$.

The dual approach to scaling costs is to instead scale the capacities. The prototype capacity scaling min cost flow algorithm is the Edmonds–Karp [4] algorithm that scales and then rounds the capacities. Generalizing this to submodular flow is difficult because rounding a scaled submodular function does not preserve submodularity. To obviate this difficulty, the algorithm in [15] adds a small multiple of the strictly submodular function b (which is the symmetric cut function of a complete directed graph with unit arc capacity) before rounding. However, the scaling scheme discussed in that paper calls maximum submodular flow algorithms repeatedly, and hence requires $\mathrm{O}(n^7 h \log U)$ time.

An easier way to do capacity scaling for submodular flow was introduced in [18]. There the submodular bounds are again relaxed by a multiple of b, but now only the relaxation parameter is scaled instead of the data. The result was the first strongly polynomial capacity scaling algorithm for submodular flow, but it again calls maximum submodular flow algorithms repeatedly and so runs in $\mathrm{O}(n^6 h \min\{\log(nU), n^2 \log n\})$ time.

In this paper, we modify and extend the Edmonds–Karp capacity scaling algorithm to an algorithm for submodular flow. A key innovation of our algorithm is a modification of Dijkstra's algorithm to efficiently handle exchange capacity arcs inherent in submodular flow. As in [15], capacities are relaxed by attaching a complete directed graph with uniform arc capacity δ in each scaling phase. As in [18], we relax the problem by a parameter δ, which is scaled. We then modify our current feasible submodular flow so that complementary slackness is satisfied. This creates discrepancies between the boundary of the flow and a base in the base polytope. To reduce these discrepancies, we use a variant of the successive shortest path algorithm that augments flow along min cost paths of residual capacity at least δ. Our shortest path algorithm is further specialized to dynamically modify flows so that we can handle exchange capacity arcs efficiently. This yields a weakly polynomial algorithm with a new run time, $\mathrm{O}(n^4 h \log U)$, that is faster than all other known submodular flow algorithms when U is smaller than C and $n^{\mathrm{O}(n^2)}$.

As in [18] we also show how to use (approximations of) max mean cuts to get a strongly polynomial version of the algorithm. The resulting algorithm runs in $\mathrm{O}(n^6 h \log n)$ time, matching the best current strongly polynomial bound.

The first strongly polynomial algorithm for the submodular flow problem is due to Frank and Tardos [7]. This is an application of simultaneous Diophantine approximation and substantially generalizes the first strongly polynomial minimum cost flow algorithm of Tardos [25]. A more direct generalization of the Tardos algorithm to the submodular flow problem is described by Fujishige, Röck, and Zimmermann [13] with the aid of the tree-projection method of

Fujishige [11]. The latter algorithm combined with the improved primal-dual algorithm in [19] runs in $O(n^6 h \log n)$ time, the current best strongly polynomial bound, which is also achieved by the lexicographic cycle canceling algorithm in [17].

Our algorithm depends crucially on the sophisticated new technique of modifying Dijkstra's algorithm to deal with the exchange arcs inherent in submodular flow, while ignoring small residual capacity arcs. This novel technique is further extended in [16] to obtain the first combinatorial, polynomial-time algorithm for the fundamental problem of minimizing a submodular function.

In the same way that cost scaling is associated with cycle canceling, capacity scaling is associated with cut canceling (see [24] for a survey of this for min cost flow). Although our algorithm was developed as an extension of the Edmonds–Karp algorithm, it is also possible to see it as a way of speeding up the cut canceling algorithm for submodular flow described in [18]. That is, it is well-known that updating dual variables via shortest path distances from Dijkstra's algorithm can be seen as simultaneously canceling several cuts at once. From this perspective, our new algorithm is a much faster version of the algorithm of [18], thus answering the open question of finding such an algorithm. Thus the new algorithm merges the two seemingly different strands of research of [15] and [18].

2 Preliminaries on Submodular Functions

In this section, we review some well-known facts about submodular functions, base polyhedra, and exchange capacities. We re-prove many of the results for completeness, and cite the relevant results in the monograph by Fujishige [12] for readers interested in more details.

2.1 Base Polyhedra

Let V be a finite set, and $f : 2^V \to \mathbf{R}$ be a set function on V. Then f is *submodular* if

$$f(X) + f(Y) \geq f(X \cup Y) + f(X \cap Y) \tag{1}$$

holds for all $X, Y \subseteq V$. The *base polyhedron* of f is

$$\mathrm{B}(f) = \{x \mid x \in \mathbf{R}^V,\ x(V) = f(V),\ \forall X \subseteq V : x(X) \leq f(X)\},$$

and $x \in \mathrm{B}(f)$ is called a *base*. For $p \in \mathbf{R}^V$, consider the linear program to maximize $\sum_{v \in V} p(v)x(v)$ on the base polyhedron $\mathrm{B}(f)$. An optimal solution is called a *p-maximum base*. Let $p_1 > \cdots > p_k$ be the distinct values of $p(v)$, and put $L_i = \{v \mid p(v) \geq p_i\}$, the ith *level set* of p. Let $f_p : 2^V \to \mathbf{R}$ be defined by

$$f_p(X) = \sum_{i=1}^{k} \{f((X \cap L_i) \cup L_{i-1}) - f(L_{i-1})\},$$

where $L_0 = \emptyset$. The following theorems give properties of $\mathrm{B}(f_p)$ that we will need.

Theorem 2.1 ([1, Theorem 6]) *This f_p is a submodular function, and $x \in \mathrm{B}(f_p)$ iff $x \in \mathrm{B}(f)$ and $x(L_i) = f(L_i)$ for every i.* ■

Theorem 2.2 ([12, Theorem 3.15]) *A base $x \in \mathrm{B}(f)$ is p-maximum if and only if $x \in \mathrm{B}(f_p)$.* ■

2.2 Exchange Capacity

For each $v \in V$, we define $\vec{v} \in \mathbf{R}^V$ such that $\vec{v}(v) = 1$ and $\vec{v}(w) = 0$ for $w \neq v$. For a base $x \in \mathrm{B}(f)$ and a pair of distinct elements $v, w \in V$, we define the *exchange capacity* $\sigma(x, v, w)$ by

$$\sigma(x, v, w) = \max\{\alpha \mid \alpha \in \mathbf{R},\ x + \alpha(\vec{v} - \vec{w}) \in \mathrm{B}(f)\}.$$

Equivalently, the exchange capacity can be written as

$$\sigma(x, v, w) = \min\{f(X) - x(X) \mid v \in X,\ w \notin X\}.$$

Given $x \in \mathrm{B}(f)$, the *exchangeability graph* w.r.t. x is the directed graph with node set V and arc set $E_x = \{(w, v) \mid \sigma(x, v, w) > 0\}$. An arc $(w, v) \in E_x$ is called an *exchange arc*. For $0 < \alpha \leq \sigma(x, v, w)$ the operation $x = x + \alpha(\vec{v} - \vec{w})$ is called *augmenting* (w, v) by α. If $\alpha = \sigma(x, v, w)$, then we say that this operation *saturates* (w, v).

Lemma 2.3 ([12, Theorem 3.16]) *A base $x \in \mathrm{B}(f)$ is p-maximum if and only if $p(w) \geq p(v)$ holds for any $(w, v) \in E_x$.*

Proof. The "only if" part is immediate from the definition of the exchangeability graph. To prove the "if" part, Theorem 2.1 says that it suffices to show that $x(L_i) = f(L_i)$ holds. For any pair of $w \in V - L_i$ and $v \in L_i$, it follows from $(w, v) \notin E_x$ that there exists a subset $X_{wv} \subseteq V$ such that $v \in X_{wv}$, $w \notin X_{wv}$, and $x(X_{wv}) = f(X_{wv})$. Then by the submodularity of f and

$$L_i = \bigcup_{v \in L_i} \bigcap_{w \in V - L_i} X_{wv},$$

we obtain $x(L_i) = f(L_i)$. ■

Lemma 2.4 *The exchangeability graph is transitive. Namely, $(w, v) \in E_x$ and $(v, z) \in E_x$ imply $(w, z) \in E_x$.*

Proof. Suppose to the contrary $(w, z) \notin E_x$. Then there must be an $X \subseteq V$ that satisfies $z \in X$, $w \notin X$, and $x(X) = f(X)$. If $v \in X$, this contradicts $(w, v) \in E_x$. On the other hand, if $v \notin X$, it contradicts $(v, z) \in E_x$. ■

Lemma 2.5 ([12, Figure 4.4]) *Suppose $(t, z) \notin E_x$ and $(t, z) \in E_y$ hold for a base y obtained from $x \in \mathrm{B}(f)$ by $y = x + \alpha(\vec{v} - \vec{w}) \in \mathrm{B}(f)$ with $0 < \alpha \leq \sigma(x, v, w)$. Then $t \neq v$ implies $(t, v) \in E_x$ and $z \neq w$ implies $(w, z) \in E_x$.*

Proof. Since $(t, z) \notin E_x$, there must be an $X \subseteq V$ that satisfies $z \in X$, $t \notin X$, and $x(X) = f(X)$. If $t \neq v$ and $(t, v) \notin E_x$, then there exists an $Y \subseteq V$ such that $v \in Y$, $t \notin Y$, and $x(Y) = f(Y)$. Submodularity of f and $v \in X \cup Y$ imply that $y(X \cup Y) = x(X \cup Y) = f(X \cup Y)$, which together with $z \in X \cup Y$ and $t \notin X \cup Y$ contradicts $(t, z) \in E_y$. The $z \neq w$ case is similar. ■

Lemma 2.6 ([12, Theorem 4.4]) *For distinct nodes $t, v, w, z \in V$ and a base $x \in \mathrm{B}(f)$, suppose $(w, v) \in E_x$, $(t, z) \in E_x$, and $(w, z) \notin E_x$. Then $\sigma(y, z, t) = \sigma(x, z, t)$ holds for a base y obtained by $y = x + \alpha(\vec{v} - \vec{w}) \in \mathrm{B}(f)$ with $0 < \alpha \leq \min\{\sigma(x, z, t), \sigma(x, v, w)\}$.*

Proof. Since $f(X) - y(X) < f(X) - x(X)$ only if $v \in X$ and $w \notin X$, and $\sigma(\cdot, z, t)$ only depends on sets X with $z \in X$, $t \notin X$, it suffices to show that $v, z \in X$ and $t, w \notin X$ imply $f(X) - x(X) \geq \sigma(x, z, t) + \alpha$. Since $(w, z) \notin E_x$, there exists a subset $Y \subseteq V$ such that $z \in Y$, $w \notin Y$, and $x(Y) = f(Y)$. Then $z \in X \cap Y$ and $t \notin X \cap Y$ imply $f(X \cap Y) - x(X \cap Y) \geq \sigma(x, z, t)$. It follows from $v \in X \cup Y$ and $w \notin X \cup Y$ that $f(X \cup Y) - x(X \cup Y) \geq \sigma(x, v, w)$ holds. Therefore we obtain $f(X) - x(X) \geq \sigma(x, z, t) + \sigma(x, v, w) \geq \sigma(x, z, t) + \alpha$ by the submodularity of f. ■

Lemma 2.7 ([12, Theorem 4.5]) *Let x be a base of f. Suppose $(w_1, v_1), (w_2, v_2), \cdots, (w_k, v_k) \in E_x$ have distinct end-nodes and that $i < j$ implies $(w_i, v_j) \notin E_x$. If $\alpha \leq \sigma(x, v_i, w_i)$ holds for every $i = 1, 2, \ldots, k$, then $y = x + \alpha \sum_{i=1}^{k} (\vec{v}_i - \vec{w}_i)$ is also a base of f.*

Proof. By Lemma 2.6, $x' = x + \alpha(\vec{v}_1 - \vec{w}_1) \in \mathrm{B}(f)$ satisfies $\sigma(x', v_i, w_i) = \sigma(x, v_i, w_i)$ for $i = 2, \ldots, k$. It follows from Lemma 2.5 that $i < j$ implies $(w_i, v_j) \notin E_{x'}$. Therefore the result follows by induction on k. ■

3 The Submodular Flow Problem

Let $G = (V, A)$ be a directed graph with *lower* and *upper bounds* (*capacities*) $l : A \to \mathbf{R} \cup \{-\infty\}$ and $u : A \to \mathbf{R} \cup \{+\infty\}$ with $l \leq u$, and costs $c \in \mathbf{R}^A$. A *flow* is a function φ obeying $l(a) \leq \varphi(a) \leq u(a)$, $\forall a \in A$. We assume there are no parallel arcs, nor loops, nor isolated nodes in G, so that $\Theta(n) \leq m \leq \Theta(n^2)$. Given arc $(w, v) = a \in A$, define $\partial^+ a = w$ and $\partial^- a = v$. Given a set $X \subset V$, define $\Delta^+ X := \{a \in A \mid \partial^+ a \in X, \partial^- a \notin X\}$ (the set of arcs leaving X), and $\Delta^- X := \{a \in A \mid \partial^+ a \notin X, \partial^- a \in X\}$ (the set of arcs entering X). Given a flow φ and $X \subseteq V$, the *boundary* of φ on X is $\partial\varphi(X) = \varphi(\Delta^+ X) - \varphi(\Delta^- X)$, the net flow leaving X. Let f be a submodular function on V with $f(\emptyset) = f(V) = 0$. Then the *submodular flow problem* is:

$$\begin{array}{ll} \text{Minimize} & c^\top \varphi \\ \text{subject to} & l \leq \varphi \leq u \\ & \partial\varphi \in \mathrm{B}(f). \end{array}$$

A feasible solution is called a *submodular flow*.

Theorem 3.1 (Frank [6]) *There exists a submodular flow if and only if*

$$l(\Delta^+ X) - u(\Delta^- X) \leq f(X) \tag{2}$$

holds for every $X \subseteq V$. If in addition l, u, f are integer valued, then there exists an integral submodular flow. ■

Given a *price function* (or *node potentials*) $p \in \mathbf{R}^V$, we define the *reduced cost* w.r.t. p as $c_p(a) = c(a) + p(\partial^+ a) - p(\partial^- a)$ for each $a \in A$. The following optimality conditions for the submodular flow problem are due to Fujishige [12, Theorem 5.3]. Edmonds and Giles [3] show that the submodular flow problem is totally dual integral, which implies the integrality claim.

Theorem 3.2 *A submodular flow φ is optimal if and only if there exists $p \in \mathbf{R}^V$ such that:*
⟨a⟩ *For any $a \in A$, $c_p(a) > 0$ implies $\varphi(a) = l(a)$, and $c_p(a) < 0$ implies $\varphi(a) = u(a)$, and*
⟨b⟩ *The boundary $\partial\varphi$ is a p-maximum base in $\mathrm{B}(f)$.*
Moreover, if c is integral, then we may restrict the above p to be integral. ■

Edmonds and Giles [3] originally formulated the submodular flow problem in terms of crossing submodular functions, which are not necessarily defined on all subsets of V. On the other hand, Fujishige [10] shows that a base polyhedron of a crossing submodular function is identical to a base polyhedron of a submodular function defined on all subsets (fully submodular). Since most of the algorithms presented in this paper rely only on geometric properties of the base polyhedron, we deal with fully submodular functions for the sake of simplicity.

4 The Successive Shortest Path Algorithm

This section describes the successive shortest path algorithm for the submodular flow problem, which is the basis for our algorithms. As with successive shortest path algorithm for minimum cost flows, this algorithm is not polynomial, but only pseudopolynomial since the number of iterations is linear in U. The algorithm was originally presented by Fujishige [9] for solving the minimum cost independent flow problem, which turns out to be equivalent to the submodular flow problem [12, §5.2].

We start with a flow φ and a price vector p that satisfy condition $\langle$a$\rangle$ of Theorem 3.2. First we check whether there exists a primal feasible solution, which can be done in $\mathrm{O}(n^3h)$ time [14, 19]. Next we check whether there exists a dual feasible solution p. To do this we construct a directed graph $G^\circ = (V, A^\circ)$ with the arc set $A^\circ = F^\circ \cup B^\circ$ defined by $F^\circ = \{a \mid u(a) = +\infty\}$ and $B^\circ = \{\overline{a} \mid l(a) = -\infty\}$, where $\overline{a}$ is the reversal of a. Each arc in A° has length $c(a)$ if $a \in F^\circ$ and $-c(\overline{a})$ if $a \in B^\circ$. We use the $\mathrm{O}(mn)$ Bellman–Ford–Moore shortest path algorithm to try to compute shortest path distances p in G°. If the algorithm finds a negative cycle in G°, then the instance is dual infeasible [12, Theorem 5.5]. Otherwise, the shortest path distances p satisfy $c_p(a) > 0$ implies $l(a) > -\infty$ and $c_p(a) < 0$ implies $u(a) < +\infty$. Then we can construct a flow φ that satisfies condition $\langle$a$\rangle$ of Theorem 3.2 in $\mathrm{O}(m)$ time. Since we now know that the instance is both primal and dual feasible, we know that an optimal solution exists.

We also start with a base $x \in \mathrm{B}(f_p)$, which can be obtained using the greedy algorithm [2]. If $x = \partial\varphi$, then x satisfies condition $\langle$b$\rangle$ of Theorem 3.2, proving that φ is optimal (and that p is dual optimal). The algorithm will iteratively modify φ, x, and p so that $\langle$a$\rangle$ is preserved for φ and p, $\langle$b$\rangle$ is preserved for x and p, and so that $\partial\varphi$ and x eventually converge.

We construct an auxiliary network $\widehat{G}(\varphi, x) = (V, A_\varphi \cup E_x)$ with respect to flow φ and base $x \in \mathrm{B}(f_p)$. The arc set $A_\varphi := F_\varphi \cup B_\varphi$ is defined by $F_\varphi := \{a \mid a \in A,\ \varphi(a) < u(a)\}$ and $B_\varphi := \{\overline{a} \mid a \in A,\ \varphi(a) > l(a)\}$ (in other words, A_φ is just the ordinary residual graph of flow φ). As defined in Section 2.2, E_x is the arc set of the exchangeability graph w.r.t. $\mathrm{B}(f)$ (we emphasize that this is $\mathrm{B}(f)$ rather than $\mathrm{B}(f_p)$ to allow for changes to p). For each arc $a \in A_\varphi \cup E_x$, we define the residual capacity $r(a)$ and the cost $c(a)$ by

$$r(a) := \begin{cases} u(a) - \varphi(a) & (a \in F_\varphi) \\ \varphi(a) - l(a) & (a \in B_\varphi) \\ \sigma(x, v, w) & (a = (w, v) \in E_x), \end{cases} \qquad c(a) := \begin{cases} c(a) & (a \in F_\varphi) \\ -c(\overline{a}) & (a \in B_\varphi) \\ 0 & (a \in E_x). \end{cases}$$

The algorithm computes augmentations of φ and x w.r.t. r on $A_\varphi \cup E_x$. This will ensure that φ remains a flow. Together with the corresponding update to p, it will also ensure that x remains in $\mathrm{B}(f_p)$.

We regard the reduced cost $c_p(a) = c(a) + p(\partial^+ a) - p(\partial^- a)$ as the length of $a \in A_\varphi \cup E_x$. Since φ and p satisfy $\langle$a$\rangle$, which is equivalent to having $c_p(a) \geq 0$ for all $a \in A_\varphi$, we can compute shortest path distances using Dijkstra's algorithm. Define $S^+ = \{v \mid x(v) > \partial\varphi(v)\}$ and $S^- = \{v \mid x(v) < \partial\varphi(v)\}$ as the sets of nodes where $\partial\varphi$ differs from x. The idea of the algorithm

SSP (x, φ, p)

Repeat the following Steps (1-1)–(1-4), until $S^+ = \emptyset$.

(1-1) Compute the shortest distance $d(v)$ from S^+ to each $v \in V$ in $\widehat{G}(\varphi, x)$ with respect to the arc length c_p. Among the shortest paths from S^+ to S^-, let P be one with a minimum number of arcs. Let s be the initial node of P and t be the final node of P.

(1-2) For each $v \in V$, put $p(v) := p(v) + \min\{d(v), c_p(P)\}$.

(1-3) Put $\alpha := \min\{x(s) - \partial\varphi(s),\ \partial\varphi(t) - x(t),\ \min\{r(a) \mid a \in P\}\}$.

(1-4) For each arc $a \in P$,

- if $a \in F_\varphi$, then $\varphi(a) := \varphi(a) + \alpha$;
- if $a \in B_\varphi$, then $\varphi(\overline{a}) := \varphi(\overline{a}) - \alpha$;
- if $a \in E_x$, then $x(\partial^+ a) := x(\partial^+ a) - \alpha$, $x(\partial^- a) := x(\partial^- a) + \alpha$.

Figure 1: Successive Shortest Path algorithm.

is to compute shortest path distances $d(v)$ (w.r.t. c_p) from any node in S^+ to each $v \in V$. Among the nodes in S^- with minimum $d(v)$ we choose a path P with a minimum number of arcs. Since we will augment φ only on P and not on any longer paths from S^+ to S^-, we update $p(v)$ by $\min(d(v), c_p(P))$. Then augmenting φ on P and updating p preserves $\langle$a$\rangle$ and $\langle$b$\rangle$, will decrease $|\partial\varphi(v) - x(v)|$ by the augmentation amount for the initial and final nodes of P, and will not change $|\partial\varphi(v) - x(v)|$ for any other node. The detailed algorithm is in Figure 1.

If there is no path from S^+ to S^-, then the problem is infeasible. In fact, the set X of nodes not reachable from S^+ satisfies $l(\Delta^+ X) - u(\Delta^- X) = \varphi(\Delta^+ X) - \varphi(\Delta^- X) = \partial\varphi(X) > x(X) = f(X)$, which proves infeasibility by Theorem 3.1.

Since P has a minimum number of arcs among the shortest paths, it follows from Lemma 2.4 that there is no consecutive pair of exchangeability arcs in P. The reduced cost c_p is nonnegative after Step (1-2). In particular, each arc a in P satisfies $c_p(a) = 0$. The following Lemma 4.1 makes it possible to apply Lemma 2.7 to show that x is a base after Step (1-3). Furthermore, Lemma 2.5 can be shown to ensure that c_p is nonnegative.

Lemma 4.1 *There exists a numbering $(w_1, v_1), (w_2, v_2), \cdots, (w_k, v_k)$ of the arcs in $P \cap E_x$ such that $i < j$ implies $(w_i, v_j) \notin E_x$.*

Proof. If no such numbering exists, there is a subset $\{(w_j, v_j) \mid j = 1, \ldots, q\}$ of $P \cap E_x$ such that $e_j = (w_{j-1}, v_j) \in E_x$ holds for each j, where $w_0 \equiv w_q$. After Step (1-2), we have $p(w_j) = p(v_j)$, which implies $\sum_{j=1}^{q} c_p(e_j) = 0$. Therefore $c_p(e_j) = 0$ holds, which contradicts the minimality of P. ■

If l, u and f are integer valued, then $\sum_{v \in V} |x(v) - \partial\varphi(v)|$ is always a nonnegative integer and decreases by at least two with each augmentation. Hence the algorithm terminates in a finite number of iterations. Since the number of arcs in $\widehat{G}(\varphi, x)$ is $|A_\varphi| + |E_x| = \mathrm{O}(m) + \mathrm{O}(n^2) = \mathrm{O}(n^2)$, the bottleneck Dijkstra computation takes $\mathrm{O}(n^2 h)$ time, so each augmentation requires $\mathrm{O}(n^2 h)$ time.

5 A Capacity Scaling Algorithm

This section develops our new capacity scaling algorithm for solving submodular flow problems. We assume throughout this section that l, u and f are integer valued.

5.1 Algorithm Overview

To obtain a polynomial time algorithm, we embed a variant of the successive shortest path algorithm within a capacity scaling framework. For the minimum cost flow problem, Edmonds and Karp [4] propose a capacity scaling algorithm that maintains reduced cost optimality conditions while relaxing node conservation constraints. We modify the Edmonds–Karp algorithm by relaxing the submodular constraints and by attaching a complete directed graph on the node set (thus also relaxing capacity constraints).

Without loss of generality, we assume that the graph restricted to infinite capacity arcs is strongly connected. With this assumption, between any pair of nodes there is always an infinite capacity path in the residual graph of any flow with non-negative reduced cost. This assumption may be enforced by introducing high cost, infinite capacity arcs between all node pairs. If these arcs are used in the final solution, then the problem is infeasible.

The standard minimum cost flow version of SSP is not guaranteed to run in polynomial time, and the submodular flow version is not either, because the number of iterations might depend linearly on U. The Edmonds–Karp algorithm effectively deals with this by scaling capacities in the residual graph by a parameter $\delta > 0$, so that each augmentation is by at least δ units of flow. For our modification of this for submodular flow, we relax capacities by δ. We do this by adding the arc set of a complete directed graph (V, D) with arc set $D = \{(w, v) \mid w \neq v \in V\}$ with capacity δ to our initial graph. For all $a \in D$ we extend c, l, and u to D by setting $c(a) = 0$, $l(a) = 0$, and $u(a) = \delta$. Define the submodular function $b : 2^V \to \mathbf{R}$ by $b(X) = |X| \cdot |V - X|$. Equivalently, $\delta b(X)$ is the capacity of the cut X in (V, D). This relaxation can be thought of as either relaxing the condition that $\partial\varphi \in \mathrm{B}(f)$ to $\partial\varphi \in \mathrm{B}(f + \delta b)$, or as relaxing the capacities l and u by δ, or both. In the following we will treat the arcs of D as having their own separate flow.

To avoid confusion we will always use φ for a flow on the original set of *flow arcs* A, we will always use ψ for a flow on the *relaxation arcs* D, and we will always use ξ for a flow on the combined arc set $I := A \cup D$. Note that there is also a third arc set, the *exchange* arcs. There is no actual flow on an exchange arc, but augmenting along an exchange arc (w, v) by δ corresponds to changing base x by $x(w) = x(w) - \delta$ and $x(v) = x(v) + \delta$.

At any given point in the algorithm we will have a flow φ on A, and flow ψ on D, a price vector p, and a base $x \in \mathrm{B}(f_p + \delta b)$. We maintain x as the sum of $y \in \mathrm{B}(f_p)$ and $-\partial\psi \in \mathrm{B}(\delta b)$. We compute exchange capacities w.r.t. y and w.r.t. the original function f, not f_p (as in SSP), so we denote the set of exchange arcs by E_y, as defined in Section 2.2. We measure progress in the algorithm via the *discrepancy* between x and $\partial\varphi$, which is defined by the *discrepancy function* $\Phi = \sum_v |x(v) - \partial\varphi(v)|$. In this sense, we relax the submodular constraints on φ twice: once in relaxing f to $f + \delta b$, and once in letting $\partial\varphi$ and x differ. Recall that, because the algorithm maintains $\langle$a$\rangle$ of Theorem 3.2 for φ and p, and $\langle$b$\rangle$ of Theorem 3.2 for x and p, if we can drive Φ to zero we are optimal.

The algorithm MCSF is described in Figure 2. It embeds a modified version of SSP in a scaling framework that depends on parameter δ. We call the iterations with a fixed value of δ a *δ-scaling phase*, or just a phase for short. A node v has *excess discrepancy*, or simply *excess*, if $x(v) > \partial\varphi(v)$ $(v \in S^+)$, and *deficit discrepancy*, or *deficit*, if $x(v) < \partial\varphi(v)$ $(v \in S^-)$. Since φ is a

flow and x is a base, the sum of excesses equals the sum of deficits. Inside a phase, we find shortest paths among paths from $S^+(\delta) := \{v \mid x(v) - \partial\varphi(v) \geq \delta\}$ to $S^-(\delta) := \{v \mid \partial\varphi(v) - x(v) \geq \delta\}$ with residual capacity at least δ (by our assumption on the subgraph on infinite capacity arcs, there is always an augmenting path between any two such nodes). The lower bound of δ on the residual capacity of an augmenting path will ensure a small number of iterations within a phase, since we are making "large" changes at each augmentation. At the end of a phase we set $\delta := \delta/2$ and continue until $\delta = 1$, at which point we can finish using SSP.

We start with the same initial flow φ, base x, and price function p as in SSP. Define U as

$$U = \max\{\max\{|u(a)| \mid u(a) < +\infty\}, \max\{|l(a)| \mid l(a) > -\infty\}, \max\{f(\{v\}) \mid v \in V\}\}.$$

Since $\varphi(a) \leq U$ for all $a \in A$, $|\partial\varphi(v)| \leq nU$ for all $v \in V$. For any p it is easily checked that $f_p(\{v\}) \leq f(\{v\})$ by using the submodularity of f and $f(\emptyset) = 0$. Since $x(V) = f(V) = 0$, we have that $|x(v)| \leq (n-1)U$. Thus the initial discrepancy between x and φ is at most $2n^2U$.

For any distinct $w, v \in V$, we may assume that at least one of $\psi(w,v)$ and $\psi(v,w)$ is zero, so that at least one of $r(v,w)$ and $r(w,v)$ equals δ. Given such a flow ψ, let D_ψ denote the set of relaxation arcs with residual capacity equal to δ, i.e., $D_\psi = \{(w,v) \mid \psi(w,v) = 0\}$. Define $A_\varphi(\delta) := \{a \mid a \in A_\varphi,\ r(a) \geq \delta\}$. Subroutine SubmodularDijkstra of Figure 3 adapts Dijkstra's algorithm to find a shortest path from any node in $S^+(\delta)$ to any node in $S^-(\delta)$ in the network with arc set $A_\varphi(\delta) \cup D_\psi$, while maintaining the δ reduced cost optimality condition that

$$c_p(a) \geq 0 \text{ holds for every } a \in A_\varphi(\delta) \cup D_\psi \cup E_y. \tag{3}$$

We call such a path a *δ-augmenting* path. It might seem that the algorithm should also include the subset of arcs of E_y with exchange capacity at least δ in its list of arcs eligible for a δ-augmenting path. Instead, SubmodularDijkstra does this implicitly by using the relaxation arcs to dynamically *swap* flow ψ on the D arcs for exchange capacity on the E_y arcs. Since δ-augmenting paths contains only flow and relaxation arcs, the only time that flow is "augmented" on exchange arcs is through swapping. It might also seem that on E_y, (3) should be maintained only for the subset of arcs of E_y with exchange capacity at least δ. However, swapping exchange capacity for flow makes it possible to include *all* arcs of E_y in (3), which is important for proving that $y \in \mathrm{B}(f_p)$ is preserved, see Corollary 5.7.

Since $x = y - \partial\psi$ and we augment by exactly δ, each augmentation decreases the discrepancy by δ at both endpoints of the augmenting path, and maintains the discrepancy of all other nodes. A phase ends when one of $S^+(\delta)$, $S^-(\delta)$ is empty. Since the net excess is 0, this implies that at the end of a phase the total discrepancy is $\Phi \leq 2n\delta$.

At the start of a new phase, we modify ψ to satisfy the capacity constraints for the new value of δ, and modify φ to satisfy the reduced cost optimality constraints (3) for arcs a with residual capacity $\delta \leq r(a) < 2\delta$. After these adjustments, the initial discrepancy in a δ-phase is at most $4n\delta + 2n^2\delta + 4m\delta$. Thus the total number of augmentations in any δ phase is at most $3n^2 + 2n = \mathrm{O}(n^2)$, since each augmentation decreases the discrepancy by 2δ. Note that the total discrepancy is roughly halved in each phase.

Initially, we set $\delta = 2^i$ for $i = \lceil \log_2(\Phi/2n^2) \rceil \leq \lceil \log(U) \rceil$. Thus, after at most $\log(U)$ scaling phases, $\delta = 1$ and $\partial\varphi = x \in \mathrm{B}(f_p + b)$. We then remove the flow ψ and its contribution to x so that $x = y \in \mathrm{B}(f_p)$ and the discrepancy Φ is at most $2n^2$. At this point, we can call SSP to obtain an optimal solution in at most n^2 additional augmentations. The following theorem summarizes this discussion.

Theorem 5.1 *Algorithm* MCSF *calls* SubmodularDijkstra $\mathrm{O}(n^2 \log U)$ *times.* ■

MCSF(G, f)

`Input:` Directed graph $G = (V, A, l, u, c)$ and submodular function $f : 2^V \to \mathbf{Z}$
`Output:` Flow φ minimizing $c^\top \varphi$ over all $l \le \varphi \le u$ and $\partial\varphi \in \mathrm{B}(f)$

Initialize
  $p =$ shortest path distances in G°
  $\varphi =$ flow that satisfies condition $\langle$a$\rangle$ of Theorem 3.2
  $x = p$-maximum base in $\mathrm{B}(f)$
  $y = x$, $\psi \equiv 0$
  $\delta = 2^{\lceil \log_2(\Phi/2n^2) \rceil}$
while $\delta \ge 1$
  for all $a \in D$, $\psi(a) = \min\{\psi(a), \delta\}$.
  $x = y - \partial\psi$
  for all $a \in A_\varphi(\delta)$,
    if $c_p(a) < 0$, **then** send $r(a)$ units of flow through a, and update φ accordingly.
  while $S^+(\delta) \ne \emptyset$ and $S^-(\delta) \ne \emptyset$
    $(P, d) =$ SubmodularDijkstra$(\widehat{G}(\varphi, \psi), \delta, p, y)$
      [*P shortest path from $S^+(\delta)$ to $S^-(\delta)$ of residual capacity $\ge \delta$.*]
      [*d shortest distance labels*]
    Augment δ on P. [*Update φ, ψ, x.*]
    for all $v \in V$, $p(v) = p(v) + \min\{d(v), c_p(P)\}$
  $\delta = \delta/2$
$\varphi = \mathrm{SSP}(y, \varphi, p)$
return φ.

Invariants
  $c_p(a) \ge 0$ for all $a \in A_\varphi(\delta) \cup D_\psi \cup E_y$.
  $x = y - \partial\psi$.
  $y \in \mathrm{B}(f_p)$.

Figure 2: High level capacity scaling algorithm.

5.2 Maintaining Optimality Conditions

The keys to the correctness of the algorithm are maintaining the reduced cost optimality conditions (3) for all flow arcs with residual capacity at least δ and all exchange capacity arcs, and maintaining $y \in \mathrm{B}(f_p)$ for the current value of p. This requires extra care in choosing an augmenting path. In the subroutine SubmodularDijkstra (see Figure 3), we modify Dijkstra's algorithm to maintain (3) for exchange capacity arcs. The modification entails a more careful treatment of the arcs emanating from the node w that is added to the set of permanently labeled nodes in an iteration.

For each zero reduced cost exchange capacity arc $a = (w, v)$ with tail w, we swap exchange capacity on a for flow on its parallel relaxation arc so as to preserve the invariant $x = y - \partial\psi$. This has two possible outcomes, both of which make a irrelevant to future shortest path operations. One outcome happens when $\psi(w, v) \le \sigma(y, v, w)$. In this case we drive $\psi(w, v)$ to zero, so that the residual capacity of (w, v) becomes δ, and we can use $(w, v) \in D$ instead of a in any δ-augmenting

SubmodularDijkstra($\widehat{G}(\varphi,\psi), \delta, p, y$)

`Input:` Residual graph $\widehat{G}(\varphi,\psi)$ of φ and ψ; scale factor δ; dual prices p; and $y \in \mathrm{B}(f_p)$.
`Output:` A δ-augmenting path P from $S^+(\delta)$ to $S^-(\delta)$.

Initialization
  $d(v) = +\infty \quad \forall v \in V - S^+(\delta)$
  Select $w \in S^+(\delta)$, and set $d(w) = 0$.
  $R = \emptyset$ [*set of permanently labeled nodes*]
while $R \cap S^- = \emptyset$
  $w =$ node in $V - R$ with smallest label.
  $R := R \cup \{w\}$
  for all $a = (w, v) \in E_y$ with $v \in V - R$, $c_p(a) = 0$, and $\psi(w, v) > 0$,
    [*swap zero reduced cost, positive exchange capacity arc* (w, v) *leaving* R]
    $\alpha = \min\{\psi(w, v), \sigma(y, v, w)\}$.
    $y = y + \alpha(\vec{v} - \vec{w})$
    $\psi(w, v) = \psi(w, v) - \alpha$
  for all $a = (w, v) \in A_\varphi(\delta) \cup D_\psi$ with $v \in V - R$,
    $d(v) := \min\{d(v), d(w) + c_p(a)\}$
return path P from s to w on nodes in R.

Invariants
  $d(\partial^- a) \le d(\partial^+ a) + c_p(a)$ for all $a \in A_\varphi(\delta) \cup D_\psi \cup E_y$
  $x = y - \partial\psi$

Figure 3: Subroutine to find a δ-augmenting path.

path. The other outcome happens when $\psi(w, v) \ge \sigma(y, v, w)$. In this case we saturate a, making $\sigma(y, v, w) = 0$, so that a cannot be used in any shortest path. This swap operation is in some sense a generalization of the AdjustFlow routine from [18].

The swap operation is applied to potential exchange capacity arcs with 0 reduced costs. If the reduced cost of an exchange capacity arc is positive, then its opposite relaxation arc cannot have residual capacity δ (otherwise it would be a negative reduced cost arc in D_ψ, violating (3)). Hence its parallel relaxation arc has residual capacity δ, with the same reduced cost. This relaxation arc can always be chosen instead of the exchange capacity arc in any δ-augmenting path.

Swapping on an exchange capacity arc can lead to the creation of new exchange capacity arcs (see Lemma 2.5). We must therefore argue that after swapping on a 0-reduced cost exchange capacity arc: $\langle$A$\rangle$ the new exchange capacity graph maintains the non-negative reduced cost condition (3); $\langle$B$\rangle$ the labels assigned to nodes by Dijkstra's algorithm continue to reflect shortest path distances; and $\langle$C$\rangle$ the algorithm terminates, i.e., we do not pursue a cycle of exchange capacity arc swappings. In the following sequence of lemmas, R is the set of permanently labeled nodes from SubmodularDijkstra.

Lemma 5.2 *When* MCSF *swaps a 0-reduced cost exchange capacity arc,* $\langle$A$\rangle$ *the reduced cost optimality condition* (3) *is maintained, and* $\langle$B$\rangle$ *the distance labels* d *computed so far remain valid.*

Proof. We must check (3) for any new exchange arcs created by swapping exchange capacity arc (w, v). Let y and r denote values *before* the swap.

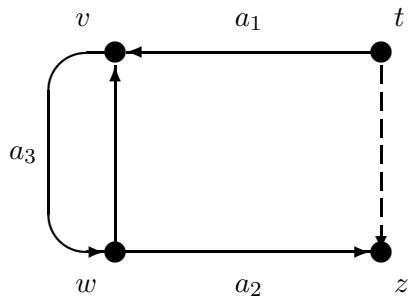

Figure 4: The path a_1, a_3, a_2 in the proof of Lemma 5.2.

The reverse exchange arc (v, w) clearly has reduced cost zero. If exchange arc (t, z) appears, then, by Lemma 2.5, exchange arcs $a_1 = (t, v)$ (if $t \neq v$) and $a_2 = (w, z)$ (if $w \neq z$) existed before swapping (w, v). Since SUBMODULARDIJKSTRA swaps (w, v), the relaxation arc (w, v) had residual capacity $r(w, v) < \delta$, which implies that the reverse relaxation arc, $a_3 = (v, w)$, had residual capacity $r(v, w) = \delta$. Hence there is a directed path a_1, a_3, a_2 from t to z in the auxiliary graph $\widehat{G}(\varphi, \psi, y) = (V, A_\varphi(\delta) \cup D_\psi \cup E_y)$ before swapping (w, v) (see Figure 4), and this path must have non-negative reduced cost. Since all arcs in this path have original cost 0, and the new exchange arc also has original cost 0, its reduced cost equals the reduced cost of this path, which is nonnegative. Thus $\langle$A$\rangle$ is shown.

The new exchange capacity arc has reduced cost equal to the length of the path a_1, a_3, a_2, thus its creation cannot make an existing label invalid, showing $\langle$B$\rangle$. ■

Lemma 5.3 *If we swap exchange capacity arcs leaving w, then no additional exchange capacity arcs leaving w are introduced, hence $\langle$C$\rangle$ we do not pursue a cycle of exchange capacity arc swaps.*

Proof. This follows trivially from Lemma 2.5. ■

The above two lemmas ensure that the path P found by SUBMODULARDIJKSTRA is a shortest path from a node in $S^+(\delta)$ to $S^-(\delta)$ in the auxiliary graph $\widehat{G}(\varphi, \psi, y) = (V, A_\varphi(\delta) \cup D_\psi \cup E_y)$ for the current ψ and y.

Theorem 5.4 *Subroutine* SUBMODULARDIJKSTRA *runs in* $O(n^2 h)$ *time.*

Proof. Each iteration adds one node to R and calls the exchange capacity oracle at most n times. ■

To complete the proof of correctness, we show that the algorithm maintains the claimed invariants that $x = y - \partial\psi$, the reduced cost optimality condition (3), and $y \in B(f_p)$.

Lemma 5.5 *The algorithm maintains $x = y - \partial\psi$.*

Proof. Every time ψ changes, x is updated. The only changes to y occur in SUBMODULARDIJKSTRA, after which ψ is modified to compensate for the corresponding change in y. ■

Lemma 5.6 *The reduced cost δ-optimality condition* (3) *is maintained after each augmentation.*

Proof. We must show that all flow arcs with residual capacity at least δ, and all exchange capacity arcs with exchange capacity at least δ have non-negative reduced cost after the augmentation and update of p using the distance labels d obtained from SubmodularDijkstra. Let p be the old potential, and q the new potential. For any previously existing arc a with non-negative reduced cost, we have $c_q(a) = c_p(a) + \min\{d(\partial^+ a), c_p(P)\} - \min\{d(\partial^- a), c_p(P)\}$. Since the distance labels satisfy $c_p(a) + d(\partial^+ a) \geq d(\partial^- a)$, we have $c_p(a) + \min\{d(\partial^+ a), c_p(P)\} \geq \min\{c_p(a) + d(\partial^+ a), c_p(P)\} \geq \min\{d(\partial^- a), c_p(P)\}$, which implies $c_q(a) \geq 0$. If there is an arc with residual capacity newly at least δ, it is the reverse arc of an arc a on P. Since P is a shortest path, $c_q(a) = 0$, so $c_q(\overline{a}) = 0$ also. ■

Corollary 5.7 *The condition that $y \in \mathrm{B}(f_p)$ is maintained after each augmentation.*

Proof. Lemma 2.3 says that $y \in \mathrm{B}(f_p)$ if $p_w \geq p_v$ for all $(w, v) \in E_y$, which is equivalent to saying that $c_p(w, v) \geq 0$ for all $(w, v) \in E_y$, which Lemma 5.6 established. ■

Lemma 5.8 *The reduced cost optimality condition* (3) *is maintained at the start of each δ-phase.*

Proof. At the start of each δ-phase, we modify both ψ and φ. The modification of ψ does not create any new arcs, so (3) is maintained. We modify φ precisely to enforce (3) for all flow arcs, by augmenting arcs for which (3) does not hold. ■

The above lemmas and theorems imply the main result of this section:

Theorem 5.9 *Algorithm* MCSF *computes a minimum cost submodular flow in $\mathrm{O}(n^4 h \log U)$ time.* ■

6 A Strongly Polynomial Algorithm

To make the capacity scaling algorithm described in the preceding section strongly polynomial we need to create large decreases in δ to avoid the $\Theta(\log U)$ number of scaling phases. To do this, we obtain a sequence of feasible flows and a geometrically decreasing sequence of values of δ for which these flows are respectively *approximately δ-optimal*: A triple $(\xi = (\varphi, \psi), y, p)$ is called approximately δ-optimal if $c_p(a) \geq 0$ for all arcs $a \in A_\varphi(\delta) \cup D_\psi \cup E_y$ and the discrepancy $\Phi \leq 2n\delta$. This definition exactly matches the conditions at the end of a δ-scaling phase.

Given a price vector p, the minimum value of δ for which there exists an approximately δ-optimal triple with prices p can be obtained by computing a maximum mean cut, from which we can also obtain the corresponding flow ξ and base x in the triple. The definition and computation of maximum mean cuts is discussed below in Section 6.1.

The broad outline of the strongly polynomial version of our algorithm is similar to other strongly polynomial algorithms [5, 18]. We divide the phases of the capacity scaling algorithm into *blocks*, so that each block contains $\mathrm{O}(\log n)$ phases. We start each block by computing an approximately δ-optimal triple using our current price vector p. In Section 6.2, we show that after performing $\mathrm{O}(\log n)$ scaling phases, this value of δ decreases enough so that we can *restrict the sign* of at least one reduced cost in the optimal solution ("restrict the sign" of $c_p(a)$ means that we can prove that $c_{p^*}(a) \geq 0$ or $c_{p^*}(a) \leq 0$ for every optimal p^*). Since there are only $m' := m + n(n-1) = \mathrm{O}(n^2)$ reduced costs (one for each arc of $I = A \cup E$), all reduced costs are fixed to zero after $2m' = O(n^2)$ blocks. At this point, we can compute an optimal flow with one maximum submodular flow computation. Thus the algorithm requires only $\mathrm{O}(n^2 \log n)$ phases overall. As in the weakly polynomial algorithm described in Section 5, the time per phase is strongly polynomial, $\mathrm{O}(n^4 h)$.

6.1 Maximum Mean Cuts and Approximations

In this section, we define maximum mean cuts, explain their usefulness in obtaining approximately δ-optimal triples, and describe how to compute them. Finally, we define a most positive cut, which is easier to compute and can be used to approximate a maximum mean cut. First some preliminary definitions.

Given a price vector p, we define *modified bounds* l_p and u_p as follows: If $c_p(a) > 0$ then $l_p(a) = u_p(a) = l(a)$; if $c_p(a) = 0$ then $l_p(a) = l(a)$ and $u_p(a) = u(a)$; and if $c_p(a) < 0$ then $l_p(a) = u_p(a) = u(a)$. This definition is consistent with Theorem 3.2 in that φ and p are primal and dual optimal iff $l_p \le \varphi \le u_p$ (equivalent to condition $\langle$a$\rangle$ of the theorem) and $\partial\varphi \in \mathrm{B}(f_p)$. A *cut* is a proper nonempty subset $X \subset V$. The *value* of cut X w.r.t. prices p is defined as $\kappa_p(X) := l_p(\Delta^+X) - u_p(\Delta^-X) - f_p(X)$. Theorems 3.1 and 3.2 imply that p is dual optimal if and only if there is no cut having positive value.

We will sometimes refer to the network restricted to arc set A, and at other times to the network with arc set $I = A \cup D$. We subscript Δ^+, Δ^-, and ∂ to indicate the relevant arc set. Since the relaxation uses arcs in D, the cut value should reflect this, and so we always discuss κ_p relative to I. Since δ modifies only the capacities of arcs in D, we define the mean value of a cut S using a denominator of $b(X) := |X| \cdot |V - X|$, the number of arcs of D leaving X. Thus the *mean value* of cut X is $\overline{\kappa}_p(X) = \kappa_p(X)/b(X)$. A *maximum mean cut* T satisfies $\overline{\kappa}_p(T) = \max_X \overline{\kappa}_p(X)$. For fixed p, we denote this value by $\overline{\kappa}_p$. This definition of maximum mean cut differs from the definition in [18] because we relax capacities differently.

The algorithm contains $\mathrm{O}(n^2)$ blocks, and hence computes $\mathrm{O}(n^2)$ max mean cuts. A maximum mean cut can be computed by discrete Newton's algorithm [21, 22] using a maximum submodular flow subroutine. This algorithm requires $\mathrm{O}(m') = \mathrm{O}(n^2)$ calls to submodular maximum flow. Each maximum submodular flow can be computed in $\mathrm{O}(n^3h)$ time using Fujishige and Zhang's algorithm [14, 19], for a total of $\mathrm{O}(n^5h)$ time per max mean cut, and $\mathrm{O}(n^7h)$ time for computing max mean cuts overall. This would dominate the time for the rest of the algorithm, which we will show is only $\mathrm{O}(n^6h \log n)$, so we use a trick of Radzik from [5] to remove this bottleneck: Just one call to submodular max flow produces a cut T° maximizing $\kappa_p(T)$, a *most positive cut*. If T^* is a max mean cut, then $\kappa_p(T^\circ) \ge \kappa_p(T^*)$, which implies that $\overline{\kappa}_p = \overline{\kappa}_p(T^*) \le \overline{\kappa}_p(T^\circ)\frac{b(T^\circ)}{b(T^*)} \le n\overline{\kappa}_p(T^\circ)$, so that $\overline{\kappa}_p(T^\circ) \ge \overline{\kappa}_p/n$. Thus T° approximates a max mean cut within a factor of n. Using most positive cuts in place of max mean cuts slightly increases the number of iterations in a block (but not enough to inflate the bound), and decreases the total time spent computing approximations to max mean cuts from $\mathrm{O}(n^7h)$ to $\mathrm{O}(n^5h)$, so that this is no longer a bottleneck.

6.2 Proximity Lemmas

In this section we prove the proximity lemmas which enable us to show that after a large enough decrease in δ, the sign of the reduced cost of some arc in the optimal solution is determined. Since the number of phases required to obtain this sufficiently large decrease is strongly polynomial, and the run time of each phase is strongly polynomial, this implies a strongly polynomial algorithm.

Lemma 6.1 (Primal Proximity) *Suppose that $(\xi = (\varphi, \psi), y, p)$ is approximately δ-optimal. Then there exists an optimal flow φ^* to the original instance such that*
$\langle$a$\rangle$ *$|\varphi(a) - \varphi^*(a)| \le 3n^2\delta$ for all $a \in A$, and*
$\langle$b$\rangle$ *for any cut X, $|\partial\varphi^*(X) - \partial\xi(X)| \le 3n^3\delta$.*

Proof. First we convert ξ to a flow on A exactly satisfying complementary slackness. We do this by saturating all residual A arcs with $c_p(a) < 0$; since all such arcs have $r(a) < \delta$, this

increases the discrepancy Φ by at most $2m\delta$. Then we delete all the D arcs and their flows ψ. Since such arcs have flow at most δ, this increases Φ by at most $n(n-1)\delta$. Since we started out with $\Phi \leq 2n\delta$, we now have $\Phi \leq (3n^2 - n)\delta$ (all such bounds in this section are computed much looser than optimal to make the arguments easier to follow). Call the resulting flow φ'. Note that $|\varphi'(a) - \varphi(a)| \leq \delta$ for all $a \in A$. In addition $|\partial_A \varphi'(X) - \partial_I \xi(X)| \leq n^2\delta$ for all $X \subset V$, since the maximum number of arcs that crosses a set X is n^2.

Now use SSP to convert φ' into an optimal flow φ^*. (As originally stated, SSP is finite only for integer data, but the lexicographic techniques of Schönsleben [23] and Lawler and Martel [20] show that it can be made finite for any data.) Since $\Phi \leq (3n^2 - n)\delta$, the total amount of flow pushed by SSP is at most $(3n^2-n)\delta$. This implies that for each $a \in A$, $|\varphi'(a)-\varphi^*(a)| \leq (3n^2-n)\delta$. Since $|\varphi'(a) - \xi(a)| \leq \delta$ we get $|\xi(a) - \varphi^*(a)| \leq 3n^2\delta$, verifying $\langle$a$\rangle$.

Suppose that P is an augmenting path chosen by the algorithm, and that flow is augmented by amount τ_P along P. By transitivity (Lemma 2.4), P contains at most $n/2$ arcs in E_y. Thus the boundary of any $X \subseteq V$ changes by at most $\frac{n}{2}\tau_P$ due to P. Since $\sum_P \tau_P \leq (3n^2 - n)\delta$, we have that $|\partial_A \varphi'(X) - \partial_A \varphi^*(X)| \leq \frac{3n^3}{2}\delta$. Since $|\partial_A \varphi'(X) - \partial_I \xi(X)| \leq n^2\delta$, this implies that $|\partial_I \xi(X) - \partial_A \varphi^*(X)| \leq 3n^3\delta$, verifying $\langle$b$\rangle$. ■

Corollary 6.2 *Suppose that (ξ, x, p) is approximately δ-optimal. Then for all optimal prices p^*,*
$\langle$a1$\rangle$ if there is an arc $a \in A$ with $\varphi(a) < u(a) - 3n^2\delta$, then $c_{p^*}(a) \geq 0$;
$\langle$a2$\rangle$ if there is an arc $a \in A$ with $\varphi(a) > l(a) + 3n^2\delta$, then $c_{p^*}(a) \leq 0$;
$\langle$b$\rangle$ if there is a cut T° and prices p° such that $\partial\xi(T^\circ) > f_{p^\circ}(T^\circ) + 3n^3\delta$, then there exists a pair (w, v) with w contained in some level set L° of p°, $v \notin L^\circ$, such that $p^*(w) \leq p^*(v)$.

Proof. For $\langle$a1$\rangle$, $\varphi(a) < u(a) - 3n^2\delta$ and Lemma 6.1 $\langle$a$\rangle$ imply that $\varphi^*(a) < u(a)$. Then complementary slackness implies that $c_{p^*}(a) \leq 0$ for all optimal prices p^*. A symmetric argument implies $\langle$a2$\rangle$.

For $\langle$b$\rangle$, $\partial\xi(T^\circ) > f_{p^\circ}(T^\circ) + 3n^3\delta$ and Lemma 6.1 $\langle$b$\rangle$ imply that $\partial\varphi^*(T^\circ) > f_{p^\circ}(T^\circ)$. We claim that this implies that there must be some level set L° of p° such that $\partial\varphi^*(L^\circ) < f(L^\circ)$: If not, $\partial\varphi^*(L) = f(L)$ for all level sets L of p°. In this case, Theorem 2.1 implies that $\partial\varphi^* \in \mathrm{B}(f_{p^\circ})$, which in turn implies that $\partial\varphi^*(T^\circ) \leq f_{p^\circ}(T^\circ)$, a contradiction. Given L° with $\partial\varphi^*(L^\circ) < f(L^\circ)$, suppose $w \in L^\circ$ and $v \notin L^\circ$ implies $p^*(w) > p^*(v)$. Then L° is a level set also of p^*, implying by Theorem 3.2 $\langle$b$\rangle$, Theorem 2.1, and Theorem 2.2 that $\partial\varphi^*(L^\circ) = f(L^\circ)$, again a contradiction. Thus $p^*(w) \leq p^*(v)$ for at least one pair (w, v) with $w \in L^\circ$ and $v \notin L^\circ$. ■

Let p° be the prices at the beginning of a block. We first compute a most positive cut T° with mean value $\delta^\circ > 0$. Since $\delta^\circ = \overline{\kappa}_p(T^\circ) \geq \overline{\kappa}_p/n$, we can compute a feasible submodular flow ξ° in the network with data l_{p°, u_{p°, and f° defined as $f^\circ(X) := f_{p^\circ}(X) + n\delta^\circ b(X)$. With $x^\circ = \partial\xi^\circ$, we have that ξ°, x°, and p° are approximately $n\delta^\circ$-optimal, and we start a block of $\mathrm{O}(\log n)$ phases with these values. Let p', ξ', and δ' be the corresponding values after the last phase in the block.

Theorem 6.3 *After a block of $\log_2(5n^5) = \mathrm{O}(\log n)$ scaling phases, at least one of the conditions of* Corollary 6.2 *is satisfied for $\delta = \delta'$, and thus the sign of one additional reduced cost can be newly restricted.*

Proof. Since each scaling phase cuts δ in half,

$$\delta' < n\delta^\circ/5n^5 = \delta^\circ/5n^4. \tag{4}$$

Conservation of flow for ξ' implies that $\xi'(\Delta_I^+ T^\circ) - \xi'(\Delta_I^- T^\circ) - \partial_I \xi'(T^\circ) = 0$. Subtracting this from $\kappa_{p^\circ}(T^\circ)$ yields

$$\begin{aligned} \delta^\circ b(T^\circ) &= \kappa_{p^\circ}(T^\circ) \\ &= (l_{p^\circ} - \xi')(\Delta_I^+ T^\circ) + (\xi' - u_{p^\circ})(\Delta_I^- T^\circ) \\ &\quad + (\partial_I \xi'(T^\circ) - f_{p^\circ}(T^\circ)). \end{aligned} \tag{5}$$

We claim that one of the following holds:

- $\exists\ a \in \Delta_A^+ T^\circ$ with $l_{p^\circ}(a) - \varphi'(a) \geq 3n^2\delta'$,
- $\exists\ a \in \Delta_A^- T^\circ$ with $\varphi'(a) - u_{p^\circ}(a) \geq 3n^2\delta'$,
- $\exists\ \partial_I \xi'(T^\circ) - f_{p^\circ}(T^\circ) \geq 3n^3\delta'$.

If not, then from (5) (noting that $|l_{p^\circ}(a) - \psi'(a)| \leq \delta'$ and $|\psi'(a) - u_{p^\circ}(a)| \leq \delta'$ for $a \in D$ and using very rough estimates) we get $\delta^\circ < \delta^\circ b(T^\circ) = \kappa_{p^\circ}(T^\circ) < \delta'(3n^2|\Delta_I(T^\circ)| + 3n^3) < 5n^4\delta'$, contradicting (4).

If $a \in \Delta_A^+ T^\circ$ and $l_{p^\circ}(a) - \varphi'(a) \geq 3n^2$, then we must have that $l_{p^\circ}(a) = u(a)$, implying that $c_{p^\circ}(a) < 0$. Thus $\varphi'(a) < u(a) - 3n^2\delta'$, and from Corollary 6.2 $\langle$a1$\rangle$ we can conclude that $c_{p^*}(a) \geq 0$. Since we had $c_{p^\circ}(a) < 0$, this is a new sign restriction on c_p. Symmetric arguments show that if $\exists a \in \Delta_A^- T^\circ$ with $\varphi'(a) - u_{p^\circ}(a) \geq 3n^2\delta'$, then $c_{p^*}(a) \leq 0$ is a new sign restriction on c_p.

Finally, if $\partial_I \xi'(T^\circ) > f_{p^\circ}(T^\circ) + 3n^3\delta'$, then Corollary 6.2$\langle$b$\rangle$ implies there is a pair w, v with w in some level set L° of p° and $v \notin L^\circ$ such that all optimal p^* have $p^*(w) \leq p^*(v)$. Since L° is a level set of p°, we have $p^\circ(w) > p^\circ(v)$, so this is a new sign restriction on p. ■

Theorem 6.4 *The algorithm described in this section computes a minimum cost submodular flow in* $O(n^6 h \log n)$ *time.*

Proof. Each block imposes at least one new sign restriction on $c_{p^*}(a)$ for some $a \in I$. At most $2m'$ such sign restrictions can be imposed before p^* is completely determined, so after at most $2m' = O(n^2)$ blocks we know the optimal p^*. Each block requires $O(\log n)$ scaling phases, each of which takes $O(n^4 h)$ time. The time to compute a most positive cut, initial ξ°, and final submodular flow given p^* is $O(n^3 h)$, which is not a bottleneck. ■

Acknowledgments

We would like to thank the IPCO VII Organizing Committee for helping to arrange the opportunity for us to work together in Graz, and Maiko Shigeno for many valuable conversations and discussions on related topics.

References

[1] W. H. Cunningham and A. Frank. A primal-dual algorithm for submodular flows. *Math. Oper. Res.*, 10 (1985), 251–262, .

[2] J. Edmonds. Submodular functions, matroids, and certain polyhedra. *Combinatorial Structures and their Applications*, R. Guy, H. Hanani, N. Sauer, and J. Schönheim, eds., Gordon and Breach, 1970, 69–87.

[3] J. Edmonds and R. Giles. A min-max relation for submodular functions on graphs. *Ann. Discrete Math.*, 1 (1977), 185–204, .

[4] J. Edmonds and R. M. Karp. Theoretical improvements in algorithmic efficiency for network flow problems. *J. ACM*, 19 (1972), 248–264, .

[5] T. R. Ervolina and S. T. McCormick. Two strongly polynomial cut canceling algorithms for minimum cost network flow. *Discrete Appl. Math.*, 46 (1993) 133–165.

[6] A. Frank. Finding feasible vectors of Edmonds–Giles polyhedra. *J. Combin. Theory*, 36 (1984), 221–239, .

[7] A. Frank and É. Tardos. An application of simultaneous Diophantine approximation in combinatorial optimization. *Combinatorica*, 7 (1987), 49–65.

[8] A. Frank and É. Tardos. Generalized polymatroids and submodular flows. *Math. Programming*, 42 (1988), 489–563.

[9] S. Fujishige. Algorithms for solving the independent-flow problems, *J. Oper. Res. Soc. Japan*, 21 (1978), 189–204.

[10] S. Fujishige. Structures of polyhedra determined by submodular functions on crossing families. *Math. Programming*, 29 (1984), 125–141.

[11] S. Fujishige. A capacity-rounding algorithm for the minimum cost circulation problem — A dual framework of the Tardos algorithm. *Math. Programming*, 35 (1986), 298–308.

[12] S. Fujishige. *Submodular Functions and Optimization.* North-Holland, 1991.

[13] S. Fujishige, H. Röck, and U. Zimmermann. A strongly polynomial algorithm for minimum cost submodular flow problems. *Math. Oper. Res.*, 14 (1989), 60–69.

[14] S. Fujishige and X. Zhang. New algorithms for the intersection problem of submodular systems. *Japan J. Indust. Appl. Math.*, 9 (1992), 369–382.

[15] S. Iwata. A capacity scaling algorithm for convex cost submodular flows, *Math. Programming*, 76 (1997), 299–308.

[16] S. Iwata, L. Fleischer, and S. Fujishige. A strongly polynomial-time algorithm for minimizing submodular functions. Research Report 99-07, Discrete Mathematics and Systems Science, Osaka University, 1999.

[17] S. Iwata, S. T. McCormick, and M. Shigeno. A faster algorithm for minimum cost submodular flows. *Proceedings of the Ninth ACM/SIAM Symposium on Discrete Algorithms* (1998), 167–174.

[18] S. Iwata, S. T. McCormick, and M. Shigeno. A strongly polynomial cut canceling algorithm for the submodular flow problem. *Proceedings of the Seventh MPS Conference on Integer Programming and Combinatorial Optimization* (1999), 259–272.

[19] S. Iwata, S. T. McCormick, and M. Shigeno. A faster cost scaling algorithm for submodular flow. 1999.

[20] E. L. Lawler and C. U. Martel. Computing maximal polymatroidal network flows. *Math. Oper. Res.* **7** (1982) 334–347.

[21] S. T. McCormick and T. R. Ervolina. Computing maximum mean cuts. *Discrete Appl. Math.*, 52 (1994) 53–70.

[22] T. Radzik. Newton's method for fractional combinatorial optimization. *Proceedings of the 33rd IEEE Annual Symposium on Foundations of Computer Science* (1992), 659–669; Parametric flows, weighted means of cuts, and fractional combinatorial optimization. *Complexity in Numerical Optimization*, P. Pardalos, ed., World Scientific, 1993, 351–386.

[23] P. Schönsleben. *Ganzzahlige Polymatroid-Intersektions Algorithmen.* Dissertation, ETH Zürich, 1980.

[24] M. Shigeno, S. Iwata, and S.T. McCormick. Relaxed most negative cycle and most positive cut canceling algorithms for minimum cost flow. *Math. Oper. Res.*, submitted.

[25] É. Tardos. A strongly polynomial minimum cost circulation algorithm. *Combinatorica*, 5 (1985), 247–256.

[26] É. Tardos, C. A. Tovey, and M. A. Trick, Layered augmenting path algorithms. *Math. Oper. Res.*, 11 (1986), 362–370.