# A Faster Capacity Scaling Algorithm for Minimum Cost Submodular Flow 

Lisa Fleischer *

Satoru Iwata ${ }^{\dagger}$

S. Thomas McCormick ${ }^{\ddagger}$

August 1999; July 2001


#### Abstract

We describe an $\mathrm{O}\left(n^{4} h \min \left\{\log U, n^{2} \log n\right\}\right)$ 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 $\delta$. Capacities are relaxed by attaching a complete directed graph with uniform arc capacity $\delta$ 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 polyhedron of a relaxed submodular function. To reduce these discrepancies, we use a variant of the successive shortest path algorithm that augments flow along minimum cost paths of residual capacity at least $\delta$. 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 maximum 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.


[^0]
## 1 Introduction

The submodular flow problem, introduced by Edmonds and Giles [4], 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 [2]. It generalizes a cost-scaling primal-dual method. A variant of this algorithm, combined with the push/relabel algorithm of Fujishige and Zhang [18], runs in $\mathrm{O}\left(n^{4} h \log C\right)$ time [24]. This is the previous best weakly polynomial bound for submodular flow. A different cost-scaling approach is based on cycle canceling [22], and achieves a running time of $\mathrm{O}\left(n^{4} h \min \left\{\log (n C), n^{2} \log n\right\}\right)$. See [15] for a recent survey of submodular flow algorithms.

The dual approach to scaling costs is to instead scale the capacities. The prototype capacity scaling min cost flow algorithm is the Edmonds-Karp [5] 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 [19] 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 $\mathrm{O}\left(n^{7} h \log U\right)$ time.

An easier way to do capacity scaling for submodular flow was introduced in [23]. 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 $\mathrm{O}\left(n^{6} h \min \left\{\log (n U), n^{2} \log n\right\}\right)$ 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 [19], capacities are relaxed by attaching a complete directed graph with uniform arc capacity $\delta$ in each scaling phase. As in [23], we relax the problem by a parameter $\delta$, 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 polyhedron. 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 $\delta$. 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, $\mathrm{O}\left(n^{4} h \log U\right)$, that is faster than all other known submodular flow algorithms when $U$ is smaller than $C$ and $n^{\mathrm{O}\left(n^{2}\right)}$.

As in [23] we also show how to use (approximations of) maximum mean cuts to get a strongly polynomial version of the algorithm. The resulting algorithm runs in $\mathrm{O}\left(n^{6} h \log n\right)$ time, matching the best current strongly polynomial bound. We also show how our algorithms can be extended to work for the case of crossing submodular functions.

The first strongly polynomial algorithm for the submodular flow problem is due to Frank and Tardos [9]. This is an application of simultaneous Diophantine approximation and substantially generalizes the first strongly polynomial minimum cost flow algorithm of Tardos [33]. A
more direct generalization of the Tardos algorithm to the submodular flow problem is described by Fujishige, Röck, and Zimmermann [16] with the aid of the tree-projection method of Fujishige [13]. The latter algorithm combined with the improved primal-dual algorithm in [24] runs in $\mathrm{O}\left(n^{6} h \log n\right)$ time, the current best strongly polynomial bound, which is also achieved by the lexicographic cycle canceling algorithm in [22].

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 [21] to obtain the first combinatorial, polynomial-time algorithm for the fundamental problem of minimizing a submodular function, and used again in [20] to get the first fully combinatorial algorithm for the same problem.

In the same way that cost scaling is associated with cycle canceling, capacity scaling is associated with cut canceling (see [32] 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 [23]. 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 [23], thus answering the open question of finding such an algorithm. Thus the new algorithm merges the two seemingly different strands of research of [19] and [23].

## 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 [14] for readers interested in more details.

### 2.1 Base Polyhedra

Let $V$ be a finite set, and $\mathcal{D} \subseteq 2^{V}$ be a distributive lattice with $\emptyset, V \in \mathcal{D}$. A function $f: \mathcal{D} \rightarrow \mathbf{R}$ is submodular if

$$
\begin{equation*}
f(X)+f(Y) \geq f(X \cup Y)+f(X \cap Y) \tag{1}
\end{equation*}
$$

holds for all $X, Y \in \mathcal{D}$. We assume throughout the paper that $f(\emptyset)=0$. Then the base polyhedron $\mathrm{B}(f)$ of $f$ is defined by

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

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 $L_{i}=\left\{v \mid p(v) \geq p_{i}\right\}$, the $i$ th level set of $p$. A $p$-maximum base exists if and only if $L_{i} \in \mathcal{D}$ for every $i$. Let $f_{p}: \mathcal{D} \rightarrow \mathbf{R}$ be defined by

$$
f_{p}(X)=\sum_{i=1}^{k}\left\{f\left(\left(X \cap L_{i}\right) \cup L_{i-1}\right)-f\left(L_{i-1}\right)\right\},
$$

where $L_{0}=\emptyset$. The following theorems give properties of $\mathrm{B}\left(f_{p}\right)$ that we will need.
Theorem 2.1 ([2, Theorem 6]) This $f_{p}$ is a submodular function with $f_{p} \leq f$, and $x \in \mathrm{~B}\left(f_{p}\right)$ if and only if $x \in \mathrm{~B}(f)$ and $x\left(L_{i}\right)=f\left(L_{i}\right)$ for every $i$.

Theorem 2.2 ([14, Theorem 3.15]) A base $x \in \mathrm{~B}(f)$ is p-maximum if and only if $x \in \mathrm{~B}\left(f_{p}\right)$.

### 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

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

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 $\alpha$. If $\alpha=\sigma(x, v, w)$, then we say that this operation saturates $(w, v)$.

In common with most other submodular flow algorithms we assume that we have an exchange capacity oracle: the oracle takes input $x, v, w$ and outputs $\sigma(x, v, w)$ in $h$ time. Thus our running times will include a term in $h$ for the number of calls to this oracle. Recent combinatorial algorithms for submodular function minimization (see Schrijver [31] or Iwata, Fleischer, and Fujishige [21]) have shown that an exchange oracle can be implemented using only an evaluation oracle (whose input is $X \in \mathcal{D}$ and whose output is $f(X)$ ). This has in turn led to a new submodular flow algorithm which needs only an evaluation oracle, see Fleischer and Iwata [7].

Lemma 2.3 ([14, 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 $x\left(L_{i}\right)=f\left(L_{i}\right)$ 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 $X_{w v} \in \mathcal{D}$ such that $v \in X_{w v}, w \notin X_{w v}$, and $x\left(X_{w v}\right)=f\left(X_{w v}\right)$. Then by the submodularity of $f$ and

$$
L_{i}=\bigcup_{v \in L_{i}} \bigcap_{w \in V-L_{i}} X_{w v}
$$

we obtain $x\left(L_{i}\right)=f\left(L_{i}\right)$.

Lemma 2.4 The exchangeability graph is transitive. Namely, $(s, v) \in E_{x}$ and $(v, t) \in E_{x}$ imply $(s, t) \in E_{x}$.

Proof. Suppose to the contrary $(s, t) \notin E_{x}$. Then there must be an $X \in \mathcal{D}$ that satisfies $t \in X$, $s \notin X$, and $x(X)=f(X)$. If $v \in X$ this contradicts $(s, v) \in E_{x}$, but if $v \notin X$ it contradicts $(v, t) \in E_{x}$.

Lemma 2.5 ([14, Figure 4.4]) Suppose $(s, t) \notin E_{x}$ and $(s, t) \in E_{y}$ 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 $s \neq v$ implies $(s, v) \in E_{x}$ and $t \neq w$ implies $(w, t) \in E_{x}$.

Proof. Since $(s, t) \notin E_{x}$, there must be an $X \in \mathcal{D}$ that satisfies $t \in X, s \notin X$, and $x(X)=f(X)$. If $s \neq v$ and $(s, v) \notin E_{x}$, then there exists an $Y \in \mathcal{D}$ such that $v \in Y, s \notin Y$, 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 $t \in X \cup Y$ and $s \notin X \cup Y$ contradicts $(s, t) \in E_{y}$. The $t \neq w$ case is similar.

Lemma 2.6 ([14, Theorem 4.4]) For distinct nodes $s, t, v, w \in V$ and a base $x \in \mathrm{~B}(f)$, suppose $(w, v) \in E_{x},(s, t) \in E_{x}$, and $(w, t) \notin E_{x}$. Then $\sigma(y, t, s)=\sigma(x, t, s)$ holds for a base $y$ obtained by $y=x+\alpha(\vec{v}-\vec{w}) \in \mathrm{B}(f)$ with $0<\alpha \leq \min \{\sigma(x, t, s), \sigma(x, v, w)\}$.

Proof. Since $f(X)-y(X)<f(X)-x(X)$ only if $v \in X$ and $w \notin X$, and $\sigma(\cdot, t, s)$ only depends on sets $X$ with $t \in X, s \notin X$, it suffices to show that $v, t \in X$ and $s, w \notin X$ imply $f(X)-x(X) \geq$ $\sigma(x, t, s)+\alpha$. Since $(w, t) \notin E_{x}$, there exists a subset $Y \subseteq V$ such that $t \in Y, w \notin Y$, and $x(Y)=f(Y)$. Then $t \in X \cap Y$ and $s \notin X \cap Y$ imply $f(X \cap Y)-x(X \cap Y) \geq \sigma(x, t, s)$. 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 $f(X)-x(X) \geq \sigma(x, t, s)+\sigma(x, v, w) \geq \sigma(x, t, s)+\alpha$ by the submodularity of $f$.

Lemma 2.7 ([14, Theorem 4.5]) Let $x$ be a base of $f$. Suppose $\left(w_{1}, v_{1}\right),\left(w_{2}, v_{2}\right), \cdots,\left(w_{k}, v_{k}\right) \in$ $E_{x}$ have distinct end-nodes and that $i<j$ implies $\left(w_{i}, v_{j}\right) \notin E_{x}$. If $\alpha \leq \sigma\left(x, v_{i}, w_{i}\right)$ holds for every $i=1,2, \ldots, k$, then $y=x+\alpha \sum_{i=1}^{k}\left(\vec{v}_{i}-\vec{w}_{i}\right)$ is also a base of $f$.

Proof. By Lemma 2.6, $x^{\prime}=x+\alpha\left(\vec{v}_{1}-\vec{w}_{1}\right) \in \mathrm{B}(f)$ satisfies $\sigma\left(x^{\prime}, v_{i}, w_{i}\right)=\sigma\left(x, v_{i}, w_{i}\right)$ for $i=2, \ldots, k$. It follows from Lemma 2.5 that $i<j \operatorname{implies}\left(w_{i}, v_{j}\right) \notin E_{x^{\prime}}$. 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) $l: A \rightarrow \mathbf{R} \cup\{-\infty\}$ and $u: A \rightarrow \mathbf{R} \cup\{+\infty\}$ with $l \leq u$, and $\operatorname{costs} c \in \mathbf{R}^{A}$. A flow is a function $\varphi$ 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 $\Theta(n) \leq m \leq n(n-1)$. Given $\operatorname{arc}(w, v)=a \in A$, define $\partial^{+} a=w$ and $\partial^{-} a=v$. Given a set $X \subset V$, define $\Delta^{+} X:=\left\{a \in A \mid \partial^{+} a \in X, \partial^{-} a \notin X\right\}$ (the set of arcs leaving $X$ ), and $\Delta^{-} X:=\left\{a \in A \mid \partial^{+} a \notin X, \partial^{-} a \in X\right\}$ (the set of arcs entering $X$ ). Given a flow $\varphi$ and $X \subseteq V$, the boundary of $\varphi$ on $X$ is $\partial \varphi(X)=\varphi\left(\Delta^{+} X\right)-\varphi\left(\Delta^{-} X\right)$, the net flow leaving $X$. Let $f$ be a submodular function on a distributive lattice $\mathcal{D} \subseteq 2^{V}$ 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 [8]) There exists a submodular flow if and only if

$$
\begin{equation*}
l\left(\Delta^{+} X\right)-u\left(\Delta^{-} X\right) \leq f(X) \tag{2}
\end{equation*}
$$

holds for every $X \in \mathcal{D}$. 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 $c_{p}(a)=c(a)+p\left(\partial^{+} a\right)-p\left(\partial^{-} a\right)$ for each $a \in A$. The following optimality conditions for the submodular flow problem are due to Fujishige [14, Theorem 5.3]. Edmonds and Giles [4] show that the submodular flow problem is totally dual integral, which implies the integrality claim.

Theorem 3.2 A submodular flow $\varphi$ is optimal if and only if there exists $p \in \mathbf{R}^{V}$ such that:
$\langle\mathrm{a}\rangle$ 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 $\langle\mathrm{b}\rangle$ 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 [4] originally formulated the submodular flow problem for crossing submodular functions (defined below), which are more general than what we have assumed here. They used crossing submodular functions because one of their aims was to generalize the LucchesiYounger theorem on the directed cut covering problem [26] , and crossing functions provide a natural model in that case.

A pair of subsets $X, Y \subseteq V$ is called crossing if $X \cap Y \neq \emptyset, X-Y \neq \emptyset, Y-X \neq \emptyset$, and $X \cup Y \neq V$. A family $\mathcal{F} \subseteq 2^{V}$ is a crossing family if $X \cap Y \in \mathcal{F}$ and $X \cup Y \in \mathcal{F}$ hold for any crossing pair $X, Y \in \mathcal{F}$. A crossing submodular function is a function $f: \mathcal{F} \rightarrow \mathbf{R}$ that satisfies the submodular inequality (1) for any crossing pair $X, Y \in \mathcal{F}$.

For a crossing submodular function $f: \mathcal{F} \rightarrow \mathbf{R}$ with $\emptyset, V \in \mathcal{F}$ and $f(\emptyset)=0$, the base polyhedron $\mathrm{B}(f)$ is defined by

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

This base polyhedron may be empty. Fujishige [12] has provided a necessary and sufficient condition for $\mathrm{B}(f)$ to be nonempty, and showed that a nonempty base polyhedron $\mathrm{B}(f)$ of a crossing submodular function $f$ is identical to a base polyhedron $\mathrm{B}(\widetilde{f})$ of a submodular function $\tilde{f}$ defined on a distributive lattice that contains $\mathcal{F}$, also called a fully submodular function.

Thus, theorems and algorithms that rely only on geometric properties of the base polyhedron carry over to crossing submodular functions. For example, Theorem 3.2 is valid for the submodular flow problem for crossing submodular functions since it refers only to base polyhedra, and does not require an algorithm to compute a value for $\tilde{f}$, whereas Theorem 3.1 is not geometric because it requires a suitable modification that replaces $f$ in (2) by the fully submodular function $\widetilde{f}$. The equivalent theorem for the crossing case was obtained by Frank [8].

An exchange oracle must minimize $f(X)-x(X)$ over sets $X$ containing $v$ but not $w$. For crossing submodular functions these sets form a distributive lattice, so it is reasonable for us to assume that we have an exchange oracle in this case as well.

In the rest of this paper, we assume $f$ to be fully submodular for the sake of simplicity, but we indicate at the appropriate places that our results in fact extend also to the crossing case.

## 4 The Successive Shortest Path Algorithm

This section describes the successive shortest path (SSP) algorithm for the submodular flow problem, which is the basis for our algorithms. As with the 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 [11] for solving the minimum cost independent flow problem, which turns out to be equivalent to the submodular flow problem [14, §5.2].

In order to deal with the distributive lattice $\mathcal{D}$ algorithmically, we need to know a compact representation of $\mathcal{D}$. As Birkhoff's representation theorem [1] suggests, we represent $\mathcal{D}$ by a directed graph with the vertex $V$ and the arc set $E^{\circ}=\{(w, v) \mid \forall X \in \mathcal{D}: v \in X \Rightarrow w \in X\}$. Note that $\left(V, E^{\circ}\right)$ is transitive, i.e., $(s, v) \in E^{\circ},(v, t) \in E^{\circ}$ implies $(s, t) \in E^{\circ}$, and that $(w, v) \in E^{\circ}$ implies that for any $x \in \mathrm{~B}(f), \sigma(x, v, w)=\infty$, so that $(w, v) \in E_{x}$.

If $\left(V, E^{\circ}\right)$ had a directed cycle, by transitivity there would be $w, v \in V$ with both $(w, v)$ and $(v, w)$ in $E^{\circ}$, and so in $E_{x}$ for any $x \in \mathrm{~B}(f)$. Hence by Lemma 2.3 and Theorem 3.2, any optimal price $p^{*}$ must satisfy $p^{*}(v)=p^{*}(w)$, implying that if $a=(w, v) \in A$, any optimal flow $\varphi^{*}$ must satisfy $\varphi^{*}(a)=l(a)$ if $c(a)>0$, or $\varphi^{*}(a)=u(a)$ if $c(a)<0$. Therefore, shrinking such $v$ and $w$ to a single node does not change the problem. Thus we will assume in the rest of this paper that $\left(V, E^{\circ}\right)$ is acyclic.

We initialize by looking for a flow $\varphi$ and a price vector $p$ that satisfy condition $\langle\mathrm{a}\rangle$ of Theorem 3.2. We first check whether there exists a dual feasible solution $p$. To do this we construct a directed graph $G^{\circ}=\left(V, A^{\circ}\right)$ with the arc set $A^{\circ}=F^{\circ} \cup B^{\circ} \cup E^{\circ}$ defined by $F^{\circ}=\{a \mid u(a)=+\infty\}$ and $B^{\circ}=\{\bar{a} \mid l(a)=-\infty\}$, where $\bar{a}$ is the reversal of $a$. Each arc in $A^{\circ}$ has length $c(a)$ if $a \in F^{\circ}$, $-c(\bar{a})$ if $a \in B^{\circ}$, and 0 if $a \in E^{\circ}$. We use the $\mathrm{O}(m n)$ Bellman-Ford-Moore shortest path algorithm to try to compute shortest path distances $p$ in $G^{\circ}$. If the algorithm finds a negative cycle in $G^{\circ}$, then the instance is dual infeasible [14, 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 $\varphi$ that satisfies condition $\langle\mathrm{a}\rangle$ of Theorem 3.2 in $\mathrm{O}(m)$ time.

We also start with a base $x \in \mathrm{~B}\left(f_{p}\right)$, which can be obtained using the greedy algorithm [3, 17]. If $x=\partial \varphi$, then $x$ satisfies condition $\langle\mathrm{b}\rangle$ of Theorem 3.2, proving that $\varphi$ is optimal (and that $p$ is dual optimal). The algorithm will iteratively modify $\varphi, x$, and $p$ so that $\langle\mathrm{a}\rangle$ is preserved for $\varphi$ and $p,\langle\mathrm{~b}\rangle$ is preserved for $x$ and $p$, and so that $\partial \varphi$ and $x$ eventually converge.

We construct an auxiliary network $\widehat{G}(\varphi, x)=\left(V, A_{\varphi} \cup E_{x}\right)$ with respect to flow $\varphi$ and base $x \in \mathrm{~B}\left(f_{p}\right)$. 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 $B_{\varphi}:=\{\bar{a} \mid a \in A, \varphi(a)>l(a)\}$ (i.e., $A_{\varphi}$ is just the ordinary residual graph of flow $\varphi$ ). As defined in Section 2.2, $E_{x}$ is the arc set of the exchangeability graph with respect to $\mathrm{B}(f)$ (we emphasize that this is $\mathrm{B}(f)$ rather than $\mathrm{B}\left(f_{p}\right)$ 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

$$
r(a):=\left\{\begin{array}{cl}
u(a)-\varphi(a) & \left(a \in F_{\varphi}\right) \\
\varphi(a)-l(a) & \left(a \in B_{\varphi}\right) \\
\sigma(x, v, w) & \left(a=(w, v) \in E_{x}\right),
\end{array} \quad c(a):=\left\{\begin{array}{cl}
c(a) & \left(a \in F_{\varphi}\right) \\
-c(\bar{a}) & \left(a \in B_{\varphi}\right) \\
0 & \left(a \in E_{x}\right) .
\end{array}\right.\right.
$$

The algorithm computes augmentations of $\varphi$ and $x$ with respect to $r$ on $A_{\varphi} \cup E_{x}$. This will ensure that $\varphi$ remains a flow. Together with the corresponding update to $p$, it will also ensure that $x$ remains in $\mathrm{B}\left(f_{p}\right)$.

Each augmentation is determined as follows. We regard the reduced cost $c_{p}(a)=c(a)+$ $p\left(\partial^{+} a\right)-p\left(\partial^{-} a\right)$ as the length of $a \in A_{\varphi} \cup E_{x}$. 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 algorithm augments along a shortest path (with respect to $c_{p}$ ) from $S^{+}$to $S^{-}$with a minimum number of arcs. Since $\varphi$ and $p$ satisfy $\langle\mathrm{a}\rangle$, which is equivalent to having $c_{p}(a) \geq 0$ for all $a \in A_{\varphi}$, and since $x \in \mathrm{~B}\left(f_{p}\right)$, which is equivalent to having $c_{p}(a) \geq 0$ for $a \in E_{x}$, we can compute shortest path distances $d(v)$ from $S^{+}$to each $v \in V$ using Dijkstra's algorithm. Among the nodes in $S^{-}$with minimum $d(v)$, the algorithm selects a path $P$ with a minimum number of arcs, and augments $\varphi$ on $P$ by $\alpha:=\min \{x(s)-\partial \varphi(s), \partial \varphi(t)-x(t), \min \{r(a) \mid a \in P\}\}$. Since $\varphi$ is augmented only on $P$ and not on any longer paths from $S^{+}$to $S^{-}$, the algorithm updates $p(v)$ by $\min \left(d(v), c_{p}(P)\right)$ for each
$\operatorname{SSP}(x, \varphi, p)$
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 $p(v):=p(v)+\min \left\{d(v), c_{p}(P)\right\}$.
(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 $\operatorname{arc} a \in P$,

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

Figure 1: Successive Shortest Path algorithm.
$v \in V$. This augmentation of $\varphi$ and update of $p$ preserves $\langle\mathrm{a}\rangle$ and $\langle\mathrm{b}\rangle$, decreases $|\partial \varphi(v)-x(v)|$ by the augmentation amount for the initial and final nodes of $P$, and does not change $|\partial \varphi(v)-x(v)|$ for any other node. The details of this algorithm SSP are 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 $l\left(\Delta^{+} X\right)-u\left(\Delta^{-} X\right)=\varphi\left(\Delta^{+} X\right)-\varphi\left(\Delta^{-} X\right)=\partial \varphi(X)>x(X)=f(X)$, 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) of Figure 1. 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) of Figure 1. Furthermore, Lemma 2.5 can be used to show that $c_{p}$ remains nonnegative.

Lemma 4.1 There exists a numbering $\left(w_{1}, v_{1}\right),\left(w_{2}, v_{2}\right), \cdots,\left(w_{k}, v_{k}\right)$ of the arcs in $P \cap E_{x}$ such that $i<j$ implies $\left(w_{i}, v_{j}\right) \notin E_{x}$.

Proof. If no such numbering exists, there is a subset $\left\{\left(w_{j}, v_{j}\right) \mid j=1, \ldots, q\right\}$ of $P \cap E_{x}$ such that $e_{j}=\left(w_{j-1}, v_{j}\right) \in E_{x}$ holds for each $j$, where $w_{0} \equiv w_{q}$. After Step (1-2) of Figure 1, we have $p\left(w_{j}\right)=p\left(v_{j}\right)$, which implies $\sum_{j=1}^{q} c_{p}\left(e_{j}\right)=0$. Therefore $c_{p}\left(e_{j}\right)=0$ holds for each $j=1, \ldots, q$, which contradicts the minimality of $P$.

If $l, u$ and $f$ are integer valued, then initial $x$ and $\varphi$ are integer, hence $\alpha$ is always integer and at least 1 , so that $\sum_{v \in V}|x(v)-\partial \varphi(v)|$ 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 $\left|A_{\varphi}\right|+\left|E_{x}\right|=O(m)+\mathrm{O}\left(n^{2}\right)=\mathrm{O}\left(n^{2}\right)$, the bottleneck Dijkstra computation takes $\mathrm{O}\left(n^{2} h\right)$ time, so each augmentation requires $\mathrm{O}\left(n^{2} h\right)$ time.

Most of the successive shortest path algorithm works for crossing submodular functions. The only difference is we need to compute an initial base $x$ using the bi-truncation algorithm of Frank and Tardos [10] (cf. [28]) instead of the greedy algorithm.

## 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 [5] propose a capacity scaling algorithm that maintains reduced cost optimality conditions while rounding the arc capacities. We modify the Edmonds-Karp algorithm by attaching a complete directed graph on the node set, which effectively relaxes the submodular 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.

Neither the minimum cost flow version of SSP, nor our submodular flow version described in Section 4, runs in polynomial time, 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 $\delta$ units of flow. To modify this for submodular flow, we instead relax capacities by $\delta$. We do this by adding to our initial graph the arc set of a complete directed graph $(V, D)$ with arc set $D=\{(w, v) \mid w \neq v \in V\}$ with capacity $\delta$. 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 $b: 2^{V} \rightarrow \mathbf{R}$ by $b(X)=|X| \cdot|V-X|$. Thus $\delta b(X)$ is the capacity of the cut $X$ in $(V, D)$. This can be thought of as either relaxing the condition that $\partial \varphi \in \mathrm{B}(f)$ to $\partial \varphi \in \mathrm{B}(f+\delta b)$, as relaxing the capacities $l$ and $u$ by $\delta$ (viewing $A$ as a complete directed graph with some 0-capacity arcs), or both.

To avoid confusion we will always use $\varphi$ for a flow on the original set of flow arcs $A$, we will always use $\psi$ for a flow on the relaxation arcs $D$, and we will always use $\xi$ 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 $\delta$ 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 $\varphi$ on $A$, and flow $\psi$ on $D$, a price vector $p$, and a base $x \in \mathrm{~B}\left(f_{p}+\delta b\right)$. We maintain $x$ as the sum of $y \in \mathrm{~B}\left(f_{p}\right)$ and $-\partial \psi \in \mathrm{B}(\delta b)$. We compute exchange capacities with respect to $y$ and with respect to 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 $\partial \varphi$, defined by $\Phi=\sum_{v}|x(v)-\partial \varphi(v)|$. In a sense, we relax the submodular constraints on $\varphi$ twice: once in relaxing $f$ to $f+\delta b$, and once in letting $\partial \varphi$ and $x$ differ. Recall that, because the algorithm maintains $\langle\mathrm{a}\rangle$ of Theorem 3.2 for $\varphi$ and $p$, and $\langle\mathrm{b}\rangle$ of Theorem 3.2 for $y$ and $p$, if we can drive $\Phi$ to zero and $\psi$ to zero so that $x=y=\partial \varphi$, we are optimal.

The algorithm RCS-MCSF is described in Figure 2. It embeds a modified version of SSP in a scaling framework that depends on parameter $\delta$. We call the iterations with a fixed value of $\delta$ a $\delta$-scaling phase, or just a phase for short. 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 $\delta$ (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 $\delta$ 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 $\varphi$, base $x$, and price function $p$ as in SSP. For each $v \in V$, let $R(v)=\left\{w \mid(w, v) \in E^{\circ}\right\}$. Then $R(v) \cup\{v\}$ is the unique minimal member of $\mathcal{D}$ containing $v$. Define $U$ as

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

Since $|\varphi(a)| \leq U$ for all $a \in A,|\partial \varphi(v)| \leq(n-1) U$ for all $v \in V$. A $p$-maximum base $x$ comes from the greedy algorithm, which by Lemma 2.3 can order $R(v)$ before $v$, so we have that $x(v) \leq f(R(v) \cup\{v\})-f(R(v)) \leq U$ for each $v \in V$. Since $x(V)=f(V)=0$, we have $|x(v)| \leq(n-1) U$. Thus the initial discrepancy between $x$ and $\partial \varphi$ is at most $2 n^{2} U$.

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 $\delta$. Given such a flow $\psi$, let $D_{\psi}$ denote the set of relaxation arcs with residual capacity equal to $\delta$, i.e., $D_{\psi}=\{(w, v) \mid \psi(w, v)=0\}$. Define $A_{\varphi}(\delta):=\left\{a \mid a \in A_{\varphi}, r(a) \geq \delta\right\}$. 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 $\delta$ reduced cost optimality condition that

$$
\begin{equation*}
c_{p}(a) \geq 0 \text { holds for every } a \in A_{\varphi}(\delta) \cup D_{\psi} \cup E_{y} . \tag{3}
\end{equation*}
$$

We call such a path a $\delta$-augmenting path. It might seem that the algorithm should also include the subset of arcs of $E_{y}$ with exchange capacity at least $\delta$ in its list of arcs eligible for a $\delta$-augmenting path. Instead, SubmodularDijkstra performs exchange operations on the $E_{y}$ arcs in the search procedure, adjusting flow $\psi$ on the parallel $D$ arcs to maintain $x=y-\partial \psi$ without changing $x$, so that $\delta$-augmenting paths contain only flow and relaxation arcs. This procedure is described in detail in Section 5.2. Thus the only time that flow is "augmented" on exchange arcs is through this search procedure. 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 $\delta$. However, the exchanges between parallel $D$ and $E_{y}$ arcs makes it possible to include all arcs of $E_{y}$ in (3), which is important for proving that $y \in \mathrm{~B}\left(f_{p}\right)$ is preserved, see Corollary 5.7.

Since $x=y-\partial \psi$ and we augment by exactly $\delta$, each augmentation decreases the discrepancy by $\delta$ 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 $\sum_{v} x(v)-\partial \varphi(v)=0$, at the end of a phase we have $\Phi \leq 2 n \delta$.

At the start of a new phase, we modify $\psi$ to satisfy the capacity constraints for the new value of $\delta$, and modify $\varphi$ 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 $\delta$-phase is at most $4 n \delta+2 n^{2} \delta+4 m \delta$. Thus the total number of augmentations in any $\delta$ phase is at most $3 n^{2}+2 n=\mathrm{O}\left(n^{2}\right)$, since each augmentation decreases the discrepancy by $2 \delta$. Note that the total discrepancy is roughly halved in each phase.

Initially, we set $\delta=2^{i}$ for $i=\left\lceil\log _{2}\left(\Phi / 2 n^{2}\right)\right\rceil \leq\lceil\log U\rceil$. Thus, after at most $\log U$ scaling phases, $\delta=1$ and $\partial \varphi=x \in \mathrm{~B}\left(f_{p}+b\right)$. We then remove the flow $\psi$ and its contribution to $x$ so that $x=y \in \mathrm{~B}\left(f_{p}\right)$ and the discrepancy $\Phi$ is at most $2 n^{2}$. 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 RCS-MCSF calls SubmodularDijkstra $\mathrm{O}\left(n^{2} \log U\right)$ times.

```
Relaxed Capacity Scaling for Minimum Cost Submodular Flow
\(\operatorname{RCS}-\operatorname{MCSF}(G, f)\)
Input: Directed graph \(G=(V, A, l, u, c)\) and submodular function \(f: \mathcal{D} \rightarrow \mathbf{Z}\)
Output: Flow \(\varphi\) minimizing \(c^{\top} \varphi\) over all \(l \leq \varphi \leq u\) and \(\partial \varphi \in \mathrm{B}(f)\)
    Initialization:
        \(p:=\) shortest path distances in \(G^{\circ}\)
        \(\varphi:=\) flow that satisfies condition \(\langle\mathrm{a}\rangle\) of Theorem 3.2
        \(x:=p\)-maximum base in \(\mathrm{B}(f)\), i.e., \(x \in \mathrm{~B}\left(f_{p}\right)\)
        \(y:=x, \psi \equiv 0\)
        \(\delta:=2^{\left\lceil\log _{2}\left(\Phi / 2 n^{2}\right)\right\rceil}\)
    while \(\delta \geq 1\)
        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 \(\varphi\) accordingly.
        while \(S^{+}(\delta) \neq \emptyset\) and \(S^{-}(\delta) \neq \emptyset\)
            \((P, d):=\operatorname{SubmodularDijkstra}(\widehat{G}(\varphi, \psi), \delta, p, y)\)
            [ \(P\) : shortest path from \(S^{+}(\delta)\) to \(S^{-}(\delta)\) of residual capacity \(\geq \delta\). ]
            [ \(d\) : shortest distance labels. ]
            Augment by \(\delta\) on \(P\). [ Update \(\varphi, \psi, x\).]
            for all \(v \in V, p(v):=p(v)+\min \left\{d(v), c_{p}(P)\right\}\)
        \(\delta:=\delta / 2\)
    \(\varphi:=\operatorname{SSP}(y, \varphi, p)\)
    return \(\varphi\).
    Invariants:
        \(c_{p}(a) \geq 0\) for all \(a \in A_{\varphi}(\delta) \cup D_{\psi} \cup E_{y}\).
        \(x=y-\partial \psi\).
        \(y \in \mathrm{~B}\left(f_{p}\right)\).
```

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 $\delta$ and all exchange capacity arcs, and maintaining $y \in \mathrm{~B}\left(f_{p}\right)$ 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 $R$ of permanently labeled nodes in an iteration.

For each zero reduced cost exchange capacity arc $a=(w, v)$ with $v \notin R$, we do a doubleexchange that reduces the exchange capacity on $a$ and the flow $\psi$ on its parallel relaxation arc so as to preserve the invariant $x=y-\partial \psi$ without changing $x$. This has two possible outcomes, both of which make $a$ irrelevant to future shortest path operations in this iteration.

SUBMODULARDiJKstra $(\widehat{G}(\varphi, \psi), \delta, p, y)$
Input: Residual graph $\widehat{G}(\varphi, \psi)$ of $\varphi$ and $\psi$; scale factor $\delta$; dual prices $p$; and $y \in \mathrm{~B}\left(f_{p}\right)$.
Output: A $\delta$-augmenting path $P$ from $S^{+}(\delta)$ to $S^{-}(\delta)$.
Initialization:
$d(v):=+\infty \quad \forall v \in V-S^{+}(\delta)$
$d(v):=0 \quad \forall v \in S^{+}(\delta)$
$R:=\emptyset \quad$ [ set of permanently labeled nodes ]
while $R \cap S^{-}=\emptyset$ do
$w:=$ node in $V-R$ with smallest label.
$R:=R \cup\{w\}$
[ $\operatorname{scan} w$ ]
for all $a=(w, v) \in E_{y}$ with $v \in V-R, c_{p}(a)=0$, and $\psi(w, v)>0$ do
[ do double exchange on parallel $(w, v)$ arcs in $E_{y}$ and $D$ ]
$\alpha:=\min \{\psi(w, v), \sigma(y, v, w)\}$ $y:=y+\alpha(\vec{v}-\vec{w})$ $\psi(w, v):=\psi(w, v)-\alpha$
for all $a=(w, v) \in A_{\varphi}(\delta) \cup D_{\psi}$ with $v \in V-R$ do $d(v):=\min \left\{d(v), d(w)+c_{p}(a)\right\}$
return path $P$ from $s$ to $w$ on nodes in $R$.
Invariants:
$d\left(\partial^{-} a\right) \leq d\left(\partial^{+} a\right)+c_{p}(a)$ for all $a \in A_{\varphi}(\delta) \cup D_{\psi} \cup E_{y}$
$x=y-\partial \psi$

Figure 3: Subroutine to find a $\delta$-augmenting path.

One outcome happens when $\psi(w, v) \leq \sigma(y, v, w)$. In this case we drive $\psi(w, v)$ to zero, so that the residual capacity of $(w, v) \in D$ becomes $\delta$, and we can use $(w, v) \in D$ instead of $a$ in any $\delta$ augmenting path. The other outcome happens when $\psi(w, v) \geq \sigma(y, v, w)$. In this case we saturate $a$, making $\sigma(y, v, w)=0$, so that $a$ cannot be used in any shortest path. The double-exchange operation is in some sense a variant of the AdjustFlow routine from [23]. The correctness of SubmodularDijkstra mostly follows in the same way as the use of Dijkstra's Algorithm for ordinary SSP, except that we need to be careful about exchange arcs.

The double-exchange operation is applied only to 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 $\delta$ (otherwise it would be a negative reduced cost arc in $D_{\psi}$, violating (3)). Hence its parallel relaxation arc has residual capacity $\delta$, with the same reduced cost. This relaxation arc can always be chosen instead of the exchange capacity arc in any $\delta$-augmenting path.

The double-exchange operation can lead to the creation of new exchange capacity arcs (see Lemma 2.5). We must therefore argue that after the operation: $\langle\mathrm{A}\rangle$ the new exchangeability graph maintains the nonnegative reduced cost condition (3); $\langle\mathrm{B}\rangle$ the labels assigned to nodes by Dijkstra's algorithm continue to reflect shortest path distances; and $\langle\mathrm{C}\rangle$ the algorithm terminates, i.e., we do not pursue a cycle of double-exchange operations. In the following sequence of lemmas, $R$ is the set of permanently labeled nodes from SubmodularDiskstra.


Figure 4: The path $a_{1}, a_{3}, a_{2}$ in the proof of Lemma 5.2.

Lemma 5.2 When RCS-MCSF applies the double-exchange operation on a 0 -reduced cost exchange arc, $\langle\mathrm{A}\rangle$ the reduced cost optimality condition (3) is maintained, and $\langle\mathrm{B}\rangle$ the distance labels $d$ computed so far remain valid.

Proof. We must check (3) for any new exchange arcs created by the double-exchange applied to $(w, v)$. Let $y$ and $r$ denote values before the double-exchange.

The reverse exchange arc $(v, w)$ clearly has reduced cost zero. If exchange arc $(s, t)$ appears, then, by Lemma 2.5, exchange $\operatorname{arcs} a_{1}=(s, v)$ (if $s \neq v$ ) and $a_{2}=(w, t)$ (if $w \neq t$ ) existed before the double-exchange. Since SubmodularDijkstra applied double-exchange to $(w, v)$, 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 $s$ to $t$ in the auxiliary graph $\widehat{G}(\varphi, \psi, y)=\left(V, A_{\varphi}(\delta) \cup D_{\psi} \cup E_{y}\right)$ before the double-exchange (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 $\langle\mathrm{A}\rangle$ 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 $\langle\mathrm{B}\rangle$.

Lemma 5.3 In a call to SUBMODULAR DiJKStra, a double-exchange operation is applied to each arc at most once. Hence $\langle\mathrm{C}\rangle$ we do not pursue a cycle of double-exchange operations.

Proof. Arc $(w, v)$ participates in a double exchange only when its tail $w$ is being scanned in the while loop of Submodular Dijkstra. The only way for $(w, v)$ to participate twice is if the first double exchange causes $(w, v)$ to leave $E_{y}$ and a later double exchange of $(w, t)$ causes $(w, v)$ to re-enter $E_{y}$. But Lemma 2.5 shows that this cannot happen.

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

Theorem 5.4 Subroutine SubmodularDijkstra runs in $\mathrm{O}\left(n^{2} h\right)$ 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 $y \in \mathrm{~B}\left(f_{p}\right)$.

Lemma 5.5 The algorithm maintains $x=y-\partial \psi$.
Proof. Every time $\psi$ changes, $x$ is updated. The only changes to $y$ occur in the double-exchange operation in SUBMODULARDIJKSTRA, which includes a modification to $\psi$ to compensate for the corresponding change in $y$.

Lemma 5.6 The reduced cost $\delta$-optimality condition (3) is maintained after each augmentation.
Proof. We must show that all flow arcs with residual capacity at least $\delta$, and all exchange capacity arcs with exchange capacity at least $\delta$ 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 $c_{q}(a)=c_{p}(a)+\min \left\{d\left(\partial^{+} a\right), c_{p}(P)\right\}-\min \left\{d\left(\partial^{-} a\right), c_{p}(P)\right\}$. Since the distance labels satisfy $c_{p}(a)+d\left(\partial^{+} a\right) \geq d\left(\partial^{-} a\right)$, we have $c_{p}(a)+\min \left\{d\left(\partial^{+} a\right), c_{p}(P)\right\} \geq \min \left\{c_{p}(a)+\right.$ $\left.d\left(\partial^{+} a\right), c_{p}(P)\right\} \geq \min \left\{d\left(\partial^{-} a\right), c_{p}(P)\right\}$, which implies $c_{q}(a) \geq 0$. If there is an arc with residual capacity newly at least $\delta$, it is the reverse arc of an arc $a$ on $P$. Since $P$ is a shortest path, $c_{q}(a)=0$, so $c_{q}(\bar{a})=0$ also.

Corollary 5.7 The condition that $y \in \mathrm{~B}\left(f_{p}\right)$ is maintained after each augmentation.
Proof. Lemma 2.3 says that $y \in \mathrm{~B}\left(f_{p}\right)$ 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 $\delta$-phase.
Proof. At the start of each $\delta$-phase, we modify both $\psi$ and $\varphi$. The modification of $\psi$ does not create any new arcs, so (3) is maintained. We modify $\varphi$ 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 RCS-MCSF computes a minimum cost submodular flow in $\mathrm{O}\left(n^{4} h \log U\right)$ time.

The capacity scaling algorithm applies to crossing submodular flow problems. The only required modification is again the use of the bi-truncation algorithm for an initial base $x$.

## 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 $\delta$ to avoid the $\Theta(\log U)$ number of scaling phases. To do this, we obtain a sequence of feasible flows and a decreasing sequence of values of $\delta$ for which these flows are respectively $\delta$-optimal: A triple $(\xi=(\varphi, \psi), y, p)$ is called $\delta$-optimal if $c_{p}(a) \geq 0$ for all $\operatorname{arcs} a \in A_{\varphi}(\delta) \cup D_{\psi} \cup E_{y}$ (i.e. (3) holds) and the discrepancy $\Phi \leq 2 n \delta$. This definition exactly matches the conditions at the end of a $\delta$-scaling phase.

Given a price vector $p$, the minimum value of $\delta$ for which there exists a $\delta$-optimal triple with prices $p$ can be obtained by computing a maximum mean cut, from which we can also obtain the corresponding flow $\xi$ and base $y$ 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 that use maximum mean cuts [6, 23]. We divide the $\delta$-scaling phases of RCS-MCSF into blocks, so that each block contains $\mathrm{O}(\log n)$ phases. We start each block by computing a $\delta$-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 $\delta$ 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 $\left.p^{*}\right)$. Since there are only $m^{\prime}:=m+n(n-1)=\mathrm{O}\left(n^{2}\right)$ reduced costs (one for each arc of $I=A \cup E$ ), all reduced costs are fixed to zero after $2 m^{\prime}=\mathrm{O}\left(n^{2}\right)$ blocks. At this point, we can compute an optimal flow with one maximum submodular flow computation. Thus the algorithm requires only $\mathrm{O}\left(n^{2} \log n\right)$ phases overall. As in the weakly polynomial algorithm described in Section 5, the time per phase is strongly polynomial, $\mathrm{O}\left(n^{4} h\right)$.

### 6.1 Maximum Mean Cuts and Approximations

In this section, we define maximum mean cuts, explain their usefulness in obtaining $\delta$-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 $p$, for the flow arcs we define modified bounds $l_{p}$ and $u_{p}$ as follows: For $a \in A$, 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 $l_{p}(a)=u_{p}(a)=u(a)$. (For $a \in D, l_{p}(a)=l(a)$ and $u_{p}(a)=u(a)$.) This definition is consistent with Theorem 3.2 in that $\varphi$ and $p$ are primal and dual optimal if and only if $l_{p} \leq \varphi \leq u_{p}$ (equivalent to condition $\langle\mathrm{a}\rangle$ of the theorem) and $\partial \varphi \in \mathrm{B}\left(f_{p}\right)$. 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 $\Delta^{+}, \Delta^{-}$, and $\partial$ to indicate the relevant arc set.

A cut is a proper nonempty subset $X \in \mathcal{D}$ of $V$. The value of cut $X$ with respect to prices $p$ is defined as $\kappa_{p}(X):=l_{p}\left(\Delta_{I}^{+} X\right)-u_{p}\left(\Delta_{I}^{-} X\right)-f_{p}(X)$. If $\delta$ were zero, then Theorems 3.1 and 3.2 would imply that $p$ is dual optimal if and only if there is no cut having positive value, so positive cut value represents the amount of dual infeasibility. Note that as $\delta$ increases, $\kappa_{p}(X)$ decreases, making dual infeasibility less likely. We want the smallest $\delta$ such that no positive cuts remain; as is well-known, this ends up being the ratio between cut value and the amount of relaxation crossing the cut. Thus we define the mean value of a cut $X$ using a denominator of $b(X):=|X| \cdot|V-X|$, or as $\bar{\kappa}_{p}(X)=\kappa_{p}(X) / b(X)$. A maximum mean cut $T$ satisfies $\bar{\kappa}_{p}(T)=\max _{X \in \mathcal{D}} \bar{\kappa}_{p}(X)$. For fixed $p$, we denote this value by $\bar{\kappa}_{p}$. This definition of maximum mean cut differs from the definition in [23] because we relax capacities differently.

The algorithm contains $\mathrm{O}\left(n^{2}\right)$ blocks, and hence computes $\mathrm{O}\left(n^{2}\right)$ maximum mean cuts. A maximum mean cut can be computed by discrete Newton's algorithm [27, 29] using a maximum submodular flow subroutine for the data $l_{p}, u_{p}$, and $f_{p}$, with variable $\delta$. (Cunningham and Frank [2, Theorem 7] show how to simulate an exchange oracle for $f_{p}$ in $\mathrm{O}(1)$ calls to an exchange oracle for $f$.) This algorithm requires $\mathrm{O}\left(m^{\prime}\right)=\mathrm{O}\left(n^{2}\right)$ calls to submodular maximum flow. Each maximum submodular flow can be computed in $\mathrm{O}\left(n^{3} h\right)$ time using Fujishige and Zhang's algorithm [18, 24], for a total of $\mathrm{O}\left(n^{5} h\right)$ time per max mean cut, and $\mathrm{O}\left(n^{7} h\right)$ time for computing max mean cuts overall. This would dominate the time for the rest of the algorithm, which we will show is only $\mathrm{O}\left(n^{6} h \log n\right)$, so we use a trick of Radzik from [6] to remove this bottleneck: Just one call to submodular max flow (with $\delta$ fixed to zero) produces a cut $T^{\circ}$ maximizing $\kappa_{p}(T)$, a most positive cut. If $T^{*}$ is a max mean cut, then $\kappa_{p}\left(T^{\circ}\right) \geq \kappa_{p}\left(T^{*}\right)$, which implies that $\bar{\kappa}_{p}=\bar{\kappa}_{p}\left(T^{*}\right) \leq \bar{\kappa}_{p}\left(T^{\circ}\right) \frac{b\left(T^{\circ}\right)}{b\left(T^{*}\right)} \leq n \bar{\kappa}_{p}\left(T^{\circ}\right)$, so that $\bar{\kappa}_{p}\left(T^{\circ}\right) \geq \bar{\kappa}_{p} / n$. Thus $T^{\circ}$ 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 $\mathrm{O}\left(n^{7} h\right)$ to $\mathrm{O}\left(n^{5} h\right)$, 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 $\delta$, 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 $\delta$-optimal. Then there exists an optimal flow $\varphi^{*}$ to the original instance such that
$\langle\mathrm{a}\rangle\left|\varphi(a)-\varphi^{*}(a)\right| \leq 3 n^{2} \delta$ for all $a \in A$, and
〈b) $\left|\partial \varphi^{*}(X)-y(X)\right| \leq 3 n^{3} \delta$ for any cut $X$.
Proof. First we convert $\xi$ 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 $\Phi$ by at most $2 m \delta \leq 2 n(n-1) \delta$. Then we delete all the $D$ arcs and their flows $\psi$. Since such arcs have flow at most $\delta$, this increases $\Phi$ by at most $n(n-1) \delta$. Since we started out with $\Phi \leq 2 n \delta$, we now have $\Phi \leq\left(3 n^{2}-n\right) \delta$. Call the resulting flow $\varphi^{\prime}$. Note that $\left|\varphi^{\prime}(a)-\varphi(a)\right| \leq \delta$ for all $a \in A$. For any set $X$, saturating negative reduced cost $\operatorname{arcs}$ in $A$ and removing flow $\psi$ from $D$ implies that $\left|\partial_{I} \xi(X)-\partial_{A} \varphi^{\prime}(X)\right| \leq \frac{n^{2}}{2} \delta$. Since we had $|\partial \varphi(X)-y(X)+\partial \psi(X)| \leq \Phi \leq 2 n \delta$, this implies that $\left|\partial_{A} \varphi^{\prime}(X)-y(X)\right| \leq n^{2} \delta$ for all $X \subset V$.

Now use SSP to convert $\varphi^{\prime}$ into an optimal flow $\varphi^{*}$. (As originally stated, SSP is finite only for integer data, but the lexicographic techniques of Schönsleben [30] and Lawler and Martel [25] show that it can be made finite for any data.) Since $\Phi \leq\left(3 n^{2}-n\right) \delta$, the total amount of flow pushed by SSP is at most $\left(3 n^{2}-n\right) \delta$. This implies that for each $a \in A,\left|\varphi^{\prime}(a)-\varphi^{*}(a)\right| \leq\left(3 n^{2}-n\right) \delta$. Since $\left|\varphi^{\prime}(a)-\varphi(a)\right| \leq \delta$ we get $\left|\varphi(a)-\varphi^{*}(a)\right| \leq 3 n^{2} \delta$, verifying $\langle\mathrm{a}\rangle$.

Suppose that $P$ is an augmenting path chosen by the algorithm, and that flow is augmented by amount $\tau_{P}$ along $P$. By transitivity (Lemma 2.4 ), $P$ contains at most $n / 2 \operatorname{arcs}$ in $E_{y}$. Thus the flow boundary of any $X \subset V$ changes by at most $\frac{n}{2} \tau_{P}$ due to $P$. Since $\sum_{P} \tau_{P} \leq\left(3 n^{2}-n\right) \delta$, we have that $\left|\partial_{A} \varphi^{\prime}(X)-\partial_{A} \varphi^{*}(X)\right| \leq \frac{3 n^{3}}{2} \delta$. Since $\left|\partial_{A} \varphi^{\prime}(X)-y(X)\right| \leq n^{2} \delta$, this implies that $\left|y(X)-\partial_{A} \varphi^{*}(X)\right| \leq 3 n^{3} \delta$, verifying $\langle\mathrm{b}\rangle$.

Corollary 6.2 Suppose that $(\xi, y, p)$ is $\delta$-optimal. Then for all optimal prices $p^{*}$,
$\langle\mathrm{a}\rangle$ if there is an arc $a \in A$ with $\varphi(a)<u(a)-3 n^{2} \delta$, then $c_{p^{*}}(a) \geq 0$;
$\langle\mathrm{a} 2\rangle$ if there is an arc $a \in A$ with $\varphi(a)>l(a)+3 n^{2} \delta$, then $c_{p^{*}}(a) \leq 0$;
$\langle\mathrm{b}\rangle$ if there is a pair $(w, v)$ with $\sigma(y, v, w)>3 n^{3} \delta$, then $p^{*}(v) \leq p^{*}(w)$.
Proof. For $\langle\mathrm{a} 1\rangle, \varphi(a)<u(a)-3 n^{2} \delta$ and Lemma $6.1\langle\mathrm{a}\rangle$ imply that $\varphi^{*}(a)<u(a)$. Then complementary slackness implies that $c_{p^{*}}(a) \geq 0$ for all optimal prices $p^{*}$. A symmetric argument implies $\langle\mathrm{a} 2\rangle$.

For $\langle\mathrm{b}\rangle$, suppose that $X^{*}$ determines $\sigma\left(\partial \varphi^{*}, v, w\right)$, so that $\sigma\left(\partial \varphi^{*}, v, w\right)=f\left(X^{*}\right)-\partial \varphi^{*}\left(X^{*}\right)$. Then $\sigma(y, v, w)>3 n^{3} \delta$ implies that $f\left(X^{*}\right)-y\left(X^{*}\right)>3 n^{3} \delta$, and Lemma $6.1\langle\mathrm{~b}\rangle$ implies that $\sigma\left(\partial \varphi^{*}, v, w\right)=f\left(X^{*}\right)-\partial \varphi^{*}\left(X^{*}\right)>0$. Then Theorem 3.2 implies $\langle\mathrm{b}\rangle$.

Let $p^{\circ}$ be the prices at the beginning of a block. We first compute a most positive cut $T^{\circ}$. If $\kappa_{p}\left(T^{\circ}\right) \leq 0$, then we are optimal. Otherwise, $\bar{\kappa}_{p}\left(T^{\circ}\right):=\delta^{\circ}>0$. Since $\delta^{\circ}=\bar{\kappa}_{p^{\circ}}\left(T^{\circ}\right) \geq \bar{\kappa}_{p^{\circ}} / n$, we can compute a feasible submodular flow $\xi^{\circ}=\left(\varphi^{\circ}, \psi^{\circ}\right)$ in the network with data $l_{p^{\circ}}, u_{p^{\circ}}, f_{p^{\circ}}$, using $\delta=n \delta^{\circ}$. With $y^{\circ}=\partial \varphi^{\circ}$ and $x^{\circ}=y^{\circ}-\partial \psi^{\circ}$, we have that $\xi^{\circ}, y^{\circ}$, and $p^{\circ}$ are almost $n \delta^{\circ}$-optimal, except that the discrepancy $\Phi=\sum_{v}\left|x^{\circ}(v)-\partial \varphi^{\circ}(v)\right|=\sum_{v}\left|\psi^{\circ}(v)\right|$ could be as large as $n^{2}\left(n \delta^{\circ}\right)$. So we start out our block of $\mathrm{O}(\log n)$ phases with $\delta=n^{2} \delta^{\circ}$ and these values. Let $\xi^{\prime}, y^{\prime}$, and $p^{\prime}$ be the corresponding values that are $\delta^{\prime}$-optimal after the last phase in the block. We first need a lemma on the proximity of $f_{p^{\circ}}$ and $f$.

Theorem 6.4 below will show that at the end of each block, there is some $a \in \Delta_{A}^{+} T^{\circ}$ satisfying Corollary $6.2\langle\mathrm{a} 1\rangle$, or some $a \in \Delta_{A}^{-} T^{\circ}$ satisfying $\langle\mathrm{a} 2\rangle$, or some $a=(w, v)$ and a level set $L^{\circ}$ of $p^{\circ}$ with $v \in L^{\circ}, w \notin L^{\circ}$ satisfying $\langle\mathrm{b}\rangle$. We restrict the sign of $a^{\prime}$ 's reduced cost as follows: In case $\langle\mathrm{a} 1\rangle$, replace $u(a)$ by $+\infty$. In case $\langle\mathrm{a} 2\rangle$, replace $l(a)$ by $-\infty$. In case $\langle\mathrm{b}\rangle$, add a new arc $\widehat{a}=(w, v)$ with $l(\widehat{a})=0, u(\widehat{a})=\infty$, and $c(\widehat{a})=0$. Since an optimal submodular flow $\varphi^{*}$ for the original problem remains optimal (with $\varphi^{*}(\widehat{a})=0$ ), these changes do not affect the set of optimal prices. We first need a lemma on the proximity of $f_{p^{\circ}}$ and $f$.

Lemma 6.3 Suppose that we have scalar $\gamma>0, y^{\prime} \in \mathrm{B}(f)$ and prices $p^{\circ}$ such that for all level sets $L$ of $p^{\circ}$, and all $(w, v)$ with $v \in L$ and $w \notin L$ we have $\sigma\left(y^{\prime}, v, w\right) \leq \gamma$. Then for any $T^{\circ}$, $y^{\prime}\left(T^{\circ}\right)-f_{p^{\circ}}\left(T^{\circ}\right) \leq n^{3} \gamma$.

Proof. First, $\sigma\left(y^{\prime}, v, w\right) \leq \gamma$ implies that there exists a set $X_{w v}$ with $v \in X_{w v}, w \notin X_{w v}$, and $y^{\prime}\left(X_{w v}\right) \geq f\left(X_{v w}\right)-\gamma$. Since $L=\bigcup_{v \in L} \bigcap_{w \notin L} X_{w v}$, by submodularity of $f-y^{\prime}, y^{\prime} \in \mathrm{B}(f)$, and $f(V)=0$, we get $y^{\prime}(L) \geq f(L)-b(L) \gamma \geq f(L)-n^{2} \gamma$. Then modularity of $y^{\prime}$ and $y^{\prime} \in \mathrm{B}(f)$ imply

$$
\begin{aligned}
y^{\prime}\left(T^{\circ}\right) & =\sum_{i=1}^{k}\left[y^{\prime}\left(\left(T^{\circ} \cap L_{i}\right) \cup L_{i-1}\right)-y^{\prime}\left(L_{i-1}\right)\right] \\
& \leq \sum_{i=1}^{k}\left[f\left(\left(T^{\circ} \cap L_{i}\right) \cup L_{i-1}\right)-\left(f\left(L_{i-1}\right)-n^{2} \gamma\right)\right] \leq f_{p^{\circ}}\left(T^{\circ}\right)+n^{3} \gamma
\end{aligned}
$$

Theorem 6.4 After a block of $\left\lceil\log _{2}\left(5 n^{8}\right)\right\rceil=\mathrm{O}(\log n)$ scaling phases, at least one of the following holds:
(a1) $\exists a \in \Delta_{A}^{+} T^{\circ}$ with $l_{p^{\circ}}(a)-\varphi^{\prime}(a)>3 n^{2} \delta^{\prime}$,
(a2) $\exists a \in \Delta_{A}^{-} T^{\circ}$ with $\varphi^{\prime}(a)-u_{p^{\circ}}(a)>3 n^{2} \delta^{\prime}$,
(b) $\exists(w, v)$ with $\sigma\left(y^{\prime}, v, w\right)>3 n^{3} \delta^{\prime}$ and with $v \in L^{\circ}, w \notin L^{\circ}$ for some level set $L^{\circ}$ of $p^{\circ}$.

The reduced cost sign restriction resulting from applying Corollary 6.2 is new.
Proof. Since each scaling phase cuts $\delta$ in half,

$$
\begin{equation*}
\delta^{\prime} \leq n^{2} \delta^{\circ} / 5 n^{8}=\delta^{\circ} / 5 n^{6} \tag{4}
\end{equation*}
$$

Conservation of flow for $\xi^{\prime}$ implies that $\xi^{\prime}\left(\Delta_{I}^{+} T^{\circ}\right)-\xi^{\prime}\left(\Delta_{I}^{-} T^{\circ}\right)-\partial_{I} \xi^{\prime}\left(T^{\circ}\right)=0$. Subtracting this from $\kappa_{p^{\circ}}\left(T^{\circ}\right)$ yields

$$
\begin{equation*}
\delta^{\circ} b\left(T^{\circ}\right)=\kappa_{p^{\circ}}\left(T^{\circ}\right)=\left(l_{p^{\circ}}-\xi^{\prime}\right)\left(\Delta_{I}^{+} T^{\circ}\right)+\left(\xi^{\prime}-u_{p^{\circ}}\right)\left(\Delta_{I}^{-} T^{\circ}\right)+\left(\partial_{I} \xi^{\prime}\left(T^{\circ}\right)-f_{p^{\circ}}\left(T^{\circ}\right)\right) \tag{5}
\end{equation*}
$$

If none of (a1), (a2), or (b) holds, then from (5), the bound on $\Phi$, and Lemma 6.3 with $\gamma=3 n^{3} \delta^{\prime}$ (noting that for $a \in D,\left|l_{p^{\circ}}(a)-\psi^{\prime}(a)\right| \leq \delta^{\prime}$ and $\left|\psi^{\prime}(a)-u_{p^{\circ}}(a)\right| \leq \delta^{\prime}$ ), we get $\delta^{\circ}<\delta^{\circ} b\left(T^{\circ}\right)=$ $\kappa_{p^{\circ}}\left(T^{\circ}\right) \leq 3 n^{2}\left|\Delta_{A}\left(T^{\circ}\right)\right| \delta^{\prime}+n^{2} \delta^{\prime}+\partial_{I} \xi^{\prime}\left(T^{\circ}\right)-y^{\prime}\left(T^{\circ}\right)+y^{\prime}\left(T^{\circ}\right)-f_{p^{\circ}}\left(T^{\circ}\right) \leq 4 n^{4} \delta^{\prime}+2 n \delta^{\prime}+3 n^{6} \delta^{\prime} \leq 5 n^{6} \delta^{\prime}$, contradicting (4).

In case (a1), we must have that $l_{p^{\circ}}(a)=u(a)$, implying that $c_{p^{\circ}}(a)<0$. Thus $\varphi^{\prime}(a)<$ $u(a)-3 n^{2} \delta^{\prime}$, and from Corollary $6.2\langle\mathrm{a} 1\rangle$ 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 in case (a2), $c_{p^{*}}(a) \leq 0$ is a new sign restriction on $c_{p}$.

Finally, in case (b), from Corollary $6.2\langle\mathrm{~b}\rangle$ we have $p^{*}(v) \leq p^{*}(w)$ for all optimal $p^{*}$. Since $L^{\circ}$ is a level set of $p^{\circ}$, we had $p^{\circ}(v)>p^{\circ}(w)$, so this is a new sign restriction on $p$.

Theorem 6.5 The algorithm described in this section computes a minimum cost submodular flow in $\mathrm{O}\left(n^{6} h \log n\right)$ time.

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

The strongly polynomial algorithm can be modified to work for crossing submodular functions: In the definition of the cut value $\kappa_{p}(X)$, the last term $f_{p}(X)$ should be replaced by $\widetilde{f}_{p}(X)$, where $\tilde{\widetilde{f}}$ is the fully submodular function with $\mathrm{B}(f)=\mathrm{B}(\widetilde{f})$. An implementation of Fujishige and Zhang's algorithm [24] finds a most positive cut for the new definition of $\kappa_{p}$. The algorithm does not need the function values of $\tilde{f}$, which is used only for its analysis.

## Acknowledgments

We would like to thank the IPCO VII Organizing Committee for helping to arrange the opportunity for us to work together in Graz, Maiko Shigeno for many valuable conversations and discussions on related topics, and a referee for prompting us to consider the crossing submodular function case.

## References

[1] G. Birkhoff. Lattice Theory. Amer. Math. Soc., 1967.
[2] W. H. Cunningham and A. Frank. A primal-dual algorithm for submodular flows. Math. Oper. Res., 10 (1985), 251-262, .
[3] 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.
[4] J. Edmonds and R. Giles. A min-max relation for submodular functions on graphs. Ann. Discrete Math., 1 (1977), 185-204, .
[5] J. Edmonds and R. M. Karp. Theoretical improvements in algorithmic efficiency for network flow problems. J. ACM, 19 (1972), 248-264, .
[6] 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.
[7] L. Fleischer and S. Iwata. Improved algorithms for submodular function minimization and submodular flow. Proceedings of the 32th Annual ACM Symposium on Theory of Computing, 107-116.
[8] A. Frank. Finding feasible vectors of Edmonds-Giles polyhedra. J. Combin. Theory, 36 (1984), 221-239, .
[9] A. Frank and É. Tardos. An application of simultaneous Diophantine approximation in combinatorial optimization. Combinatorica, 7 (1987), 49-65.
[10] A. Frank and É. Tardos. Generalized polymatroids and submodular flows. Math. Programming, 42 (1988), 489-563.
[11] S. Fujishige. Algorithms for solving the independent-flow problems, J. Oper. Res. Soc. Japan, 21 (1978), 189-204.
[12] S. Fujishige. Structures of polyhedra determined by submodular functions on crossing families. Math. Programming, 29 (1984), 125-141.
[13] 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.
[14] S. Fujishige. Submodular Functions and Optimization. North-Holland, 1991.
[15] S. Fujishige and S. Iwata: Algorithms for submodular flows, IEICE Trans. Inform. Syst., E83-D (2000), 322-329.
[16] S. Fujishige, H. Röck, and U. Zimmermann. A strongly polynomial algorithm for minimum cost submodular flow problems. Math. Oper. Res., 14 (1989), 60-69.
[17] S. Fujishige and N. Tomizawa. A note on submodular functions on distributive lattices. $J$. Oper. Res. Soc. Japan, 26 (1983), 309-318.
[18] S. Fujishige and X. Zhang. New algorithms for the intersection problem of submodular systems. Japan J. Indust. Appl. Math., 9 (1992), 369-382.
[19] S. Iwata. A capacity scaling algorithm for convex cost submodular flows. Math. Programming, 76 (1997), 299-308.
[20] S. Iwata. A fully combinatorial algorithm for submodular function minimization. METR 00-09, University of Tokyo, October 2000. To appear in J. Combin. Theory Ser. B.
[21] S. Iwata, L. Fleischer, and S. Fujishige. A combinatorial, strongly polynomial-time algorithm for minimizing submodular functions. Proceedings of the 32th Annual ACM Symposium on Theory of Computing (2000), 97-106. To appear in J. ACM.
[22] S. Iwata, S. T. McCormick, and M. Shigeno. A faster algorithm for minimum cost submodular flows. Proceedings of the Ninth Annual ACM-SIAM Symposium on Discrete Algorithms (1998), 167-174. To appear in Combinatorica.
[23] 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.
[24] S. Iwata, S. T. McCormick, and M. Shigeno. A fast cost scaling algorithm for submodular flow. Inform. Process. Lett., 74 (2000), 123-128.
[25] E. L. Lawler and C. U. Martel. Computing maximal polymatroidal network flows. Math. Oper. Res., 7 (1982), 334-347.
[26] C. L. Lucchesi and D. H. Younger. A minimax relation for directed graphs. J. London Math. Soc., 17 (1978), 369-374.
[27] S. T. McCormick and T. R. Ervolina. Computing maximum mean cuts. Discrete Appl. Math., 52 (1994), 53-70.
[28] T. Naitoh and S. Fujishige. A note on the Frank-Tardos bi-truncation algorithm for crossingsubmodular functions. Math. Programming, 53 (1992), 361-363.
[29] 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.
[30] P. Schönsleben. Ganzzahlige Polymatroid-Intersektions Algorithmen. Dissertation, ETH Zürich, 1980.
[31] A. Schrijver. A combinatorial algorithm minimizing submodular functions in strongly polynomial time. J. Combin. Theory Ser. B 80 (2000), no. 2, 346-355.
[32] 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., 25 (2000), 76-104.
[33] É. Tardos. A strongly polynomial minimum cost circulation algorithm. Combinatorica, 5 (1985), 247-256.
[34] É. Tardos, C. A. Tovey, and M. A. Trick, Layered augmenting path algorithms. Math. Oper. Res., 11 (1986), 362-370.


[^0]:    ${ }^{*}$ Graduate School of Industrial Administration, Carnegie Mellon University, Pittsburgh, PA 15213, USA. Email: lkf@andrew.cmu.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. Additional support from NSF through grant EIA-0049084.
    ${ }^{\dagger}$ Department of Mathematical Engineering and Information Physics, University of Tokyo, Tokyo 113-8656, Japan. Email: iwata@sr3.t.u-tokyo.ac.jp. Partially supported by Grant-in-Aid for Scientific Research No. 10750053 from Ministry of Education, Science, Sports, and Culture of Japan.
    ${ }^{\ddagger}$ 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 visits to SORIE at Cornell University and to LIMOS at Université Blaise Pascal, Clermont-Ferrand.

