

Concave Generalized Flows with Applications to Market Equilibria

László A. Végh*

Department of Management, London School of Economics.
L.Vegh@lse.ac.uk

July 24, 2012

Abstract

We consider a nonlinear extension of the generalized network flow model, with the flow leaving an arc being an increasing concave function of the flow entering it, as proposed by Truemper [40] and Shigeno [36]. We give a polynomial time combinatorial algorithm for solving corresponding flow maximization problems, finding an ε -approximate solution in $O(m(m + \log n) \log(MUm/\varepsilon))$ arithmetic operations and value oracle queries, where M and U are upper bounds on simple parameters. This also gives a new algorithm for linear generalized flows, an efficient, purely scaling variant of the Fat-Path algorithm by Goldberg, Plotkin and Tardos [13], not using any cycle cancellations.

We show that this general convex programming model serves as a common framework for several market equilibrium problems, including the linear Fisher market model and its various extensions. Our result immediately provides combinatorial algorithms for various extensions of these market models. This includes nonsymmetric Arrow-Debreu Nash bargaining, settling an open question by Vazirani [42].

1 Introduction

A classical extension of network flows is the *generalized network flow model*, with a gain factor $\gamma_e > 0$ associated with each arc e so that if α units of flow enter arc e , then $\gamma_e \alpha$ units leave it. Since first studied by Kantorovich [23], Dantzig [5] and Jewell [21], the problem has found many applications including financial analysis, transportation, management sciences, see [2, Chapter 15].

In this paper, we consider a nonlinear extension, *concave generalized flows*, studied by Truemper [40] in 1978, and by Shigeno [36] in 2006. For each arc e we are given a concave, monotone increasing function Γ_e such that if α units enter e then $\Gamma_e(\alpha)$ units leave it. We give a combinatorial algorithm for corresponding flow maximization problems, with running time polynomial in the network data and some simple parameters. We also exhibit new applications, showing that it is a general framework containing multiple convex programs for market equilibrium settings, for which combinatorial algorithms have been developed over the last decade. As an application, we also get a combinatorial algorithm for nonsymmetric Arrow-Debreu Nash bargaining, resolving an open question by Vazirani [42]. We can also extend existing results to more general settings.

Generalized flows are linear programs and thus can be solved efficiently by general linear programming techniques, the currently most efficient such algorithm being the interior-point method

*This work was done in the College of Computing, Georgia Institute of Technology, supported by NSF Grant CCF-0914732.

by Kapoor and Vaidya [24]. Combinatorial approaches have been used since the sixties (e.g. [21, 28, 39]), yet the first polynomial-time combinatorial algorithms were given only in 1991 by Goldberg, Plotkin and Tardos [13]. This inspired a line of research to develop further polynomial-time combinatorial algorithms, e.g. [4, 15, 17, 38, 10, 16, 18, 44, 33, 34]; for a survey on combinatorial generalized flow algorithms, see [35]. Despite the vast literature, no strongly polynomial algorithm is known so far. Our algorithm for this special case derives from the FAT-PATH algorithm in [13], with the remarkable difference that no cycle cancellations are needed.

Nonlinear extensions of generalized flows have also been studied, e.g. in [1, 3], minimizing a separable convex cost function for generalized flows. However, these frameworks do not contain our problem, which involves nonlinear convex constraints.

Concave generalized flows being nonlinear convex programs, they can also be solved by the ellipsoid method, yet no practically efficient methods are known for this problem. Hence finding a combinatorial algorithm is also a matter of running time efficiency. Shigeno [36] gave the first combinatorial algorithm that runs in polynomial time for some restricted classes of functions Γ_e , including piecewise linear. It is also an extension of the FAT-PATH algorithm in [13]. In spite of this development, it has remained an open problem to find a combinatorial polynomial-time algorithm for arbitrary concave increasing gain functions.

Our result settles this question by allowing arbitrary increasing concave gain functions provided via value oracle access. The running time bounds for this general problem are reasonably close to the most efficient linear generalized flow algorithms. Concave gain functions extend the applicability range of the classical generalized flow model, as they can describe e.g. diminishing marginal utilities. From the application point of view, another contribution of the paper is extending generalized flow techniques to the domain of market equilibrium computations, where this model turns out to be a concise unifying framework.

The concave optimization problem might have irrational optimal solutions: in general, we give a fully polynomial-time approximation scheme, with running time dependent on $\log(\frac{1}{\epsilon})$ for finding an ϵ -approximate solution. In the market equilibrium applications we have rational convex programs (as in [42]): the existence of a rational optimal solution is guaranteed. We show a general technique to transform a sufficiently good approximation delivered by our algorithm to an exact optimal solution under certain circumstances. We demonstrate how this technique can be applied on the example of nonsymmetric Arrow-Debreu Nash bargaining, where the existence of a combinatorial algorithm was open [42].

In Section 2, we give the precise definition of the problems considered. Thereby we introduce a new, equivalent variant of the problem, called the *symmetric formulation*, providing a more flexible algorithmic framework. Section 3 shows the applications for market equilibrium problems. Section 4 explores the background of minimum-cost circulation and generalized flow algorithms, and exhibits the main algorithmic ideas. We first present our symmetric generalized flow algorithm in Section 5 for the special case of linear gains. Based on this, Section 6 gives the algorithm for arbitrary concave gain functions. Section 7 adapts these algorithms for the more standard sink formulation of the problems. Section 8 considers the case when the existence of a rational optimal solution is guaranteed, and shows how the approximate solution provided by our algorithm can be turned to an optimal solution. The final Section 9 discusses possible further directions.

2 Problem definitions

We define two closely related variants of the linear and the concave generalized flow problems. Let $G = (V, E)$ be a directed graph. Let $n = |V|$, $m = |E|$, and for each node $i \in V$, let d_i be the total

number of incoming and outgoing arcs incident to i .

In the linear setting, we are given lower and upper arc capacities $\ell, u : E \rightarrow \mathbb{R}$ and gain factors $\gamma : E \rightarrow \mathbb{R}^+$ on the arcs, and node demands $b : V \rightarrow \mathbb{R}$. By a *pseudoflow* we mean a function $f : E \rightarrow \mathbb{R}$ with $\ell \leq f \leq u$. Given the pseudoflow f , let

$$e_i := \sum_{j:ji \in E} \gamma_{ji} f_{ji} - \sum_{j:ij \in E} f_{ij} - b_i. \quad (1)$$

In the first variant of the problem, called the *sink formulation*, there is a distinguished *sink* node $t \in V$. The objective is to maximize e_t for pseudoflows satisfying $e_i \geq 0$ for all $i \in V - t$.

This differs from the way the problem is usually defined in the literature with the more restrictive $e_i = 0$ for $i \in V - t$, and assuming $\ell \equiv 0$, $b \equiv 0$. However, this problem can easily be reduced to solving the sink formulation, see e.g. [35].

The convex extension has been proposed by Truemper [40] and Shigeno [36]. On each arc $ij \in E$, we are given lower and upper arc capacities $\ell, u : E \rightarrow \mathbb{R}$ and a monotone increasing concave function $\Gamma_{ij} : [\ell_{ij}, u_{ij}] \rightarrow \mathbb{R} \cup \{-\infty\}$; we are also given node demands $b : V \rightarrow \mathbb{R}$. As for generalized flows, a pseudoflow is a function $f : E \rightarrow \mathbb{R}$ with $\ell \leq f \leq u$. For a pseudoflow f , let

$$e_i := \sum_{j:ji \in E} \Gamma_{ji}(f_{ji}) - \sum_{j:ij \in E} f_{ij} - b_i.$$

In the *concave sink formulation*, we say that the pseudoflow f is *feasible*, if $e_i \geq 0$ for all $i \in V - t$ and $e_t > -\infty$. The objective is to maximize e_t for feasible pseudoflows.

Shigeno [36] defines this problem with $e_i = 0$ if $i \in V - t$, and $b \equiv 0$ and without explicit capacity constraints. She also discusses the version with $e_i \geq 0$, and gives a reduction from the original version to this one. Whereas capacity constraints can be simulated by the functions Γ_e , we impose them explicitly as they will be included in the running time bounds. The formulation with $e_i \geq 0$ seems more natural as it gives a convex optimization problem, which is not the case for $e_i = 0$.

In the sink formulation, the node t plays a distinguished role. It turns out to be more convenient to handle all nodes equally. For this reason, we introduce another, seemingly more general version, called the *symmetric formulation* of both problems. Ideally, we would like to find a pseudoflow satisfying $e_i \geq 0$ for every $i \in V$. The formulation will be a relaxation of this feasibility problem, allowing violation of the constraints, penalized by possibly different rates at different nodes.

For each node $i \in V$ we are given a penalty factor $M_i > 0$ and an auxiliary variable $\kappa_i \geq 0$. The objective is to minimize $\kappa_f = \sum_{i \in V} M_i \kappa_i$ for a pseudoflow f subject to $e_i + \kappa_i \geq 0$ for each $i \in V$.

The objective κ_f is called the *excess discrepancy*. $\kappa_f = 0$ means $e_i \geq 0$ for each $i \in V$. These conditions might be violated, but we have to pay penalty M_i per unit violation at i .

The sink version fits into this framework with $M_i = \infty$ for $i \neq t$ and $M_t = 1$. However, it can be shown that setting finite, polynomially bounded M_i values, the symmetric version returns an optimal (or sufficiently close approximate) solution to the sink version, both for linear and for concave gain functions. Besides the sink version, another natural setting is when $M_i = 1$ for all $i \in V$, that is, maintaining $e_i \geq 0$ has the same importance for all nodes.

While the symmetric formulation could seem more general than the sink version, it can indeed be reduced to it. For an instance of the symmetric version with graph $G = (V, E)$, let us add a new sink node t with an arc from t to every node $i \in V$ with gain factor $1/M_i$. Solving the sink version for this extended instance gives an optimal solution to the original problem. The reason for introducing the symmetric formulation is its pertinence to our algorithmic purposes.

2.1 Complexity model

The complexity setting will be different for generalized flows and for concave generalized flows. For generalized flows, we aim to find an optimal solution, while in the concave case, only an approximate one. For generalized flows, the gain functions are given explicitly as linear functions, while in the concave case, the description of the functions might be infinite. To handle this difficulty, following the approach of Hochbaum and Shantikumar [19], we assume oracle access to the Γ_{ij} 's: our running time estimation will give a bound on the number of necessary oracle calls. Two kinds of oracles are needed: (i) value oracle, returning $\Gamma_{ij}(\alpha)$ for any $\alpha \in [\ell_{ij}, u_{ij}]$; and (ii) inverse value oracle, returning a value β with $\alpha = \Gamma_{ij}(\beta)$ for any $\alpha \in [\Gamma_{ij}(\ell_{ij}), \Gamma_{ij}(u_{ij})]$.

We assume that both oracles return the exact (possibly irrational) solution, and any oracle query is done in $O(1)$ time. Also, we assume any basic arithmetic operation is performed in $O(1)$ time, regardless to size and representation of the possibly irrational numbers. We expect that our results naturally extend to the setting with only approximate oracles and computational capacities in a straightforward manner. Notice that in an approximate sense, an inverse value oracle can be simulated by a value oracle.

By an ε -approximate solution to the symmetric concave generalized flow problem we mean a feasible solution with the excess discrepancy larger than the optimum by at most ε . An ε -approximate solution to the sink version means a pseudoflow with the objective value e_t at most ε less than the optimum, and the total violation of the inequalities $e_i \geq 0$ for $i \in V - t$ is also at most ε . (Note that an ε -approximate solution is thus not necessarily feasible.)

In both cases, we assume that all M_i values are positive integers, and let M denote their maximum.

For generalized flows, we assume all ℓ, u and b are given as integers and γ as rational numbers; let B be the largest integer used in their descriptions. The running time bound will be $O(m^2(m \log B + \log M) \log n)$ for the symmetric formulation and $O(m^2(m + n \log n) \log B)$ for the sink formulation. This is the same as the complexity bound of the highest gain augmenting path algorithm by Goldfarb, Jin and Orlin [17]. The best current running time bounds are $O(m^{1.5}n^2 \log B)$ using an interior point approach by Kapoor and Vaidya [24], and $\tilde{O}(m^2n \log B)$ by Radzik [33], that is an enhanced version of [17].

For the concave setting, we allow irrational capacities as well; in the complexity estimation, we will have U as an upper bound on the absolute values on the b_i 's, the capacities ℓ_{ij}, u_{ij} and the $\Gamma_{ij}(\ell_{ij}), \Gamma_{ij}(u_{ij})$ values. For each arc ij , let us define $r_{ij} = |\Gamma_{ij}(\ell_{ij})|$ whenever $\Gamma_{ij}(\ell_{ij}) > -\infty$ and $r_{ij} = 0$ otherwise. Let

$$U = \max \{ \max\{|b_i| : i \in V\}, \max\{|\ell_{ij}|, |u_{ij}|, |\Gamma_{ij}(u_{ij})|, r_{ij} : ij \in E\} \}.$$

For the sink version, we need to introduce one further complexity parameter U^* due to difficulties arising if $\Gamma_{ij}(\ell_{ij}) = -\infty$ for certain arcs. Let U^* satisfy $U \leq U^*$, and that $e_t \leq U^*$ for any pseudoflow (it is easy to see that $U^* = d_t U$ always satisfies this property). We also require that whenever there exists a feasible solution to the problem (that is, $e_i \geq 0$ for each $i \in V - t$ and $e_t > -\infty$), there exists one with $e_t \geq -U^*$. If $\Gamma_{jt}(\ell_{jt}) > -\infty$ for each arc $jt \in E$, then $U^* = d_t U$ also satisfies this property. However, we allow $-\infty$ values as in the market applications we also have logarithmic gain functions. A bound on U^* can be given as in Section 8.

The main result is as follows:

Theorem 2.1. *For the symmetric formulation of the concave generalized flow problem, there exists a combinatorial algorithm that finds an ε -approximate solution in $O(m(m + n \log n) \log(MUm/\varepsilon))$.*

For the sink formulation, there exists a combinatorial algorithm that finds an ε -approximate solution in $O(m(m + n \log n) \log(U^*m/\varepsilon))$. In both cases, the running time bound is on the number of arithmetic operations and oracle queries.

The starting point of our investigation is the FAT-PATH algorithm [13]. The first important idea is using the symmetric formulation. This is a more flexible framework, and thus we will be able to entirely avoid cycle cancellation and use excess transportation phases only. Our (linear) generalized flow algorithm is the first generalized flow algorithm that uses a pure scaling technique, without any cycle cancellation. The key new idea here is the way ‘ Δ -positive’ and ‘ Δ -negative’ nodes are defined, maintaining a ‘security reserve’ in each node that compensates for adjustments when moving from the Δ -scaling phase to the $\Delta/2$ -phase.

We extend the linear algorithm to the concave setting using a local linear approximation of the gain functions, following Shigeno [36]. This approximation is motivated by the technique of Minoux [26] and Hochbaum and Shantikumar [19] for minimum cost flows with separable convex objectives.

3 Applications to market equilibrium and Nash-bargaining problems

Intensive research has been pursued over the last decade to develop polynomial-time combinatorial algorithms for certain market equilibrium problems. The starting point is the algorithm for computing market clearing prices in Fisher’s model with linear utilities by Devanur et al. [7], followed by a study of several variations and extensions of this model. For a survey, see [27, Chapter 5] or [42].

In the *linear Fisher market model*, we are given a set B of buyers and a set G of goods. Buyer i has a budget m_i , and there is one divisible unit of each good to be sold. For each buyer $i \in B$ and good $j \in G$, $U_{ij} \geq 0$ is the utility accrued by buyer i for one unit of good j . Let $n = |B| + |G|$ and m be the number of pairs ij with $U_{ij} > 0$. We assume there is such an edge incident to every buyer and to every good. Let $U_{\max} = \max\{U_{ij} : i \in B, j \in G\}$ and $R = \max\{m_i : i \in B\}$. An equilibrium solution consist of prices p_i on the goods and an allocation x_{ij} , so that (i) all goods are sold, (ii) all money of the buyers is spent, and (iii) each buyers i buys a best bundle of goods, that is, goods j maximizing U_{ij}/p_j .

The equilibrium solutions for linear Fisher markets were described via a convex program by Eisenberg and Gale [9] in 1959; the combinatorial algorithms for this problem and other models rely on the KKT-conditions for the corresponding convex programs. Exact optimal solutions can be found, since these problems admit rational optimal solutions.

$$\begin{aligned}
 & \max \sum_{i \in B} m_i \log z_i \\
 & z_i \leq \sum_{j \in G} U_{ij} x_{ij} \quad \forall i \in B \\
 & \sum_{i \in B} x_{ij} \leq 1 \quad \forall j \in G \\
 & z, x \geq 0
 \end{aligned} \tag{EG}$$

We show that the Eisenberg-Gale convex program, along with all extensions studied so far, falls into the broader class of concave generalized flows. Moreover, in all these extension we may replace

linear or piecewise linear concave functions by arbitrary concave ones, still solvable approximately by our algorithm.

For the Eisenberg-Gale program, let us define the graph (V, E) with $V = B \cup G \cup \{t\}$. Let $ji \in E$ whenever $j \in G$, $i \in B$, $U_{ij} > 0$, and set $\Gamma_{ji}(\alpha) = U_{ij}\alpha$ as a linear gain function. Also, let $it \in E$ for every $i \in B$ with $\Gamma_{it}(\alpha) = m_i \log \alpha$. Finally, set $b_j = -1$ for $j \in G$, and $b_i = 0$ for $i \in B$. The above program describes exactly the sink version of this concave generalized flow instance with $f_{ji} = x_{ij}$ for $i \in B$, $j \in G$ and $f_{it} = z_i$. (To formally fit into the model, we may add upper capacities $u_{ji} = 1$ and $u_{it} = \sum_{j \in G} U_{ji}$ without changing the set of feasible solutions.) Hence our general algorithm gives an ε -approximation for this problem. In Section 8, we show that for sufficiently small ε we can transform it to an exact optimal solution.

The flexibility of the concave generalized flow model enables various extensions. For example, we can replace each linear function $U_{ij}\alpha$ by an arbitrary concave increasing function, obtaining the perfect price discrimination model of Goel and Vazirani [12]. They studied piecewise linear utility functions; our model enables arbitrary functions (although the optimal solution may be irrational).

In the *Arrow-Debreu Nash bargaining (ADNB)* defined by Vazirani [42], traders arrive to the market with initial endowments of goods, giving utility c_i for player i . They want to redistribute the goods to obtain higher utilities using Nash bargaining. The disagreement point is when everyone keeps the initial endowment, guaranteeing her $c_i \geq 0$ utility. In an optimal Nash bargaining solution we maximize $\sum_{i \in B} \log(z_i - c_i)$ over the constraint set in (EG). Unlike for the linear Fisher model, equilibrium prices may not exist, corresponding to a disagreement solution. A sophisticated two phase algorithm is given in [42], first for deciding feasibility, then for finding the equilibrium solution.

The convex program for nonsymmetric ADNB can be obtained from the Eisenberg-Gale program by modifying the first set of inequalities to $z_i \leq \sum_{j \in G} U_{ij}x_{ij} - c_i$. In the formulation as a concave generalized flow, this corresponds to modifying the $b_i = 0$ values for $i \in B$ to $b_i = c_i$. Hence this problem also fits into our framework. From this general perspective, it does not seem more difficult than the linear Fisher model.

Nonsymmetric Nash-bargaining was defined by Kalai [22]. For ADNB, it corresponds to maximizing $\sum_{i \in B} m_i \log(z_i - c_i)$ over the constraint set in (EG), for some positive coefficients m_i . The algorithm in [42] heavily relies on the assumption $m_i = 1$, and does not extend to this more general setting, called *nonsymmetric ADNB*. Finding a combinatorial algorithm for this latter problem was left open in [42]. Another open question in [42] is to devise a combinatorial algorithm for (nonsymmetric) ADNB with piecewise linear, concave utility functions. Our result generalizes even further, for arbitrary concave utility functions, since the linear functions $U_{ij}\alpha$ can be replaced by arbitrary concave functions.

Let $C = \max c_i$. In Section 8, we show how our algorithm can be used to find an exact solution to the nonsymmetric ADNB problem in time $O(m(m + n \log n)(n \log(nU_{\max}R) + \log C))$. The running time bound in [42] for symmetric ADNB ($R = 1$) is $O(n^8 \log U_{\max} + n^4 \log C)$.

Let us also remark that an alternative convex program for the linear Fisher market, given by Shmyrev [37], shows that it also fits into the framework of minimum-cost circulations with a separable convex cost function, and thus can be solved by the algorithms of Hochbaum and Shantikumar [19] or Karzanov and McCormick [25]. Recently, [43] gave a strongly polynomial algorithm for a class of these problems, which includes Fisher's market with linear and with spending constraint utilities. However, this does not seem to capture perfect price discrimination or ADNB, where no alternative formulations analogous to [37] are known.

As further applications of the concave generalized flow model, we can take single-source multiple-sink markets by Jain and Vazirani [20], or concave cost matchings studied by Devanur and Jain

[6].

A distinct characteristic of the Eisenberg-Gale program and its extensions is that they are rational convex programs. We may lose this property when changing to general concave spending constraint utilities. However, for the case when the existence of a rational solution is guaranteed, one would prefer finding an exact optimal solution. Section 8 addresses the question of rationality. Theorem 8.1 shows that under certain technical conditions, our approximation algorithm can be turned into a polynomial time algorithm for finding an exact optimal solution. We shall verify these conditions for nonsymmetric ADNB.

4 Background and overview

The minimum-cost circulation problem¹ is fundamental to all problems and algorithms discussed in the paper. We give an overview in Section 4.1. We present the two main algorithmic paradigms, cycle cancelling and successive shortest paths along with their efficient variants. As already revealed by early studies of the problem (e.g. [29, 39]), there is a deep connection between generalized flows and classical minimum-cost circulations: the dual structures are quite similar, and the generalized flow algorithms stem from the classical algorithms for minimum-cost circulations. In Section 4.2, we continue with an overview of generalized flow algorithms, exhibiting some important ideas and their relation to minimum-cost circulations. We also exhibit here the main ideas of our algorithm for the linear case. Section 4.3 considers a different convex extension of minimum-cost circulations, when the linear cost function is replaced by a separable convex one. We show how the two main paradigms extend to this case, using different approximation strategies of the nonlinear functions. Finally in Section 4.4 we consider the concave generalized flow problem, discuss the algorithm by Shigeno [36] and its relation to algorithmic ideas of the previous problems. We emphasize some difficulties and outline the ideas of our solution.

4.1 Minimum-cost flows: cycle cancelling and successive shortest paths

In the minimum-cost circulation problem, given is a directed graph $G = (V, E)$ with lower and upper arc capacities $\ell, u : E \rightarrow \mathbb{R} \cup \{\infty\}$, costs $c : E \rightarrow \mathbb{R}$ on the arcs and node demands $b : V \rightarrow \mathbb{R}$ with $\sum_{i \in V} b_i = 0$. Let

$$e_i = \sum_{j:ji \in E} f_{ji} - \sum_{j:ij \in E} f_{ij} - b_i.$$

$f : E \rightarrow \mathbb{R}$ with $\ell \leq f \leq u$ is called a feasible circulation, if $e_i = 0$ for all $i \in V$. The objective is to minimize $c^T f$ for feasible circulations.

Linear programming duality provides the following characterization of optimality. For a feasible circulation f , let us define the residual graph $G_f = (V, E_f)$ with $ij \in E_f$ if $ij \in E$ and $f_{ij} < u_{ij}$, or if $ji \in E$ and $\ell_{ji} < f_{ji}$. The first type of arcs are called *forward arcs* and are assigned the original cost c_{ij} , while the latter arcs are *backward arcs* assigned cost $-c_{ji}$. For notational convenience, we will use $f_{ij} = -f_{ji}$ on backward arcs. Then f is optimal if and only if E_f contains no negative cost cycles. This is further equivalent to the existence of a potential $\pi : V \rightarrow \mathbb{R}$ with $\pi_j - \pi_i \leq c_{ij}$ for all arcs $ij \in E_f$.

Two main frameworks for minimum-cost flow algorithms are as follows. In the *cycle cancelling framework* (see e.g. [2, Chapter 9.6]), we maintain a feasible circulation in each phase, with strictly increasing objective values. If the current solution is not optimal, the above conditions guarantee

¹We shall use the term ‘circulation’ to distinguish from other flow problems in the paper.

a negative cost cycle in the residual graph; such a cycle can be found efficiently. Sending some flow around this cycle decreases the objective and maintains feasibility, providing the next solution.

In the *successive shortest path framework* (see e.g. [2, Chapter 9.7]), we waive feasibility by allowing $e_i > 0$ or $e_i < 0$; we call such nodes *positive* and *negative*, respectively. However, we maintain dual optimality in the sense that the residual graph of the current pseudoflow contains no negative cost cycles in any iteration (or equivalently, admits a feasible potential). If there exists some positive and negative nodes, we send some flow from a positive node to a negative one using a minimum-cost path in the residual graph. This maintains dual optimality, and decreases the total e_i value of positive nodes.

For rational input data, both these algorithms are finite, but may take an exponential number of steps (and might not even terminate for irrational input data). Nevertheless, using (explicit or implicit) scaling techniques, both can be implemented to run in polynomial time, and even in strongly polynomial time.

A strongly polynomial version of the cycle cancellation algorithm is due to Goldberg and Tarjan [14]. In each step, a minimum mean cycle is chosen. In dual terms, we relax primal-dual optimality conditions to $\pi_j - \pi_i \leq c_{ij} + \varepsilon$ for $ij \in E_f$, with ε being equal to the negative of the minimum mean cycle value, decreasing exponentially over time.

Polynomial implementations of the successive shortest path algorithm can be obtained by capacity scaling; the most efficient, strongly polynomial such algorithm is due to Orlin [30]. We describe here a basic capacity scaling framework by Edmonds and Karp [8]. Instead of the residual graph E_f , we consider the Δ -residual graph E_f^Δ consisting of arcs with residual capacity at least Δ (the residual capacity is $u_{ij} - f_{ij}$ on a forward arc ij and $f_{ji} - \ell_{ji}$ on a backward arc). The algorithm consists of Δ -scaling phases, with Δ decreasing by a factor of 2 between two phases. In a Δ -phase, we iteratively send Δ units of flow from a positive node s with $e_s \geq \Delta$ to a negative node t with $e_t \leq -\Delta$ on a minimum-cost path in E_f^Δ . The Δ -phase finishes when this is no longer possible, which means the total positive excess is at most $n\Delta$.

In the Δ -phase, $\pi_j - \pi_i \leq c_{ij}$ is maintained on arcs of the Δ -residual graph. When moving to the $\Delta/2$ phase, this might not hold anymore, since the $\Delta/2$ -residual graph contains more arcs, namely the ones with residual capacity between $\Delta/2$ and Δ . At the beginning of the next phase, we saturate all these arcs, thereby increasing the positive excess to at most $(2n + m)\Delta/2$. This guarantees that the next phase will consist of at most $(2n + m)$ path augmentations.

4.2 Linear generalized flows – cycle cancelling and excess transportation

In what follows, we consider the sink version of the generalized flow problem, with sink $t \in V$. For a pseudoflow $f : E \rightarrow \mathbb{R}$, let us define the residual network $G_f = (V, E_f)$ as for circulations, with gain factor $\gamma_{ij} = 1/\gamma_{ji}$ on backward arcs. Consider a cycle C in E_f . We can modify f by sending some flow around C from some $i \in V(C)$. This leaves e_j unchanged if $j \neq i$, and increases e_i by $(\gamma(C) - 1)\alpha$, where $\gamma(C) = \prod_{e \in C} \gamma_e$. If $\gamma(C) > 1$ then we call C a *flow-generating cycle*, while for $\gamma(C) < 1$, a *flow-absorbing cycle*, since we can generate or eliminate excess at an arbitrary node $i \in C$, respectively. The amount of flow that can be generated is of course bounded by the capacity constraints.

To augment the excess of the sink t , we have to send the excess generated at a flow-generating cycle C to t . Hence we call a pair (C, P) a *generalized augmenting path (GAP)*, if (a) C is a flow-generating cycle, $i \in V(C)$, and P is a path in E_f from i to t ; or (b) $C = \emptyset$, and P is a path in E_f from some node i with $e_i > 0$ to t . Clearly, an optimal solution f may admit no GAPs. This is indeed an equivalence: f is optimal if and only if no GAP exists.

The gain factors γ_e play a role analogous to the costs c_e for minimum-cost circulations. Indeed,

C is a flow generating cycle if and only if it is a negative cost cycle for the cost function $c_e = -\log \gamma_e$. The dual structure for generalized flows is also analogous to potentials. Let us call $\mu : V \rightarrow \mathbb{R}_{>0} \cup \{\infty\}$ with $\mu_t = 1$ a label function. *Relabeling* the pseudoflow f by μ means dividing the flow on each arc ij going out from i by μ_i . We get a problem equivalent to the original by replacing each arc gain by $\gamma_{ij}^\mu = \gamma_{ij}\mu_i/\mu_j$. The labeling is called *conservative* if $\gamma_{ij}^\mu \leq 1$ for all $ij \in E_f$, that is, no arc may increase the relabeled flow.

Assume we have a conservative labeling μ so that $e_i = 0$ whenever $i \in V - t$, $\mu_i < \infty$. Let $V' \subseteq V$ denote the set of nodes from which there exists a directed path to t . It follows that (i) $\mu_i < \infty$ for all $i \in V'$, and (ii) V' contains no flow-generating cycles. Consequently, given a conservative labeling, no GAP can exist, and the converse can also be shown to hold. Note that on V' , $\pi_i = -\log \mu_i$ is a feasible potential for $c_e = -\log \gamma_e$ if and only if μ is conservative.

Based on this correspondence, minimum-cost circulation algorithms can be directly applied for generalized flows as a subroutine for eliminating all flow-generating cycles. This can be indeed implemented in strongly polynomial time, see [32, 35]. The novel difficulty for generalized flows is how to transport the generated excess from various nodes of the graph to the sink t . In the algorithm of Onaga [29], flow is transported iteratively on highest gain augmenting paths, that is, from $i \in V$ with $e_i > 0$ on an $i - t$ path P that maximizes $\gamma(P) = \prod_{e \in P} \gamma_e$. It can be shown that using such paths does not create any new flow generating cycles. Thus after having eliminated all type (a) GAPs, we only have to take care of type (b). Unfortunately, this algorithm may run in exponential time (or may not even terminate for irrational inputs). This is due to the analogy between Onaga's algorithm and the successive shortest path algorithm – observe that a highest gain path is a minimum-cost path for $-\log \gamma_e$.

The first algorithms to overcome this difficulty and thus establish polynomial running time bounds were the two given by Goldberg, Plotkin and Tardos [13]. One of them, FAT-PATH, uses a method analogous to capacity scaling. A path P in E_f from a node i to t is called Δ -fat, if assuming unlimited excess at i , it is possible to send enough flow along P from i to t so that e_t increases by Δ .

The algorithm consists of Δ -phases, with Δ decreasing by a factor of 2 for the next phase. In the Δ -phase, we first cancel all flow generating cycles. Then, from nodes i with $e_i > 0$, we transport flow on highest gain ones among the Δ -fat paths. This might create new flow-generating cycles to be cancelled in the next phase. Nevertheless, it can be shown that at the beginning of a Δ -phase, $e_t^* - e_t \leq 2(n+m)\Delta$ for the optimum value e_t^* and thus the number of path augmentations in each Δ -phase can be bounded by $2(n+m)$. Arriving at a sufficiently small value of Δ , it is possible to obtain an optimal solution by a single maximum flow computation.

The basic framework of [29] and of FAT-PATH, namely using different subroutines for eliminating flow-generating cycles and for transporting excess to the sink has been adopted by most subsequent algorithms, e.g. [15, 17, 38, 10, 33]. Among them, [17] is an almost purely scaling polynomial time algorithm, but it still needs the an initial cycle-cancelling as in [29].

In contrast, our algorithm does not need any cycle-cancelling, and adapts FAT-PATH to a pure successive shortest paths framework. The successive shortest paths algorithms for minimum-cost circulations start with an infeasible pseudoflow, having both positive and negative nodes. To use an analogous method for generalized flows, we have to give up the standard framework of algorithms where $e_i \geq 0$ is always maintained for all $i \in V - t$. This is the reason why we use the more flexible symmetric model: we start with possibly several nodes having $e_i < 0$, and our aim is to eliminate them. An important property of the algorithm is that we always have to maintain $\mu_i = 1/M_i$ for $e_i < 0$; for this reason we shall avoid creating new negative nodes.

Similarly to FAT-PATH, we use a scaling algorithm. In the Δ -phase, we consider the residual

graph restricted to Δ -fat arcs, arcs that may participate in a highest gain Δ -fat-path, and maintain a conservative labeling μ with $\gamma_{ij}^\mu \leq 1$ on the Δ -fat arcs. When moving to the $\Delta/2$ -phase, this condition may get violated due to $\Delta/2$ -fat arcs that were not Δ -fat. Analogously to the Edmonds-Karp algorithm, we modify the flow by saturating each violated arc and thereby reconstitute dual feasibility. However, these changes may create new negative nodes and thus violate the condition $\mu_i = 1/M_i$ for $e_i < 0$ we must maintain.

We resolve this difficulty by maintaining a ‘security reserve’ of $d_i\Delta\mu_i$ in each node i (d_i is the number of incident arcs). This gives an upper bound on the total change caused by restoring feasibility of incident arcs in all subsequent phases. We call a node Δ -positive if $e_i > d_i\Delta\mu_i$, Δ -negative if $e_i < d_i\Delta\mu_i$ and Δ -neutral if $e_i = d_i\Delta\mu_i$. Δ -negative nodes may become negative ($e_i < 0$) at a later phase, and therefore we maintain the stronger condition $\mu_i = 1/M_i$ for them. We send flow from Δ -positive nodes to Δ -negative and Δ -neutral ones. Thereby we treat some nodes with $e_i > 0$ as sinks and increase their excess further; however, as Δ decreases, such nodes may gradually become sources.

For the sink version, described in Section 7, we perform this algorithm with $M_i = B^n + 1$ if $i \neq t$ and $M_t = 1$. We shall show that this returns an optimal solution. We remark that the highest gain path algorithm [17] can also be modified to a purely scaling algorithm using the symmetric formulation, that enables to start from an arbitrary non-feasible solution and thereby eliminate the initial cycle-cancelling phase.

4.3 Minimum-cost circulations with separable convex costs

A natural and well-studied nonlinear extension of minimum-cost circulations is replacing each arc cost c_e by a convex function C_e . We are given a directed graph $G = (V, E)$ with lower and upper arc capacities $\ell, u : E \rightarrow \mathbb{R}$, convex cost functions $C_e : [\ell_e, u_e] \rightarrow \mathbb{R}$ on the arcs, and node demands $b : V \rightarrow \mathbb{R}$ with $\sum_{i \in V} b_i = 0$. Our aim is to minimize $\sum_{e \in E} C_e(f_e)$ for feasible circulations f . This is a widely applicable framework, see [2, Chapter 14].

This is a convex optimization problem, and optimality can be described by the KKT conditions. Let $C_e^+(\alpha)$ denote the left derivative of C_e . As before, for a feasible circulation f define the auxiliary graph $G_f = (V, E_f)$. Let C_{ij} denote the original function if ij is a forward arc and let $C_{ji}(\alpha) = C_{ij}(-\alpha)$ on backward arcs. f is optimal if and only if there exists no cycle C in E_f with $\sum_{e \in C} C_e^+(f_e) < 0$. In dual terms, f is optimal if and only if there exists a potential $\pi : V \rightarrow \mathbb{R}$ with $\pi_j - \pi_i \leq C_{ij}^+(f_{ij})$ for all $ij \in E_f$.

Both the minimum mean cycle cancellation and the capacity scaling algorithms can be naturally extended to this problem with polynomial (but not strongly polynomial) running time bounds. However, these two approaches relax the optimality conditions in fundamentally different ways.

Cycle cancellation was adapted by Karzanov and McCormick [25]. The algorithm subsequently cancels cycles in E_f with minimum mean value respect to the $C_e^+(f_e)$ values. The only difference is that the flow augmentation around such a cycle might be less than what residual capacities would enable, in order to maintain

$$\pi_j - \pi_i \leq C_{ij}^+(f_{ij}) + \varepsilon \quad \forall ij \in E_f \quad (2)$$

for the current potential π and scaling parameter ε .

For capacity scaling, Hochbaum and Shanthikumar [19] developed the following framework based on previous work of Minoux [26] (see also [2, Chapter 14.5]). The algorithm consists of Δ -phases. In the Δ -phase, each C_e is linearized with granularity Δ .

Let E_f^Δ denote the Δ -residual network. We will maintain Δ -optimality, that is, there exists a

potential π such that

$$\pi_j - \pi_i \leq \frac{C_{ij}(f_{ij} + \Delta) - C_{ij}(f_{ij})}{\Delta} \quad \forall ij \in E_f^\Delta. \quad (3)$$

Let $\theta_\Delta(ij)$ denote the quantity on the right hand side. If we increase flow on some arc ij by Δ on some ij for which equality holds, the resulting pseudoflow will remain Δ -optimal. We will always send Δ units of flows from a node s with $e_s > \Delta$ to a node t with $e_t < -\Delta$ on a minimum-cost path in E_f^Δ with respect to $\theta_\Delta(ij)$. By the above observation, this maintains Δ -optimality.

When moving to the next scaling phase replacing Δ by $\Delta/2$, we change to a better linear approximation of the C_e 's. Therefore, (3) may get violated not only because $E_f^{\Delta/2}$ contains more arcs than E_f^Δ does, but also on arcs already included in E_f^Δ . Yet it turns out that modifying each f_{ij} value by at most $\Delta/2$, (3) can be re-established. This creates new (positive and negative) excesses of total at most $m\Delta$.

Recently, [43] gave a strongly polynomial capacity scaling algorithm for a class of objective functions, including convex quadratic objectives, and Fisher's market with linear and with spending constraint utilities (based on Shmyrev's formulation). Let us also remark that the results [25] and [19] actually address much more general problems: minimizing convex objectives over polyhedra given by matrices with bounded subdeterminants. The framework of [19] needs weaker assumptions on the objective function and on the oracle.

4.4 Concave generalized flows

As we have seen, both the cycle cancelling and capacity scaling approaches for minimum-cost circulations naturally extend to separable concave cost functions. Similarly, our algorithm in Section 6 for concave gain functions is a natural extension of the generalized flow algorithm in Section 5.

Nevertheless, we were not able to extend any previous generalized flow algorithm for concave gains. Shigeno's [36] approach was to extend the FAT-PATH algorithm of [13]. However, [36] obtains polynomial running time bounds only for restricted classes of gain functions. The algorithm consists of two procedures applied alternately similar to FAT-PATH: a cycle cancellation phase to generate excess on cycles with positive gains, and a path augmentation phase to transport new excess to the sink in chunks of Δ . For both phases, previous methods naturally extend: cycle cancelling is performed analogously to [25], whereas path augmentation to [19]. Unfortunately, fitting the two different methods together is problematic and does not yield polynomial running time.

The main reason is that the two approaches rely on fundamentally different kinds of approximation of the nonlinear gain functions. While for generalized flows, a cycle cancelling phase completely eliminates flow generating cycles, here we can only get an approximate solution allowing some small positive gain on cycles at termination. In terms of residual arcs, we terminate with a condition analogous to (2) for concave cost flows. However, the path augmentation phase needs a linearization of the gain functions analogous to (3). Notice that for $\varepsilon = 0$, (2) implies (3) for arbitrary Δ . Yet if some small error $\varepsilon > 0$ is allowed, then no general guarantee can be given so that (3) hold for a certain value Δ .

For this reason, our goal was to avoid using the two different frameworks simultaneously. It turns out that our scaling based linear generalized flow algorithm (outlined in Section 4.2) smoothly extends to this general setting. We use the local linearization $\theta_\Delta^\mu(ij)$ of Γ_{ij} used by Shigeno, an analogue of (3). In the Δ -phase, we consider the graph of Δ -fat arcs, and maintain $\theta_\Delta^\mu(ij) \leq 1$ on them.

When moving from a Δ -phase to a $\Delta/2$ -phase in the linear algorithm, the only reason for infeasibility is due to $\Delta/2$ -fat arcs that were not Δ -fat. In contrast, feasibility can be violated on Δ -fat arcs as well, as $\theta_\Delta^\mu(ij) \leq 1 < \theta_{\Delta/2}^\mu(ij)$ may happen due to the finer linear approximation of

the gain functions in the $\Delta/2$ -phase. Fortunately, feasibility can be restored in this case as well, by changing the flow on each arc by a small amount.

5 Linear generalized flow algorithm

In this section, we investigate the symmetric formulation of the generalized flow problem. In describing the optimality conditions, we also allow infinite M_i values to incorporate the sink version. However, in the algorithmic parts, we restrict ourselves to finite M_i values.

We describe optimality conditions in Section 5.1. Notion and results here are well-known in the generalized flow literature, thus we do not include references. Section 5.2 introduces Δ -fat arcs and Δ -conservative labelings, the feasibility framework in the Δ -phase. Section 5.3 describes the subroutine TIGHTEN-LABEL, an adaptation of Dijkstra's algorithm that finds highest gain Δ -fat paths. The main algorithm is exhibited in Section 5.4. Analysis and running time bounds are given in Section 5.5. The final step of the algorithm is deferred to Section 5.6, where we show that when the total relabeled excess is sufficiently small, an optimal solution can be found by a single maximum flow computation.

5.1 Optimality conditions

Let $G = (V, E)$ be a network with lower and upper capacities ℓ, u and node demands b . In the sequel, we assume all lower capacities are 0. Every problem instance of symmetric generalized flows can be simply transformed to an equivalent one in this form in the following way. For each arc $ij \in E$, increase the node demand b_i by ℓ_{ij} and decrease b_j by $-\gamma_{ij}\ell_{ij}$. Modify the lower capacity of ij to 0 and the upper to $u_{ij} - \ell_{ij}$.

For a pseudoflow f , we define the residual network $G_f = (V, E_f)$ as follows. Let $ij \in E_f$, if $ij \in E$ and $f_{ij} < u_{ij}$ or if $ji \in E$ and $f_{ji} > 0$. The first type of arcs are called *forward arcs*, while the second type are the *backward arcs*. For a forward arc ij , let γ_{ij} be the same as in the original graph. For a backward arc ji , let $\gamma_{ji} = 1/\gamma_{ij}$. Also, we define $f_{ji} = -\gamma_{ij}f_{ij}$ for every backward arc $ji \in E_f$. For backward arcs, the capacities are $\ell_{ji} = -\gamma_{ji}u_{ji}$ and $u_{ji} = 0$. By increasing (decreasing) f_{ji} by α , we mean decreasing (increasing) f_{ij} by α/γ_{ij} .

Let $P = i_0 \dots i_k$ be a walk in the auxiliary graph E_f . By sending α units of flow along P , we mean increasing each $f_{i_h i_{h+1}}$ by $\alpha \prod_{0 \leq t < h} \gamma_{i_t i_{t+1}}$. We assume α is chosen small enough so that no capacity gets violated. Note that this decreases e_{i_0} by α , increases e_{i_k} by $\alpha \prod_{0 \leq t < k} \gamma_{i_t i_{t+1}}$, and leaves the other e_i values unchanged.

Let $C = i_0 \dots i_{k-1} i_0$ be a cycle in E_f . Sending α units of flow on C from i_0 modifies only e_{i_0} , increasing by $(\gamma(C) - 1)\alpha$ for $\gamma(C) = \prod_{e \in C} \gamma_e$. C is called a *flow generating cycle* if $\gamma(C) > 1$. On such a cycle for any choice of $i_0 \in V(C)$, we can create an excess of $(\gamma(C) - 1)\alpha$ by sending α units around (assuming that α is sufficiently small so that no capacity constraints are violated).

The pair (C, P) is called a *generalized augmenting path (GAP)* in the following cases:

- (a) C is a flow generating cycle, $i_0 \in V(C)$, $t \in V$ is a node with $e_t < 0$, and P is a path in E_f from i_0 to t ($i_0 = t$, $P = \emptyset$ is possible);
- (b) $C = \emptyset$, and P is a path between two nodes s and t with $e_s > 0$, $e_t < 0$;
- (c) $C = \emptyset$, and P is a path between s and t with $e_s \leq 0$, $e_t < 0$ and $\gamma(P) = \prod_{e \in C} \gamma_e > M_s/M_t$.

Lemma 5.1. *If f is an optimal solution, then no GAP exists.*

Proof. In case (a), we can send some $\alpha > 0$ units of flow around C from i_0 , and then send the generated $(\gamma(C) - 1)\alpha$ excess from i_0 to t along P . For sufficiently small positive value of α , this is possible without violating the capacity constraints and it decreases the excess discrepancy. In case (b), we can decrease the excess discrepancy at t while only decreasing a positive excess at s . In case (c), although $M_s\kappa_s$ increases, $M_t\kappa_t$ decreases by a larger amount. \square

The dual description of an optimal solution is in terms of *relabelings* with a label function $\mu : V \rightarrow \mathbb{R}_{>0} \cup \{\infty\}$. For each node $i \in V$, let us rescale the flow on all arcs $ij \in E$ by μ_i : let $f_{ij}^\mu = f_{ij}/\mu_i$. We get a problem equivalent to the original one with relabeled gains $\gamma_{ij}^\mu = \gamma_{ij}\mu_i/\mu_j$. Accordingly, the relabeled demands, excesses, and capacities are $b_i^\mu = b_i/\mu_i$, $e_i^\mu = e_i/\mu_i$, and $u_{ij}^\mu = u_{ij}/\mu_i$. A relabeling is *conservative*, if for any residual arc $ij \in E_f$, $\gamma_{ij}^\mu \leq 1$, that is, no arc may increase the relabeled flow. Furthermore, for each $i \in V$, $\mu_i \geq 1/M_i$ is required and equality must hold whenever $e_i < 0$.

We use the conventions $\infty \cdot 0 = 0$ and $\infty/\infty = 0$. Accordingly, if $\mu_i = \infty$, we define $b_i^\mu = e_i^\mu = 0$, and $\gamma_{ji}^\mu = 0$ for all arcs $ji \in E_f$. If $ij \in E_f$, $\mu_i = \infty$ and $\mu_j < \infty$, then $\gamma_{ij}^\mu = \infty$. Consequently, if μ is conservative, then $\mu_i < \infty$ for any $i \in V$ for which there exists a path in E_f from i to any node $t \in V$ with $e_t < 0$. Also, if μ is conservative, there exists no flow generating cycle on the node set $\{i : \mu_i < \infty\}$. This is since for a cycle C , $\gamma(C) = \prod_{ij \in C} \gamma_{ij} = \prod_{ij \in C} \gamma_{ij}^\mu$.

Theorem 5.2. *For a pseudoflow f , the following are equivalent.*

- (i) f is an optimal solution to the symmetric generalized flow problem.
- (ii) E_f contains no generalized augmenting paths.
- (iii) There exists a conservative relabeling μ with $e_i = 0$ whenever $1/M_i < \mu_i < \infty$.

Proof. The equivalence of (i) and (iii) is by linear programming duality, with μ_i being the reciprocal of the dual variable corresponding to the inequality $e_i + \kappa_i \geq 0$. (i) implies (ii) by Lemma 5.1.

It is left to show that (ii) implies (iii). If the excess discrepancy is 0 (that is, $e_i \geq 0$ for all $i \in V$), then $\mu \equiv \infty$ is conservative. Otherwise, let $N = \{t : e_t < 0\}$. If E_f contains no directed path from $i \in V$ to N , then let $\mu_i = \infty$. For the other nodes $i \in V$, let μ_i be the smallest possible value of $1/(\gamma(P)M_t)$ for $\gamma(P) = \prod_{e \in P} \gamma_e$, where P is a walk in E_f starting from i and ending in a node $t \in N$. By (ii), this is well-defined, since all cycles can be removed from a walk P without decreasing $\gamma(P)$.

μ clearly satisfies $\gamma_{ij}\mu_i/\mu_j \leq 1$. We shall prove $\mu_i \geq 1/M_i$ for each $i \in V$ and $\mu_t = 1/M_t$ for each $t \in N$. If $e_i > 0$, then no such path P may exist as it would give a type (b) GAP. Consequently, $e_i = \infty$. If $e_i \leq 0$, then $\mu_i \geq 1/M_i$ as otherwise the optimal P defining μ_i would be a type (c) GAP. Finally, if $t \in N$, then $\mu_t \leq 1/M_t$ as the gain of the path $P = \emptyset$ is defined as 1. \square

5.2 Δ -conservative labels

The residual capacity of arc $ij \in E_f$ is $u_{ij} - f_{ij}$ (for a backward arc ij , this is $\gamma_{ji}f_{ji}$). In contrast, we define the *fatness* of $ij \in E_f$ by $s_f(ij) = \gamma_{ij}(u_{ij} - f_{ij})$ (on backward arcs, $s_f(ij) = f_{ji}$). The fatness expresses the maximum possible flow increase in j if we saturate ij . This notion enables us to identify arcs that can participate in fat paths during the algorithm. In accordance with the other variables, the relabeled fatness is defined as $s_f^\mu(ij) = s_f(ij)/\mu_j$.

For the scaling parameter $\Delta > 0$, we define the following relaxation of conservativity. Let $\mu : V \rightarrow \mathbb{R}_{>0} \cup \{\infty\}$ be a label function. Recall that d_i is the total number of arcs incident to i . A node $i \in V$ is called Δ -negative if $e_i^\mu < d_i\Delta$, Δ -neutral if $e_i^\mu = d_i\Delta$ and Δ -positive if $e_i^\mu > d_i\Delta$.

The Δ -fat graph $E_f^\mu(\Delta)$ is the set of residual arcs of relabeled fatness at least Δ :

$$E_f^\mu(\Delta) = \{ij : ij \in E_f : s_f^\mu(ij) \geq \Delta\}.$$

Arcs in $E_f^\mu(\Delta)$ will be called Δ -fat arcs. The labeling μ is called Δ -conservative, if $\gamma_{ij}^\mu \leq 1$ holds for every $ij \in E_f^\mu(\Delta)$, and $\mu_i \geq 1/M_i$ for all $i \in V$. Further, for every Δ -negative node i , we require $\mu_i = 1/M_i$.

Observe that no nodes with $\mu_i = \infty$ may be present in a Δ -conservative labeling. If $\Delta \leq \Delta'$, then a Δ -conservative labeling is also Δ' -conservative, and 0-conservativity is identical to conservativity. Let $Ex^\mu(f) = \sum_{i \in V} \max\{e_i^\mu, 0\}$ and $Ex_\Delta^\mu(f) = \sum_{i \in V} \max\{e_i^\mu - d_i\Delta, 0\}$ denote the total relabeled excess and total modified relabeled excess for Δ , respectively. Note that $Ex^\mu(f) \leq Ex_\Delta^\mu(f) + 2m\Delta$.

Lemma 5.3. *Let f be a pseudoflow with a Δ -conservative labeling μ . Let $0 \leq \Delta' < \Delta$. Then there exists a flow \bar{f} such that \bar{f} and μ are Δ' -conservative and $Ex_{\Delta'}^\mu(\bar{f}) \leq Ex_\Delta^\mu(f) + 3m(\Delta - \Delta')$*

Proof. We shall construct \bar{f} by modifying f on each arc independently. For Δ -fat arcs, Δ -conservativity guarantees $\gamma_{ij}^\mu \leq 1$. Consider an arc $ij \in E_f^\mu(\Delta') - E_f^\mu(\Delta)$, that is,

$$\Delta' \leq \gamma_{ij}^\mu(u_{ij}^\mu - f_{ij}^\mu) < \Delta$$

and assume $\gamma_{ij}^\mu > 1$. Let us set $\bar{f}_{ij} = \min\{u_{ij}, f_{ij} + \frac{(\Delta - \Delta')\mu_j}{\gamma_{ij}}\}$. Then ij cannot be Δ' -fat, since either $ij \notin E_f$, or $\gamma_{ij}^\mu(u_{ij}^\mu - \bar{f}_{ij}^\mu) < \Delta'$.

$f_{ji}^\mu - \bar{f}_{ji}^\mu = \gamma_{ij}^\mu(\bar{f}_{ij}^\mu - f_{ij}^\mu) \leq \Delta - \Delta'$. Also, $\bar{f}_{ij}^\mu - f_{ij}^\mu \leq \Delta - \Delta'$ holds because of $\gamma_{ij}^\mu > 1$. We also have to consider the possibility that ji is Δ' -fat for \bar{f} . In this case, conservativity is guaranteed since $\gamma_{ji}^\mu = 1/\gamma_{ij}^\mu < 1$.

To complete the proof of Δ' -conservativity, we show that \bar{f} has no Δ' -negative nodes with $\mu_i > 1/M_i$. As we have seen, both f_{ij} and $\gamma_{ij}f_{ij}$ change by at most $\Delta - \Delta'$. Consequently for every $i \in V$, the total possible change of the relabeled flow on arcs incident to i is $d_i(\Delta - \Delta')$. A node is nonnegative for Δ if $e_i^\mu \geq d_i\Delta$ and for Δ' if $e_i^\mu \geq d_i\Delta'$. Therefore, a nonnegative node cannot become Δ' -negative, proving the claim.

Further, $Ex_{\Delta'}^\mu(\bar{f}) \leq Ex_\Delta^\mu(f) + \sum_{i \in V} d_i(\Delta - \Delta')$, and on each arc, at most $\Delta - \Delta'$ units of new excess is created. This gives $Ex_{\Delta'}^\mu(\bar{f}) \leq Ex_\Delta^\mu(f) + 3m(\Delta - \Delta')$. \square

The proof also gives a straightforward algorithm for finding such an \bar{f} . Let $\text{ADJUST}(\Delta, \Delta')$ denote this subroutine. In particular, $\text{ADJUST}(\Delta, 0)$ finds an \bar{f} for which μ is a conservative labeling. Further, if there are no Δ -negative nodes for f and μ , then \bar{f} is a 0-discrepancy optimal solution.

5.3 Δ -canonical labels

An arc ij is called *tight* if $\gamma_{ij}^\mu = 1$, and a directed path is *tight* if it consists of tight arcs. Given a pseudoflow f , a conservative labeling μ is called *canonical*, if for each $i \in V$ with $\mu_i < \infty$, there exists a tight path in E_f from i to a negative node. Analogously, for $\Delta > 0$, a labeling is called Δ -canonical, if it is Δ -conservative, and for each $i \in V$ there exists a tight path in $E_f^\mu(\Delta)$ from i to some Δ -negative or Δ -neutral node. Such a path is a highest gain Δ -fat path as in [13]. (Note that $\mu_i < \infty$ for every $i \in V$ for Δ -canonical labelings, and also that paths are allowed to end in Δ -neutral nodes, in contrast to canonical labelings.) By 0-conservative labeling we mean a conservative one.

Given a Δ -conservative relabeling μ which is not canonical, $\text{TIGHTEN-LABEL}(f, \mu, \Delta)$ replaces μ by a Δ -canonical labeling μ' with $\mu'_i \geq \mu_i$ for each $i \in V$.

We first describe $\text{TIGHTEN-LABEL}(f, \mu, 0)$, when it is essentially a multiplicative interpretation of Dijkstra's algorithm. Let $V' \subseteq V$ be the set of nodes i with a directed path in E_f from i to a negative node. For nodes in $V \setminus V'$, let us set $\mu_i = \infty$. Let $S \subseteq V'$ be the set of nodes i for which there exists a (possibly empty) tight path for the current μ to a negative node. In each step of the algorithm, S will be extended by at least one element, and we terminate if $S = V'$, when the current relabeling is canonical.

If $V' \setminus S \neq \emptyset$, let us multiply μ_i for each $i \in V' \setminus S$ by α defined as

$$\alpha = \min \left\{ \frac{1}{\gamma_{ij}^\mu} : ij \in E_f, i \in V' \setminus S, j \in S \right\}.$$

By the definition of S , $\alpha > 1$, and after multiplying by α , at least one arc $ij \in E_f$ with $i \in V' \setminus S$, $j \in S$ will become tight. Tight arcs inside S also remain tight, hence S is extended by at least one node. Also, the choice of α guarantees that μ remains conservative. Note that arcs with $j \in V \setminus S$ may disappear from $E_f^\mu(\Delta)$ as their fatness decreases.

Let us now describe $\text{TIGHTEN-LABEL}(f, \mu, \Delta)$ for $\Delta > 0$. The main difference is that increasing μ_i may turn a Δ -positive node into Δ -neutral. We have to stop increasing μ_i at this point and add i to S . (This is in accordance with the goal that we allow tight paths to Δ -neutral nodes as well.) In each phase of the algorithm, let $S \subseteq V$ denote the subset of nodes that have a (possibly empty) tight path in $E_f^\mu(\Delta)$ to a Δ -negative or Δ -neutral node. S is initialized as the set of Δ -negative and Δ -neutral nodes, and is extended by at least one element per phase. The algorithm terminates once $S = V$.

In every phase, we multiply μ_i for every $i \in V \setminus S$ by the same factor $\alpha > 1$. Consider an arc $ij \in E_f(\Delta)$ with $i \in V - S$, $j \in S$. This must satisfy $\gamma_{ij}^\mu < 1$. Increasing μ_i increases γ_{ij}^μ . Note that the fatness $s_f^\mu(ij) = s_f(ij)/\mu_j$ is not changed as it is not dependent of μ_i . By definition, all nodes in $V - S$ are Δ -positive. When increasing μ_i , e_i^μ decreases and therefore i may become Δ -neutral. At this point, we have to stop to avoid creating new Δ -negative nodes. Let us define

$$\alpha = \min \left\{ \min \left\{ \frac{1}{\gamma_{ij}^\mu} : ij \in E_f^\mu(\Delta), i \in V \setminus S, j \in S \right\}, \min \left\{ \frac{e_i^\mu}{d_i \Delta} : i \in V - S \right\} \right\}. \quad (4)$$

Clearly, $\alpha > 1$, and after multiplying each μ_i by α for $i \in V \setminus S$, either we obtain a new Δ -neutral node in $V - S$, or at least one arc $ij \in E_f$ with $i \in V \setminus S$, $j \in S$ will become tight. Tight arcs inside S also remain tight, and their capacities is unchanged, hence S is extended by at least one node. Also, the choice of α guarantees that μ remains Δ -conservative.

5.4 Description of the algorithm

The algorithm is shown on Figure 1. We start with $\mu_i = 1/M_i$ for every node $i \in V$, $f \equiv 0$ and $\Delta = MB^2 + 1$. Once $(2n + 6m)\Delta < 1/B^m$, then an optimal solution can be found by a single maximum flow computation, as shown in Section 5.6. In this case we terminate.

The algorithm consists of Δ -scaling phases. During the Δ -phase, we maintain a pseudoflow f along with a Δ -conservative labeling μ . The μ_i values can only increase. Let N_0 denote the set of Δ -negative and Δ -neutral nodes, and D the set of nodes i with $e_i^\mu \geq (d_i + 1)\Delta$. The Δ -phase consists of a sequence of iterations until D becomes empty. In every iteration of the algorithm, we update μ to a Δ -canonical labeling by calling $\text{TIGHTEN-LABEL}(f, \mu, \Delta)$. If $D \neq \emptyset$ still holds, we

```

Algorithm SYMMETRIC FAT-PATH
for  $i \in V$  do  $\mu_i \leftarrow \frac{1}{M_i}$ ;
for  $ij \in E$  do  $f_{ij} \leftarrow 0$ ;
 $\Delta \leftarrow MB^2 + 1$ ;
while  $(2n + 6m)\Delta \geq 1/B^m$  do
  do
    TIGHTEN-LABEL( $f, \mu, \Delta$ );
     $D \leftarrow \{i \in V : e_i > (d_i + 1)\Delta\}$ ;
     $N_0 \leftarrow \{i \in V : e_i \leq d_i\Delta\}$ ;
    pick  $s \in D, t \in N$  connected by a tight path  $P$ ;
    send  $\Delta$  units of flow along  $P$ ;
  while  $D \neq \emptyset$ ;
  ADJUST( $\Delta, \frac{\Delta}{2}$ );
   $\Delta \leftarrow \frac{\Delta}{2}$ ;
  ADJUST( $\Delta, 0$ );
  TIGHTEN-LABEL( $f, \mu, 0$ );
compute a maximum flow from nodes  $\{s \in V : e_s > 0\}$ 
  to nodes  $\{t \in V : e_t < 0\}$  using tight arcs for  $\mu$ ;
return optimal primal solution  $f$  and optimal dual solution  $\mu$ 

```

Figure 1: The algorithm for symmetric linear generalized flows

pick an arbitrary $s \in D$, and send Δ units of flow from s to some Δ -negative or Δ -neutral $t \in N_0$ on a tight path P . At the end of the Δ -phase, we modify f by $\text{ADJUST}(\Delta, \Delta/2)$, and proceed to the $\Delta/2$ -phase.

5.5 Analysis

Claim 5.4. *The initial μ is Δ -conservative, and Δ -conservativity is maintained during the entire Δ -phase.*

Proof. At the beginning $s_f^\mu(ij) = \gamma_{ij}u_{ij}/\mu_j \leq MB^2 < \Delta$ for every $ij \in E$, and E_f contains no backward arcs. Consequently, $E_f^\mu(\Delta) = \emptyset$ and Δ -conservativity trivially holds. Also, $\mu_i = 1/M_i$ holds for every node i . $\text{TIGHTEN-LABEL}(f, \mu, \Delta)$ clearly maintains Δ -conservativity. This is also maintained when sending flow, as we only use tight arcs. At the end of the Δ -phase, $\text{ADJUST}(\Delta, \Delta/2)$ transforms f to a $\Delta/2$ -conservative pseudoflow. \square

Claim 5.5. *At the beginning of every Δ -phase, $Ex_\Delta^\mu(f) \leq (2n + 3m)\Delta$.*

Proof. In the first phase, $Ex_\Delta^\mu(f) \leq M(\sum_{i \in V} b_i) = MB < (2n + 3m)\Delta$. Once we finish all iterations in the Δ -phase, $D = \emptyset$ implies $Ex_\Delta^\mu(f) \leq n\Delta$. In $\text{ADJUST}(\Delta, \Delta/2)$, we increase the excess by at most $3m\Delta/2$ (Lemma 5.3.) Hence at the beginning of the $\Delta/2$ -phase, $Ex_{\Delta/2}^\mu(f) \leq (2n + 3m)\Delta/2$, proving the claim. \square

Lemma 5.6. *A Δ -phase consists of at most $2n + 3m$ iterations.*

Proof. Let $\Psi(i) = \lfloor e_i^\mu / \Delta \rfloor - d_i$ if $e_i^\mu > (d_i + 1)\Delta$ and $\Psi(i) = 0$ otherwise. Consider the potential function $\Psi = \sum_{i \in V} \Psi(i)$. Clearly, $\Psi \leq Ex_\Delta^\mu(f)/\Delta$ and therefore $\Psi \leq 2n + 3m$ holds at the beginning by Claim 5.5. In the relabeling steps, Ψ may only decrease, and in every path augmentation, it decreases by exactly 1. \square

Theorem 5.7. *The algorithm runs in $O(m(m + n \log n)(m \log B + \log M))$ time.*

Proof. Δ always decreases by a factor of 2, its initial value is MB^2 and we terminate if $Ex_\Delta^\mu(f) < 1/B^m$. Hence the number of phases can be bounded by $O(m \log B + \log M)$. The number of iterations is $O(m)$ by Lemma 5.6. The running time of an iteration is dominated by the TIGHTEN-LABEL step, that can be done in $O(m + n \log n)$ time following Fredman and Tarjan's [11] implementation of Dijkstra's algorithm. \square

5.6 Moving to an optimal solution

If $(2n + 6m)\Delta < 1/B^m$ at a certain iteration of the algorithm, then by Claim 5.5, $Ex_\Delta^\mu(f) + 3m\Delta < 1/B^m$ holds. Next, we transform f to \bar{f} by ADJUST($\Delta, 0$), so that μ is a conservative labeling for \bar{f} . By Lemma 5.3, $Ex^\mu(\bar{f}) < 1/B^m$ follows. Then TIGHTEN-LABEL(\bar{f}, μ, Δ) transforms μ into a canonical labeling. By the following lemma, a single maximum flow computation yields an optimal solution. This is the standard technique how most algorithms in the literature terminate.

Lemma 5.8. *Let μ be a canonical labeling for f , and let $\tilde{G}_f = (V, \tilde{E}_f)$ be the subgraph of G_f consisting of tight arcs in E_f . If $Ex^\mu(f) < 1/B^m$, then a single maximum flow computation on G from source set $P = \{s \in V : e_s > 0, \mu_s < \infty\}$ to sink set $N = \{t \in V : e_t < 0\}$ terminates with an optimal solution.*

Proof. Consider the flow f' resulting by the maximum flow computation. Since flow was sent only on tight arcs, μ is also conservative for f' . If there are no more nodes s with $e_s^\mu > 0$, then by Theorem 5.2, f' is optimal. Assume now P' , the set of such nodes for f' is nonempty.

Let $S \subseteq V$ be the set of nodes reachable from P' using tight residual arcs in $E_{f'}$. By optimality, S contains no node with negative excess. If an arc $ij \in E$ leaves S then either ij is saturated, that is, $f'_{ij} = u_{ij}$, or $\gamma_{ij}^\mu < 1$ and $f'_{ij} = 0$. Let F_0 denote the set of these saturated arcs. Similarly, if $ij \in E$ enters S , then $f'_{ij} = 0$ must hold. Also, on all arcs ij with $i, j \in S$, $f'_{ij} > 0$ either $\gamma_{ij}^\mu = 1$ or $f'_{ij} = u_{ij}$. Let F_1 denote the set of such arcs with $f'_{ij} < u_{ij}$ (and $\gamma_{ij}^\mu = 1$), and F_2 the set of those with $f'_{ij} = u_{ij}$. Therefore,

$$\begin{aligned}
0 &\leq Ex^\mu(f') = \sum_{i \in S} e_i^\mu(f') = \sum_{i \in S} \left(\sum_{j: j \in E} \gamma_{ji}^\mu f'_{ji} - \sum_{j: j \in E} f'_{ij} - b_i^\mu \right) \\
&= \sum_{ij \in F_1} (\gamma_{ji}^\mu f'_{ji} - f'_{ji}) + \sum_{ij \in F_2} u_{ij}^\mu (\gamma_{ij}^\mu - 1) - \sum_{ij \in F_0} u_{ij}^\mu - \sum_{i \in S} b_i^\mu \\
&= \sum_{ij \in F_2} u_{ij} \left(\frac{\gamma_{ij}}{\mu_j} - \frac{1}{\mu_i} \right) - \sum_{ij \in F_0} \frac{u_{ij}}{\mu_i} - \sum_{i \in S} \frac{b_i}{\mu_i} \tag{5}
\end{aligned}$$

Let $B^* \leq B^m$ denote the product of the denominators of the γ_{ij} 's. We claim that every term in the above expression is an integer multiple of $1/B^*$. Indeed, using that μ is a canonical labeling for f , there exists a tight path from each node i to a negative node t . $1/\mu_i$ is then the product of the integer M_t and the gain factors on such a tight path and is hence an integer multiple of $1/B^*$. Similarly, every γ_{ij}/μ_j is also an integer multiple of $1/B^*$, since if $\gamma_{ij}^\mu \neq 1$, then ij cannot be on the

tight path from i to a negative node. Since the u_{ij} 's and b_i 's are integers, this verifies the claim. By the assumption of the lemma, $0 \leq Ex^\mu(f') \leq Ex^\mu(f) < 1/B^m \leq 1/B^*$. This implies $Ex^\mu(f') = 0$, that is, $\mu_i = \infty$ whenever $e_i > 0$. By Theorem 5.2, this shows optimality. \square

6 Concave generalized flows algorithm

We describe the algorithm in the same structure as for the linear case: Section 6.1 presents the optimality conditions; Δ -conservative and Δ -canonical labels are discussed in Sections 6.2 and 6.3, respectively. Section 6.4 presents the algorithm, and Section 6.5 gives its analysis.

6.1 Optimality conditions

The characterization of optimality was given in [36]; we have to modify the results slightly as we use the symmetric formulation. Let us call an arc $ij \in E$ *immense*, if $\Gamma_{ij}(\ell_{ij}) = -\infty$, and other arcs *regular*. First, let us transform the problem to an equivalent instance with (i) $\ell = 0$ for every arc and $\Gamma_{ij}(0) = 0$ for every regular; and (ii) every gain function Γ_{ij} is strictly monotone increasing on $[0, u_{ij}]$.

For (i), on each arc $ij \in E$, let us replace u_{ij} by $u_{ij} - \ell_{ij}$ and ℓ_{ij} by 0. If ij is a regular arc, we modify the gain function to $\Gamma_{ij}(\alpha + \ell_{ij}) - \Gamma_{ij}(\ell_{ij})$, and if ij is an immense arc, to $\Gamma_{ij}(\alpha + \ell_{ij})$. Accordingly for every $i \in V$, let us increase b_i by $\sum_{j:ij \in E} \ell_{ij}$, and decrease it by the sum of $\Gamma_{ji}(\ell_{ji})$'s on regular arcs.

For (ii), let us define $\tilde{u}_{ij} = \inf\{p : 0 \leq p \leq u_{ij}, \Gamma_{ij}(p) = \Gamma_{ij}(u_{ij})\}$. By concavity, $\Gamma(\tilde{u}_{ij}) = \Gamma_{ij}(u_{ij})$, and $\Gamma_{ij}(u_{ij})$ is strictly monotone increasing on the interval $[0, \tilde{u}_{ij}]$. Let us replace u_{ij} by $\tilde{u}_{ij} \leq u_{ij}$.

For a pseudoflow $f : E \rightarrow \mathbb{R}$, we define the residual network $G_f = (V, E_f)$ identical as for the generalized flow setting: $ij \in E_f$ if $ij \in E$ and $f_{ij} < u_{ij}$ or $ji \in E$ and $f_{ji} > 0$. For notational convenience, we define $f_{ji} = -\Gamma_{ij}(f_{ij})$ on backward arcs. We also define the function $\Gamma_{ji}(\alpha) : [-\Gamma_{ij}(u_{ij}), -\Gamma_{ij}(0)] \rightarrow [-u_{ij}, 0]$ by

$$\Gamma_{ji}(\alpha) = -\Gamma_{ij}^{-1}(-\alpha).$$

Hence $\Gamma_{ji}(f_{ji}) = -f_{ij}$.

The concavity of Γ_{ij} implies that for each $0 \leq \alpha < u_{ij}$, there exists the right derivative, denoted by $\Gamma_{ij}^+(\alpha)$, and for $0 < \alpha \leq u_{ij}$, there exists the left derivative $\Gamma_{ij}^-(\alpha)$. If $0 < \Delta < \Delta'$, then

$$\frac{\Gamma_{ij}(\alpha + \Delta') - \Gamma_{ij}(\alpha)}{\Delta'} \leq \frac{\Gamma_{ij}(\alpha + \Delta) - \Gamma_{ij}(\alpha)}{\Delta} \leq \Gamma_{ij}^+(\alpha), \quad (6a)$$

$$\frac{\Gamma_{ij}(\alpha) - \Gamma_{ij}(\alpha - \Delta')}{\Delta'} \geq \frac{\Gamma_{ij}(\alpha) - \Gamma_{ij}(\alpha - \Delta)}{\Delta} \geq \Gamma_{ij}^-(\alpha) \quad (6b)$$

for $0 \leq \alpha \leq u_{ij}$ and for $\Delta' \leq \alpha \leq u_{ij}$, respectively. Furthermore, if $0 < \alpha < \alpha' < u_{ij}$, then $\Gamma_{ij}^+(\alpha') \leq \Gamma_{ij}^-(\alpha') \leq \Gamma_{ij}^+(\alpha) \leq \Gamma_{ij}^-(\alpha)$. The following claim is easy verify.

Claim 6.1. For any $ij \in E$ with $0 < f_{ij} < u_{ij}$, $\Gamma_{ij}^+(f_{ij}) = 1/\Gamma_{ji}^-(f_{ji})$, $\Gamma_{ij}^-(f_{ij}) = 1/\Gamma_{ji}^+(f_{ji})$. \square

Let $P = i_0 \dots i_k$ be a walk in the auxiliary graph E_f . By sending α units of flow along P , we mean the following. First we increase $f_{i_0 i_1}$ by α and set $\beta = \Gamma_{i_0 i_1}(f_{i_0 i_1} + \alpha) - \Gamma_{i_0 i_1}(f_{i_0 i_1})$ to be the flow arriving at i_1 . In step $h = 1, \dots, k-1$, we increase the flow on $i_h i_{h+1}$ by β and set the new value of β as $\Gamma_{i_h i_{h+1}}(f_{i_h i_{h+1}} + \beta) - \Gamma_{i_h i_{h+1}}(f_{i_h i_{h+1}})$. We assume α is chosen small enough so that no capacity gets violated. Let $f^{\alpha, P}$ denote the modified flow.

If $C = i_0 \dots i_{k-1}$ is a cycle of E_f , then by sending α units of flow around C from i_0 we mean sending α units on the walk $i_0 \dots i_{k-1} i_0$. This modifies e_i only in node i_0 : if the flow increase from $i_{k-1} i_0$ is bigger than α , then e_{i_0} increases, and if it is smaller then it decreases. The next lemma characterizes when e_{i_0} can increase. For an arbitrary walk P in E_f , let $\Gamma_f^+(P) = \prod_{e \in P} \Gamma_e^+(f_e)$.

Lemma 6.2. *Let C be a cycle in E_f with $i \in V(C)$. If $\Gamma_f^+(C) > 1$ then e_i can be increased by sending some flow around C . If $\Gamma_f^+(C) \leq 1$, then it is not possible to increase e_i by sending any amount of flow around C .*

Since this property is independent from the choice of i , we simply say that C is a *flow generating cycle* if $\Gamma_f^+(C) > 1$. The lemma is an immediate consequence of the following claim.

Claim 6.3. *Let $P = i_0 i_1 \dots i_k$ be a walk in E_f . For any value of $\alpha > 0$, the flow increase in i_k for $f^{\alpha, P}$ is at most $\Gamma_f^+(P)\alpha$. On the other hand, for any $\varepsilon > 0$ there exists a $\delta > 0$ so that for any $0 < \alpha \leq \delta$, $f^{\alpha, P}$ increases e_{i_k} by at least $(\Gamma_f^+(P) - \varepsilon)\alpha$.*

Proof. The first part is trivial by concavity. We prove the second part by induction on the subpaths $P_h = i_0 \dots i_h$ for $h = 0, \dots, k$. There is nothing to prove for $h = 0$; assume we have already proved it for P_{h-1} . We want to find a $\delta > 0$ for some an $\varepsilon > 0$ satisfying the claim for P_h . First, it is possible to pick a small enough $\varepsilon^* > 0$ such that

$$\Gamma_f^+(P_h) - \varepsilon < (\Gamma_{i_{h-1}i_h}^+(f_{i_{h-1}i_h}) - \varepsilon^*)(\Gamma_f^+(P_{h-1}) - \varepsilon^*). \quad (7)$$

By the definition of $\Gamma_{i_{h-1}i_h}^+$, there exists a $\delta^* > 0$ such that for any $0 < \beta \leq \delta^*$,

$$(\Gamma_{i_{h-1}i_h}^+(f_{i_{h-1}i_h}) - \varepsilon^*)\beta \leq \Gamma_{i_{h-1}i_h}(f_{i_{h-1}i_h} + \beta) - \Gamma_{i_{h-1}i_h}(f_{i_{h-1}i_h}). \quad (8)$$

By induction for P_{h-1} and ε^* , we can choose a small enough $\delta > 0$ such that we have the following properties for every $0 < \alpha \leq \delta$: $\Gamma_f^+(P_{h-1})\alpha \leq \delta^*$, and the increase of $e_{i_{h-1}}$ for $f^{\alpha, P_{h-1}}$ is at least $\beta = (\Gamma_f^+(P_{h-1}) - \varepsilon^*)\alpha$. Then (7) and (8) show that for $0 < \alpha \leq \delta$, $f^{\alpha, P}$ increases e_{i_h} by at least $(\Gamma_f^+(P) - \varepsilon)\alpha$. \square

The definition of GAPS is identical to the linear case, with the only difference that $\Gamma_f^+(P) > M_s/M_t$ instead of $\gamma(P) > M_s/M_t$ in case (c). The following lemma can be proved similarly as Lemma 5.1.

Lemma 6.4. *If f is an optimal solution, then no GAP may exist.* \square

Relabelings are also defined analogously as for generalized flows. Given $\mu : V \rightarrow \mathbb{R}_{>0} \cup \{\infty\}$, let us define $f_{ij}^\mu = f_{ij}/\mu_i$ for each arc $ij \in E$. We get problems equivalent to the original with relabeled functions $\Gamma_{ij}^\mu(\alpha) = \Gamma_{ij}(\mu_i\alpha)/\mu_j$. Accordingly, the relabeled demands, excesses, and capacities are $b_i^\mu = b_i/\mu_i$, $e_i^\mu = e_i/\mu_i$, and $u_{ij}^\mu = u_{ij}/\mu_i$. A relabeling is *conservative*, if for any residual arc $ij \in E_f$, $\Gamma_{ij}^{\mu+}(f_{ij}^\mu) \leq 1$, that is, no edge may increase the relabeled flow. Furthermore we require $\mu_i \geq 1/M_i$ for every $i \in V$ and equality whenever $e_i < 0$.

We use the same convention for infinite μ_i values as for generalized flows. If $\mu_i = \infty$, we define $b_i^\mu = e_i^\mu = 0$, $u_{ij}^\mu = 0$ for $ij \in E$, and furthermore $\Gamma_{ji}^{\mu+}(f_{ji}^\mu) = 0$ for all arcs $ji \in E_f$. Finally, for $ij \in E_f$ with $\mu_i = \infty$, $\mu_j < \infty$, let $\Gamma_{ij}^{\mu+}(f_{ij}^\mu) = \infty$.

If μ is conservative, then if for a node $i \in V$ there exists a path from i to a node $t \in V$ with $e_t < 0$, then $\mu_i < \infty$. The following claim is also easy to verify.

Claim 6.5. $\Gamma_{ij}^{\mu^+}(\alpha) = \frac{\mu_i}{\mu_j} \Gamma_{ij}^+(\alpha)$, and $\Gamma_{ij}^{\mu^-}(\alpha) = \frac{\mu_i}{\mu_j} \Gamma_{ij}^-(\alpha)$. \square

This claim implies that for an $s - t$ walk P , $\Gamma_f^{\mu^+}(P) = \frac{\mu_s}{\mu_t} \Gamma_f^+(P)$, and thus for a cycle C , $\Gamma_f^{\mu^+}(C) = \Gamma_f^+(C)$.

Theorem 6.6 ([36]). *For a pseudoflow f , the following are equivalent.*

(i) f is an optimal solution to the symmetric version.

(ii) E_f contains no generalized augmenting paths.

(iii) There exists a conservative labeling μ with $e_i = 0$ whenever $1/M_i < \mu_i < \infty$.

Proof. The equivalence of (i) and (iii) follows by the Karush-Kuhn-Tucker conditions, with μ_i being the reciprocal of the Lagrange multiplier corresponding to $e_i + \kappa_i \geq 0$. (i) implies (ii) by Lemma 6.4. The proof of (ii) \Rightarrow (iii) is the same as in Theorem 5.2, with $\gamma(P)$ replaced by $\Gamma_f^+(P)$. \square

6.2 Δ -conservative labelings

We define the notion of Δ -conservative labeling analogously as in Section 5.3 for the linear case. Let us define the *fatness* of $ij \in E_f$ by $s_f(ij) = \Gamma_{ij}(u_{ij}) - \Gamma_{ij}(f_{ij})$ (if ij is a backward arc, this is equivalent to $s_f(ij) = f_{ji}$.) The fatness expresses the maximum possible flow increase in j if we saturate ij . This notion enables us to identify arcs that can participate in fat paths during the algorithm. In accordance with the other variables, the relabeled fatness is defined as $s_f^\mu(ij) = s_f(ij)/\mu_j$.

Consider a scaling parameter $\Delta > 0$. The Δ -fat graph $E_f^\mu(\Delta)$ is the set of residual arcs of relabeled fatness at least Δ :

$$E_f^\mu(\Delta) = \{ij : ij \in E_f : s_f^\mu(ij) \geq \Delta\}.$$

Arcs in $E_f^\mu(\Delta)$ will be called Δ -fat arcs. As in [36], we use the following linearization on Δ -fat arcs in chunks of Δ .

$$\theta_\Delta^\mu(ij) := \frac{\Delta\mu_i}{\Gamma_{ij}^{-1}(\Gamma_{ij}(f_{ij}) + \Delta\mu_j) - f_{ij}} \quad ij \in E_f^\mu(\Delta). \quad (9)$$

This is well-defined since $\Gamma_{ij}(f_{ij}) + \Delta\mu_j \leq \Gamma_{ij}(u_{ij})$ for Δ -fat arcs. Note that $\Gamma_{ij}(\cdot)$ is linear, i.e. $\Gamma_{ij}(\alpha) = \gamma_{ij}\alpha$, we have $\theta_\Delta^\mu(ij) = \gamma_{ij}$. An equivalent way of writing (9) is

$$\theta_\Delta^\mu(ij) = \frac{\Delta}{\Gamma_{ij}^{\mu^{-1}}(\Gamma_{ij}^\mu(f_{ij}^\mu) + \Delta) - f_{ij}^\mu} \quad ij \in E_f^\mu(\Delta). \quad (10)$$

Also, if the reverse arc ji is Δ -fat, then using $\Gamma_{ji}(f_{ji}) = -f_{ij}$ and $\Gamma_{ji}^{-1} = -\Gamma_{ij}(-\alpha)$, we get

$$\theta_\Delta^\mu(ji) = \frac{\Delta\mu_j}{\Gamma_{ij}(f_{ij}) - \Gamma_{ij}(f_{ij} - \Delta\mu_i)}. \quad (11)$$

Consider a label function $\mu : V \rightarrow \mathbb{R}_{>0} \cup \{\infty\}$. A node $i \in V$ is called Δ -negative if $e_i^\mu < d_i\Delta$, Δ -neutral if $e_i^\mu = d_i\Delta$ and Δ -positive if $e_i^\mu > d_i\Delta$. The labeling μ is Δ -conservative, if $\theta_\Delta^\mu(ij) \leq 1$ holds for every $ij \in E_f^\mu(\Delta)$. Furthermore, we require $\mu_i \geq 1/M_i$ for all $i \in V$, with equality for every Δ -negative node i .

Using the convexity of Γ^{-1} , it can be shown that if μ is a Δ -conservative labeling then it is Δ' -conservative for all $\Delta' \geq \Delta$. Let $Ex^\mu(f) = \sum_{i \in V} \max\{e_i^\mu, 0\}$ and $Ex_\Delta^\mu(f) = \sum_{i \in V} \max\{e_i^\mu - d_i\Delta, 0\}$ denote the total relabeled excess of positive and Δ -positive nodes, respectively.

The key importance of Δ -conservativity is that it is maintained when sending Δ units of flow on arcs with $\theta_\Delta^\mu(ij) = 1$. This is formulated in the next simple lemma.

Lemma 6.7. *Assume μ is Δ -conservative, and let $ij \in E_f^\mu(\Delta)$ be an arc with $\theta_\Delta^\mu(ij) = 1$. If we increase f_{ij}^μ by Δ , then $\Gamma(f_{ij}^\mu)$ also increases by Δ , and Δ -conservativity is maintained.*

Proof. $\theta_\Delta^\mu(ij) = 1$ is equivalent to $\Gamma_{ij}^\mu(f_{ij}^\mu + \Delta) = \Gamma_{ij}^\mu(f_{ij}^\mu) + \Delta$, showing the first part. Let $\bar{f}_{ij} = f_{ij} + \Delta\mu_i$ be the modified flow. For the second part, $\theta_\Delta^\mu(ij) \leq 1$ for \bar{f}_{ij} easily follows from convexity. Further, observe (11) shows that we get $\theta_\Delta^\mu(ji) = 1$ for \bar{f}_{ij} . This gives Δ -conservativity for the modified flow as all other arcs are left unchanged. \square

In contrast to Lemma 5.3, the following claim only enables to transform a Δ -conservative labeling can be transformed to a $\Delta/2$ -conservative one, instead of arbitrary $\Delta' < \Delta$. The reason is that besides the set of Δ' -fat arcs being larger than the Δ -fat arcs, we may have Δ -fat arcs with $\theta_\Delta^\mu(ij) \leq 1 < \theta_{\Delta'}^\mu(ij)$.

Lemma 6.8. *Let f be a pseudoflow with a Δ -conservative labeling μ . Then there exists a flow \bar{f} such that μ is $\Delta/2$ -conservative for \bar{f} and $Ex_{\Delta/2}^\mu(\bar{f}) \leq Ex_\Delta^\mu(f) + \frac{3}{2}m\Delta$.*

Proof. Consider a $\Delta/2$ -fat arc ij with $\theta_{\Delta/2}^\mu(ij) > 1$ for f , that is,

$$\Gamma_{ij}^{-1} \left(\Gamma_{ij}(f_{ij}) + \frac{\Delta}{2}\mu_j \right) - f_{ij} < \frac{\Delta}{2}\mu_i. \quad (12)$$

There are two possible scenarios: (a) ij was not Δ -fat, that is,

$$\frac{\Delta}{2}\mu_j \leq \Gamma_{ij}(u_{ij}) - \Gamma_{ij}(f_{ij}) < \Delta\mu_j, \quad (13)$$

or (b) ij was also a Δ -fat arc. Then by Δ -conservativity,

$$\Gamma_{ij}^{-1}(\Gamma_{ij}(f_{ij}) + \Delta\mu_j) - f_{ij} \geq \Delta\mu_i. \quad (14)$$

In both cases, let us define

$$\bar{f}_{ij} = \Gamma_{ij}^{-1} \left(\Gamma_{ij}(f_{ij}) + \frac{\Delta}{2}\mu_j \right).$$

$\Delta/2$ -fatness of ij guarantees that this is well-defined. In case (a), we claim that ij is not $\Delta/2$ -fat for \bar{f} . Indeed,

$$\Gamma_{ij}(u_{ij}) - \Gamma_{ij}(\bar{f}_{ij}) = \Gamma_{ij}(u_{ij}) - \left(\Gamma_{ij}(f_{ij}) + \frac{\Delta}{2}\mu_j \right) < \frac{\Delta}{2}\mu_j.$$

The last inequality follows by the second part of (13). In case (b), we claim that if ij is a $\Delta/2$ -fat arc for \bar{f} then $\theta_{\Delta/2}^\mu(ij) \leq 1$ must hold for \bar{f} . Indeed, if we subtract (12) from (14), we get

$$\Gamma_{ij}^{-1}(\Gamma_{ij}(f_{ij}) + \Delta\mu_j) - \Gamma_{ij}^{-1} \left(\Gamma_{ij}(f_{ij}) + \frac{\Delta}{2}\mu_j \right) > \frac{\Delta}{2}\mu_i,$$

and by substituting \bar{f}_{ij} , it follows that

$$\Gamma_{ij}^{-1} \left(\Gamma_{ij}(\bar{f}_{ij}) + \frac{\Delta}{2}\mu_j \right) - \bar{f}_{ij} > \frac{\Delta}{2}\mu_i,$$

that is, $\theta_{\Delta/2}^\mu(ij) < 1$ for \bar{f} .

We next show that if ji is also a $\Delta/2$ -fat arc for \bar{f} , then $\theta_{\Delta/2}^\mu(ji) \leq 1$ holds for \bar{f} . Indeed, using (11), $\theta_{\Delta/2}^\mu(ji) \leq 1$ for \bar{f} is equivalent to

$$\Gamma_{ij}(\bar{f}_{ij}) - \Gamma_{ij}\left(\bar{f}_{ij} - \frac{\Delta}{2}\mu_i\right) \geq \frac{\Delta}{2}\mu_j.$$

Equivalently,

$$\Gamma_{ij}(f_{ij}) + \frac{\Delta}{2}\mu_j - \Gamma_{ij}\left(\bar{f}_{ij} - \frac{\Delta}{2}\mu_i\right) \geq \frac{\Delta}{2}\mu_j.$$

Subtracting $\frac{\Delta}{2}\mu_j$, rearranging and applying the monotone increasing function Γ_{ij}^{-1} , we get $f_{ij} \geq \bar{f}_{ij} - \frac{\Delta}{2}\mu_i$, that follows from (12) by substituting \bar{f}_{ij} .

We define \bar{f}_{ij} the above way whenever ij is a $\Delta/2$ -fat arc with $\theta_\mu(ij) > 1$. (As a simple consequence of concavity, this cannot be the case for both ij and ji .) If this does not hold for neither ij nor ji , then let $\bar{f}_{ij} = f_{ij}$. The next claim compares f_{ij} and $\Gamma(f_{ij})$ to \bar{f}_{ij} and $\Gamma(\bar{f}_{ij})$.

Claim 6.9. $|f_{ij}^\mu - \bar{f}_{ij}^\mu| \leq \frac{\Delta}{2}$ and $|\Gamma_{ij}^\mu(\bar{f}_{ij}^\mu) - \Gamma_{ij}^\mu(f_{ij}^\mu)| \leq \frac{\Delta}{2}$.

Proof. There is nothing to prove if $\bar{f}_{ij} = f_{ij}$. Assume f_{ij} was increased as above (the statement is equivalent for ij and ji). The first part follows by (12). By the definition of \bar{f}_{ij} ,

$$\Gamma_{ij}(\bar{f}_{ij}) - \Gamma_{ij}(f_{ij}) = \Gamma_{ij}(f_{ij}) + \frac{\Delta}{2}\mu_j - \Gamma_{ij}(f_{ij}) = \frac{\Delta}{2}\mu_j,$$

giving the second part. \square

For $\Delta/2$ -conservativity, we also need to show that \bar{f} has no $\Delta/2$ -negative nodes with $\mu_i > 1/M_i$. By the above claim, the total possible change of relabeled flow on arcs incident to i is $d_i\Delta/2$. A node is nonnegative for Δ if $e_i^\mu \geq d_i\Delta$ and for $\Delta/2$ if $e_i^\mu \geq d_i\Delta/2$. Consequently, a Δ -nonnegative node cannot become $\Delta/2$ -negative.

Finally, $Ex_{\Delta/2}^\mu(f) \leq Ex_\Delta^\mu(f) + \sum_{i \in V} d_i\Delta/2$, and each arc is responsible for creating at most $\Delta/2$ units of new excess. This gives $Ex_{\Delta/2}^\mu(\bar{f}) \leq Ex_\Delta^\mu(f) + \frac{m}{2}\Delta$, as required. \square

The subroutine $\text{ADJUST}(\Delta)$ performs the simple modifications described in the proof (in contrast to the linear case, this function has only one parameter).

6.3 Δ -canonical labelings

Given a pseudoflow f and a Δ -conservative labeling μ , the arc $ij \in E_f^\mu(\Delta)$ is called *tight* if $\theta_\Delta^\mu(ij) = 1$. A directed path in $E_f^\mu(\Delta)$ is called tight if it consists of tight arcs. μ is a Δ -canonical labeling, if from each node i there exists a tight path to a Δ -negative or to a Δ -neutral node. Such a path is approximately a highest gain Δ -fat augmenting path. The subroutine $\text{TIGHTEN-LABEL}(f, \mu, \Delta)$ returns a Δ -canonical label $\mu' \geq \mu$ for a Δ -conservative label μ . This is almost identical to the algorithm described in Section 5.3. The only difference is in the definition of the multiplier α , which is given by (4) for the linear case. Instead, we define

$$\alpha = \min \left\{ \min \left\{ \frac{1}{\theta_\Delta^\mu(ij)} : ij \in E_f^\mu(\Delta), i \in V \setminus S, j \in S \right\}, \min \left\{ \frac{e_i^\mu}{d_i\Delta} : i \in V - S \right\} \right\}.$$

In an iteration, we multiply every μ_i by α for $i \in V - S$, where S is the set of nodes from which there exists a tight path to a Δ -negative or a Δ -neutral node. We claim that as in the linear case,

this maintains Δ -conservativity, and extends S by at least one node. This is a simple consequence of the fact that multiplying μ_i by α multiplies $\theta_\Delta^\mu(ij)$ by α for every incident arc ij .

To verify that μ remains Δ -conservative, we also have to check $\theta_\Delta^\mu(ij) \leq 1$ on all arcs $ij \in E_f^\Delta(\mu)$, $j \in V - S$. This follows by the convexity of Γ_{ij}^{-1} .

6.4 The main algorithm

```

Algorithm SYMMETRIC CONCAVE FAT-PATH
for  $i \in V$  do  $\mu_i \leftarrow \frac{1}{M_i}$ ;
for  $ij \in E$  do  $f_{ij} \leftarrow u_{ij}$ ;
 $\Delta \leftarrow MU + 1$ ;
while  $(2n + 3m)\Delta \geq \varepsilon$  do
  do
    TIGHTEN-LABEL( $f, \mu, \Delta$ );
     $D \leftarrow \{i \in V : e_i > (d_i + 1)\Delta\}$ ;
     $N_0 \leftarrow \{i \in V : e_i \leq d_i\Delta\}$ ;
    pick  $s \in D, t \in N$  connected by a tight path  $P$ ;
    send  $\Delta$  units of flow along  $P$ ;
  while  $D \neq \emptyset$ ;
  ADJUST( $\Delta$ );
   $\Delta \leftarrow \frac{\Delta}{2}$ ;
return  $\varepsilon$ -approximate optimal solution  $f$ .

```

Figure 2: The algorithm for symmetric concave generalized flows

Let us initialize $\mu_i = 1/M_i$ for every $i \in V$, and $f_{ij} = u_{ij}$ for every $ij \in E$. (We set the upper bounds rather than the lower bounds 0 because $\Gamma_{ij}(0) = -\infty$ is allowed.) Let us pick the initial value $\Delta = MU + 1$.

The algorithm consists of Δ -phases, and terminates with an ε -approximate solution if $(2n + 3m)\Delta < \varepsilon$. During the Δ -phase, we maintain a pseudoflow f and a Δ -conservative labeling μ . The μ_i values may only increase. Let D denote the set of nodes i with $e_i^\mu > (d_i + 1)\Delta$. The Δ -phase consists of iterations, and terminates whenever D becomes empty. In each iteration, we update μ to a canonical labeling by calling TIGHTEN-LABEL(f, μ, Δ). If $D \neq \emptyset$ still holds, send Δ units of relabeled flow on a tight path from some $s \in D$ to a Δ -neutral or Δ -negative node t .

6.5 Analysis

Claim 6.10. *The initial μ is Δ -conservative, and Δ -conservativity is maintained during the entire Δ -phase.* □

Proof. Initially, $f \equiv u$ and hence E_f is the set of backward arcs. For an arc $ij \in E$, $s_\Delta^\mu(ji) = u_{ij}/\mu_i \leq MU < \Delta$, and hence $E_f^\mu(\Delta) = \emptyset$. Also, $\mu_i = 1/M_i$ holds for every node i . TIGHTEN-LABEL(f, μ, Δ) clearly maintains Δ -conservativity. We use only tight arcs to send flow, and Lemma 6.7 guarantees that this preserves Δ -conservativity. At the end of the Δ -phase, ADJUST(Δ) transforms f to a $\Delta/2$ -conservative pseudoflow. □

Claim 6.11. *The Δ -phase starts with $Ex_{\Delta}^{\mu}(f) \leq (2n + 3m)\Delta$.*

Proof. For the initial solution, $Ex_{\Delta}^{\mu}(f) \leq M(\sum_{i \in V} |b_i| + mU) \leq (m + n)MU$. The claim follows, since $\Delta = MU + 1$. Once we finish all iterations in the Δ -phase, $D = \emptyset$ implies $Ex_{\Delta}^{\mu}(f) \leq n\Delta$. By Lemma 6.8, ADJUST(Δ) transforms f to a $\Delta/2$ -conservative solution by increasing the excess by at most $\frac{3}{2}m\Delta$. Hence the $\Delta/2$ phase starts with $Ex_{\Delta}^{\mu}(f) \leq (2n + 3m)\Delta/2$, proving the claim. \square

Lemma 6.12. *A Δ -phase consists of at most $2n + 3m$ iterations.*

Proof. As for the linear case, let $\Psi(i) = \lfloor e_i^{\mu} / \Delta \rfloor - d_i$ if $e_i^{\mu} > (d_i + 1)\Delta$ and $\Psi(i) = 0$ otherwise. Consider the potential function $\Psi = \sum_{i \in V} \Psi(i)$. By Claim 6.11, $\Psi \leq 2n + 3m$ holds at the beginning. In the relabeling steps, Ψ may only decrease, and in every path augmentation, it decreases by exactly 1. \square

Recall that $\kappa_f = \sum_{i \in V} M_i \kappa_i = \sum_{i \in V} M_i \min\{-e_i, 0\}$ denotes the excess discrepancy. For a Δ -conservative μ , $M_i \kappa_i = e_i^{\mu}$ holds for every node i with $e_i < 0$, because of $\mu_i = 1/M_i$. Consequently, κ_f is the total relabeled deficiency of the negative nodes. The next theorem shows that if $\Delta < \varepsilon/(2n + 3m)$, then we have an ε -optimal solution at the end of the Δ -phase.

Theorem 6.13. *At the end of phase Δ , the actual f is $(2n + 3m)\Delta$ -optimal.*

Proof. Let us keep running the algorithm forever unless it finds a 0-discrepancy solution at some phase. First, consider the case when for some $\Delta' = \Delta/2^k$, we terminate with a 0-discrepancy solution. In all phases between Δ and Δ' , the total decrease of excess discrepancy is bounded by $(2n + 3m)(\Delta/2 + \Delta/4 + \dots + \Delta/2^k) < (2n + 3m)\Delta$. Since in the Δ' -phase we have a 0-discrepancy solution, the total discrepancy at the end of the Δ -phase is at most $(2n + 3m)\Delta$, proving the theorem.

Assume now the procedure runs forever. For each $i \in V$, κ_i is decreasing and thus converges to some limit κ_i^* . Let $\kappa^* = \sum_{i \in V} M_i \kappa_i^*$. As above, the total decrease of the excess discrepancy after phase Δ is bounded by $(2n + 3m)\Delta$, hence $\kappa_f \leq \kappa^* + (2n + 3m)\Delta$. The proof finishes by constructing an optimal pseudoflow f^* with discrepancy κ^* .

Let $f^{(t)}$ denote the flow at time t , for $\Delta^{(t)} = \Delta_0/2^t$, with labels $\mu_i^{(t)}$. For each node i , $\mu_i^{(t)}$ is increasing; let $\mu_i^* = \lim_{t \rightarrow \infty} \mu_i^{(t)}$. For every $ij \in E$, $f_{ij}^{(t)}$ is a bounded sequence ($0 \leq f_{ij}^{(t)} \leq u_{ij}$). Consequently, we can choose an infinite set $T' \subseteq \mathbb{N}$ so that restricted to $t \in T'$, all sequences $f_{ij}^{(t)}$ converge; let f_{ij}^* denote the limits. We shall prove that f^* is an optimal pseudoflow with optimal labeling μ_i^* , completing the proof.

Let $V_{\infty} = \{i : \mu_i^* = \infty\}$. We claim that $V - V_{\infty} \neq \emptyset$. Indeed, if i is Δ -negative in a certain phase, then $\mu_i = 1/M_i$, and once i becomes Δ -positive or neutral, it would never again become Δ -negative. Consequently, the set of Δ -negative nodes is decreasing. Once it becomes empty, we arrive at a 0-discrepancy solution. If it never becomes empty, then we have a set N^* which remains the set of Δ -negative nodes after a finite number of steps and thus $\mu_i^* = 1/M_i$ for $i \in N^*$.

Let e_i^* denote the excesses of f^* . If $e_i^* < 0$, then clearly $i \in N^*$ and $\mu_i^* = 1/M_i$. If $e_i^* > 0$, we shall prove $\mu_i^* = \infty$. For a contradiction, assume $\mu_i^* < \infty$. Then for sufficiently large $t \in T'$, $(d_i + 2n + 3m)\Delta^{(t)}\mu_i^{(t)} < e_i^{(t)}$ and thus $Ex_{\Delta^{(t)}}^{\mu}(f) > (2n + 3m)\Delta^{(t)}$, a contradiction.

We have to prove $\Gamma_{ij}^{\mu_i^*+}(f_{ij}^{*\mu_i^*}) \leq 1$ whenever $ij \in E_{f^*}$. If $\mu_j^* = \infty$, then $\Gamma_{ij}^{\mu_i^*+}(f_{ij}^{*\mu_i^*}) = 0$. If $\mu_j^* < \infty$, then the definition (9) gives

$$1 \geq \theta_{\Delta^{(t)}}^{\mu_i^{(t)}}(ij) = \frac{\Delta^{(t)}\mu_j^{(t)}}{\Gamma_{ij}^{-1}(\Gamma_{ij}(f_{ij}^{(t)}) + \Delta^{(t)}\mu_j^{(t)}) - f_{ij}^{(t)}} \cdot \frac{\mu_i^{(t)}}{\mu_j^{(t)}}.$$

Then $\Delta^{(t)} \mu_j^{(t)} \rightarrow 0$ and hence the first fraction converges to $\Gamma_{ij}^+(f_{ij}^*) = 1/\{\Gamma_{ij}^{-1}\}^+(\Gamma_{ij}^+(f_{ij}^*))$, while the second to μ_i^*/μ_j^* , leading to the conclusion using Claim 6.5. \square

Theorem 6.14. *The above algorithm finds an ε -approximate solution to the symmetric concave generalized flow problem in $O(m(m+n \log n) \log(MUm/\varepsilon))$ oracle calls.*

Proof. The initial value of Δ is $MU + 1$, and we terminate if $\Delta < \varepsilon/(2n + 3m)$. Hence the total number of scaling phases is $O(\log(MUm/\varepsilon))$. The number of iterations in a phase is $O(m)$ by Lemma 6.12, and the running time of an iteration is dominated by TIGHTEN-LABEL, a slightly modified version of Dijkstra's algorithm that can be implemented in $O(m + n \log n)$ time using Fibonacci heaps as in [11] \square

7 Sink versions of the problems

In this section, we show how the algorithms in Sections 5 and 6 can be applied to solve the sink versions of the corresponding problems. For linear generalized flows, let us set $M_t = 1$ and $M_i = B^n + 1$ for every $i \in V - t$. Let us set $b_t = \lceil \sum_{j:jt \in E} \gamma_{jt} u_{jt} - \sum_{j:tj \in E} \ell_{tj} \rceil + 1 \leq d_t B^2 + 1$. This is a strict upper bound on $\sum_{j:jt \in E} \gamma_{jt} f_{jt} - \sum_{j:tj \in E} f_{tj}$, hence $e_t < 0$ will hold for any pseudoflow.

Let us run the algorithm for the symmetric formulation with these M_i 's, returning an optimal pseudoflow f and optimal labels μ . We claim that f is also optimal for the sink formulation. If $e_i \geq 0$ for all $i \neq t$, this is clearly the case.

On the other hand, we claim that if there exists any node i with $e_i < 0$, then the sink version is infeasible. First, we show that such an i cannot be reached from t on a path in E_f . Indeed, if P were a t - i path in E_f , then $1 \geq \gamma^\mu(P) = \gamma(P) \mu_t / \mu_i$. Since both i and t are negative, $\mu_i = 1/M_i = 1/(B^n + 1)$ and $\mu_t = 1$. Consequently, $\gamma(P) \leq 1/(B^n + 1)$. This is a contradiction since $\gamma(P)$ is the product of at most n rational numbers, each with denominator at most B . Let $V' \subseteq V$ be the set of nodes j for which there exists a path in E_f from j to some $i \in V - t$ with $e_i < 0$. This set verifies that the sink version cannot be feasible.

By setting the b_t value and the M_i 's, B has increased to $d_t B^2 + 1$ and $M = B^n + 1$. This gives running time $O(m^2(m + n \log n) \log B)$.

Let us turn to concave generalized flows. An ε -approximate solution to the sink version means a pseudoflow f with $\sum_{i \in V - t} \max\{0, -e_i\} \leq \varepsilon$ and e_t being at least the optimum value minus ε .

Let us set $b_t = U^* + 1$, a strict upper bound on $\sum_{j:jt \in E} \Gamma_{jt}(f_{jt}) - \sum_{j:tj \in E} f_{tj}$ (we defined U^* in Section 2.1). Thus $e_t < 0$ is always guaranteed. Let us set $M_i = \lceil (2U^* + 1)/\varepsilon \rceil + 1$ if $i \in V - t$ and $M_t = 1$. Let us run the algorithm for the symmetric formulation to obtain an ε -optimal solution f .

If $\kappa_f > 2U^* + 1 + \varepsilon$, then no feasible solution may exist. Indeed, by the definition of U^* , if there is a feasible solution f' , then there exists one with $e_t \geq -U^*$. If f' is such a feasible solution for the sink formulation, then its excess discrepancy for the symmetric formulation is at most $\kappa_{f'} \leq b_i + U^* \leq 2U^* + 1$, a contradiction as f was ε -optimal for the symmetric formulation.

If $\kappa_f \leq 2U^* + 1 + \varepsilon$, then

$$\sum_{i \in V - t} \max\{0, -e_i\} = \frac{1}{\lceil (2U^* + 1)/\varepsilon \rceil + 1} \sum_{i \in V - t} M_i \kappa_i \leq \frac{\kappa_f}{\lceil (2U^* + 1)/\varepsilon \rceil + 1} \leq \varepsilon.$$

Also κ_t cannot be further than ε from the optimum value of e_t for the sink formulation. Indeed, let f' be the optimal solution to the sink formulation with objective value e'_t . Then $\kappa_{f'} = b_t - e'_t$. The claim follows by

$$b_t - e'_t + \varepsilon = \kappa_{f'} + \varepsilon \geq \kappa_f \geq \kappa_t = b_t - e_t,$$

and thus $e_t \geq e'_t - \varepsilon$. In the first inequality, we use that f is ε -optimal for the sink formulation. This gives a running time bound $O(m(m + n \log n) \log(U^*m/\varepsilon))$.

8 Finding the optimal solution for rational convex programs

In this section, we first give a general theorem which shows how an approximate solution to the sink version can be converted to an exact optimal solution, given that one exists. We shall verify the required technical properties with appropriate parameters for nonsymmetric Arrow-Debreu Nash bargaining. Unlike the linear Fisher model, ADNB might be infeasible. However, it can be shown that if the problem is infeasible, then for appropriate (polynomially small) ε , the ε -approximate version is also infeasible. Similar reductions should hold for all other rational convex programs discussed in Section 3 as well, giving polynomial time algorithms for finding optimal solutions.

Theorem 8.1. *Let problem \mathcal{P} be given by the sink formulation with n nodes and m arcs, and complexity parameters U, U^* . Assume \mathcal{P} is guaranteed to have a rational optimal solution, and the following conditions hold for some values ε, T and a function $\tau(n, m, U^*)$.*

- (P1) *Consider the algorithm for the sink version for an ε -approximation. Then either there exists no feasible solution, or $\mu_i \leq T$ holds for any $i \in V$, even if running the algorithm for an arbitrary number of phases.*
- (P2) *A subroutine is provided for finding an optimal solution \tilde{f} in $\tau(n, m, U^*)$ time, if the following assumptions hold. Assume that for each $ij \in E$, we are given an interval $I_{ij} \subseteq [\ell_{ij}, u_{ij}]$ with $|I_{ij}| \leq 2T\varepsilon$, with the guarantee that there exists an optimal solution f^* with $f_{ij}^* \in I_{ij}$ for all $ij \in E$.*

*Then there exist an algorithm for finding the exact optimal solution or proving that the problem is infeasible in $O(m(m + n \log n) \log(U^*m/\varepsilon)) + \tau(n, m, U^*)$.*

We remark that in (P2), $\tilde{f} = f^*$ is not required.

Proof. Let us formulate the symmetric version for ε -approximation as in Section 7. Assume we run the algorithm for this problem forever, as in the proof of Theorem 6.13. The μ_j 's shall converge to some finite values $\mu_j^* \leq T$ as otherwise infeasibility is implied by (P1). In any Δ -phase, the total change of f_{ij}^t is bounded by $\varepsilon' = (2n + 3m)\Delta$, and thus f_{ij} may change by at most $T\varepsilon'$. Therefore all f_{ij} 's converge to some values f_{ij}^* , which can be seen to give an optimal solution, as in the proof of Theorem 6.13.

The algorithm terminates whenever $\Delta < \varepsilon/(2n + 3m)$. At this point, the intervals $I_{ij} = [f_{ij} - T\varepsilon, f_{ij} + T\varepsilon]$ satisfy the conditions in (P1), since $|f_{ij} - f_{ij}^*| \leq T\varepsilon$. Running the ε -approximation algorithm and then the algorithm in (P2) gives the running time bound. \square

To ensure property (P2), a useful method is to enforce the existence of a unique optimal solution by perturbing the input data, as done by Orlin [31] for linear Fisher markets. If there is a unique rational optimal solution f^* with all entries having denominator at most Q , then setting $2T\varepsilon < 1/Q$ enables us to identify the set of arcs with $f_{ij}^* > 0$. This can be already enough to compute f^* efficiently.

8.1 Application to nonsymmetric Arrow-Debreu Nash bargaining

Let us now apply Theorem 8.1 for the nonsymmetric ADNB problem. Let us assume all utilities U_{ij} , budgets m_i and disagreement utilities c_i are nonnegative integers, with $U_{\max} = \max\{U_{ij} : i \in B, j \in G\}$, $R = \max\{m_i : i \in B\}$, and $C = \max\{c_i : i \in B\}$. Let $n = |G| + |B|$ and let m be the number of pairs ij with $U_{ij} > 0$; in the concave generalized flow instance, the number of nodes is $n + 1$ and the number of arcs is $m + |B|$. Let us assume that there exists at least one arc with positive utility incident to any buyer and to any good. The special case $c \equiv 0$ is identical to Fisher's market with linear utilities.

Consider a candidate solution with price p_j for each good $j \in G$. Let $x_{ij} \geq 0$ denote the amount of good j purchased by buyer i . It follows from the KKT-conditions (see also [42]) that (p, x) is an optimal solution if and only if (i) $\sum_{i \in B} x_{ij} = 1$ for each good j , that is, each good is fully sold; and (ii) for any buyer i and good j , $U_{ij}/p_j \leq (\sum_{j \in G} U_{ij}x_{ij} - c_i)/m_i$, and equality holds whenever $x_{ij} > 0$.

By strict concavity of the objective, the utilities $\sum_{j \in G} U_{ij}x_{ij}$ accrued by the players are the same in any optimal solution whenever there is a feasible solution. Yet these same values can be obtained by different allocations. As in [31], we assume that there exist a unique optimal allocation as well. This can be done by a lexicographic perturbation of the U_{ij} values, without significantly increasing the running time. This guarantees that the set of arcs with $x_{ij} > 0$ in an optimal solution is cycle-free.

For the concave generalized flow instance, let us set upper capacity $v_{ji} = 2$ on each arc ji with $j \in G$, $i \in B$, and $v_{it} = 2\sum_{j \in G} U_{ij}$ on each arc it (we set larger capacities so that the capacity constraints would never become tight). Hence the complexity parameter U is bounded by $2|G|U_{\max}$. We shall prove the following.

Theorem 8.2. *Let $K = nRU_{\max}$. Setting $T = U^* = \max\{C, nK \log K\}$, $\varepsilon = 1/(2K^n U^*)$ satisfies the requirements on U^* in Section 2.1 and (P1) and (P2) in Theorem 8.1. Our algorithm delivers an optimal solution in running time $O(m(m + n \log n)(n \log(nRU_{\max}) + \log C))$ for nonsymmetric ADNB.*

The running time $\tau(n, m, U^*)$ will be negligible compared to the main algorithm and therefore it does not affect the complexity bound. In the rest of the section, we shall verify the above choices.

Lemma 8.3 (see [42, Thm 2], [31, Lemma 2.1]). *Assuming that the problem is feasible and there exists a unique optimal allocation x^* , all positive x_{ij}^* values are rational numbers with a common denominator $S \leq K^n$.*

Proof. The optimal allocations x^* and prices p^* can be uniquely obtained given the set F of arcs ji with $x_{ij}^* > 0$. If we introduce the variable $q_j = 1/p_j$, then an optimal solution must satisfy the following system of linear equations.

$$\begin{aligned} \sum_{j:ji \in E} x_{ij} &= 1 & \forall i \in G & \quad (15) \\ \sum_{k:ki \in E} U_{ik}x_{ik} - U_{ij}m_iq_j &= c_i & \forall ji \in F & \end{aligned}$$

We claim that this system has a unique solution (x^*, p^*) . Fixing the price of one good j_0 as $p_{j_0} = \alpha$ in a component of F , it uniquely determines p_j for any good j in the same component as α times the product of the U_{ab} values on the unique path from j_0 to j . Similarly, this determines the best bang-per-buck values $b_i = U_{ij}/p_j = (\sum_{k:ki \in E} U_{ik}x_{ik} - c_i)/m_i$ for $ji \in F$, which are proportional

to $1/\alpha$. The optimality conditions imply that in an equilibrium, the money spent by buyer i is $r_i = m_i + c_i/b_i$. In each component of F , the sum of prices should be equal to the money spent by the buyers. This uniquely determines all prices and bang-per-buck values in the component. The x_{ij} values in the component have to sum up to 1 for each good i and $\sum_{k:ki \in E} U_{ik}x_{ik} = b_i m_i + c_i$. As F is a forest, this has a unique solution.

In the solution to (15), a common denominator is the determinant S of a largest non-singular submatrix of the constraint matrix. The Hadamard-bound gives $S \leq (nRU_{\max})^n = K^n$, using that $|F| \leq n - 1$. \square

The above proof also gives a simple linear time algorithm for finding the optimal solution, verifying (P2) if $2T\varepsilon < 1/K^n$, with $\tau(n, m, U^*)$ being negligible compared to the running time of the approximation algorithm.

Next we justify the choice of U^* . $U \leq U^*$ clearly holds. For an arbitrary pseudoflow f , $e_t \leq \sum_{i \in G} m_i \log v_{it} \leq nR \log(nU) \leq U^*$. It is left to show that if the problem is feasible, there exists a feasible solution with $e_t > -U^*$. Since the U_{ij} and c_i values are integers, whenever $\sum_{j:ji \in E} U_{ij}x_{ij} - c_i > 0$, it should be at least $1/K^n$. Consequently, if the sink version of the problem is feasible, the optimal objective value is at least $e_t \geq \sum_{i \in G} m_i \log(1/K^n) \geq -n^2 R \log K \geq -U^*$.

The next claim verifies (P1) and thus completes the proof of Theorem 8.2.

Lemma 8.4. *Either the problem is not feasible, or $\mu_k \leq U^*$ holds for any $k \in B \cup G$ in arbitrary Δ -phase.*

Proof. Recall from Section 7 that we solve the sink version by reducing it to the symmetric algorithm with $M_t = 1$ and $M_i = \lceil 2U^*/\varepsilon \rceil + 1$. Since μ_i is non-decreasing, these values converge to some limits $\mu_i^* \in \mathbb{R} \cup \{\infty\}$. We have $0 \leq f_{ij} \leq u_{ij}$ on all arcs $ij \in E$ in every phase, and therefore we can choose an infinite subset $T' \subseteq \mathbb{N}$ so that all f_{ij} 's converge if we restrict ourselves to iteration numbers in T' . As in Theorem 6.13, it can be easily verified that the limit f^* is an optimal solution to the symmetric version with conservative labeling μ^* .

As in Section 7, if $\kappa_f > 2U^* + 1$, then the sink version is not feasible, and otherwise f^* is also optimal to the sink version. In the latter case, $f_{it}^* > 0$ for arbitrary $i \in B$ must hold as otherwise $e_t = -\infty$ in the sink version gives infeasibility. Recall also that the symmetric version was defined such that $e_t < 0$ holds for every feasible solution and therefore $\mu_t^* = 1$.

Consider now an arbitrary $i \in B$. Both $it, ti \in E_f$ (it is easy to verify that $f_{it} = v_{it}$ is impossible), and therefore $\frac{\mu_i^* m_i}{\mu_t^* f_{it}^*} = 1$ must hold by the conservativity of μ^* . This means $\mu_i^* = f_{it}^*/m_i \leq U \leq U^*$.

Finally, let $j \in G$. Then for an arbitrary arc $ji \in E$, $ji \in E_f$ easily follows, and therefore conservativity gives $\frac{\mu_j^*}{\mu_i^*} U_{ji} \leq 1$, which implies $\mu_j^* \leq U^*$. \square

Finally, we remark that if we apply this algorithm to linear Fisher markets ($c \equiv 0$), the algorithm runs in a fundamentally different way as [7] or [31]. While both these algorithms increase the prices, our algorithm works the other way around: it starts with the highest possible prices, and decreases them.

9 Discussion

We have given the first polynomial time combinatorial algorithms for both the symmetric and the sink formulation of the concave generalized flow problem. Our algorithm is not strongly polynomial. In fact, no such algorithm is known already for the linear case: it is a fundamental open question to find a strongly polynomial algorithm for linear generalized flows. If resolved, a natural question

could be to devise a strongly polynomial algorithm for some class of convex generalized flow problems, analogously to the recent result [43], desirably including the market and Nash bargaining applications.

Linear Fisher market is also captured by [43]. A natural question is if there is any direct connection between our model and the convex minimum cost flow model studied in [43]. Despite certain similarities, no reduction is known in any direction. Indeed, no such reduction is known even between the linear special cases, that is, generalized flows and minimum-cost circulations. In fact, the only known market setting captured by both is linear Fisher. Perfect price discrimination and ADNB are not known to be reducible to flows with convex objective. In contrast, spending constraint utilities [41] are not known to be captured by our model, although they are captured by the other.

As discussed in Section 4.4, it seems difficult to extend any generalized flow algorithm having separate cycle cancelling and flow transportation subroutines. While this includes the majority of combinatorial algorithms, there are some exceptions. Goldberg, Plotkin and Tardos [13] gave two different algorithms: besides FAT-PATH, they also presented MCF, an algorithm that uses a minimum-cost circulation algorithm directly as a subroutine. Hence for the concave setting, it could be possible to develop a similar algorithm using a minimum concave cost circulation algorithm, for example [19] or [25] as a black box.

Another approach that avoids scaling is [44] for minimum-cost generalized flows and [34] for generalized flows: these algorithms can be seen as extensions of the cycle cancelling method, extending minimum mean cycles to GAP's in a certain sense. While it does not seem easy, it might be possible to develop such an algorithm for concave generalized flows as well.

In defining an ε -approximate solution for the sink version of concave generalized flows, we allow two types of errors, both for the objective and for feasibility. A natural question is if either of these could be avoided. While the value oracle model as we use it, seems to need feasibility error, it might be possible to avoid it using a stronger oracle model as in [25]. One might also require a feasible solution as part of the input, as a starting point to maintain feasibility (For example if all lower bounds and node demands are 0 and $\Gamma_{ij}(0) = 0$ on all arcs ij , then $f \equiv 0$ is always feasible).

Acknowledgments

The author is grateful to Vijay Vazirani for several fruitful discussions on market equilibrium problems.

References

- [1] D. P. Ahlfeld, J. M. Mulvey, R. S. Dembo, and S. A. Zenios. Nonlinear programming on generalized networks. *ACM Transactions on Mathematical Software (TOMS)*, 13(4):350–367, 1987.
- [2] R. K. Ahuja, T. L. Magnanti, and J. B. Orlin. *Network Flows: Theory, Algorithms, and Applications*. Prentice-Hall, Inc., feb 1993.
- [3] D. P. Bertsekas, L. C. Polymenakos, and P. Tseng. An ε -relaxation method for separable convex cost network flow problems. *SIAM Journal on Optimization*, 7(3):853–870, 1997.
- [4] E. Cohen and N. Megiddo. New algorithms for generalized network flows. *Mathematical Programming*, 64(1):325–336, 1994.

- [5] G. B. Dantzig. *Linear Programming and Extensions*. Princeton University Press, Princeton, 1963.
- [6] N. Devanur and K. Jain. Online matching with concave returns. In *Proceedings of the 44th Symposium on Theory of Computing (STOC)*, pages 137–144. ACM, 2012.
- [7] N. R. Devanur, C. H. Papadimitriou, A. Saberi, and V. V. Vazirani. Market equilibrium via a primal–dual algorithm for a convex program. *Journal of the ACM (JACM)*, 55(5):22, 2008.
- [8] J. Edmonds and R. M. Karp. Theoretical improvements in algorithmic efficiency for network flow problems. *Journal of the ACM (JACM)*, 19(2):248–264, 1972.
- [9] E. Eisenberg and D. Gale. Consensus of subjective probabilities: The pari-mutuel method. *The Annals of Mathematical Statistics*, 30(1):165–168, 1959.
- [10] L. K. Fleischer and K. D. Wayne. Fast and simple approximation schemes for generalized flow. *Mathematical Programming*, 91(2):215–238, 2002.
- [11] M. L. Fredman and R. E. Tarjan. Fibonacci heaps and their uses in improved network optimization algorithms. *Journal of the ACM (JACM)*, 34(3):596–615, 1987.
- [12] G. Goel and V. V. Vazirani. A perfect price discrimination market model with production, and a (rational) convex program for it. *Mathematics of Operations Research*, 36:762–782, 2011.
- [13] A. V. Goldberg, S. A. Plotkin, and É. Tardos. Combinatorial algorithms for the generalized circulation problem. *Mathematics of Operations Research*, 16(2):351, 1991.
- [14] A. V. Goldberg and R. E. Tarjan. Finding minimum-cost circulations by canceling negative cycles. *Journal of the ACM (JACM)*, 36(4):873–886, 1989.
- [15] D. Goldfarb and Z. Jin. A faster combinatorial algorithm for the generalized circulation problem. *Mathematics of Operations Research*, 21(3):529–539, 1996.
- [16] D. Goldfarb, Z. Jin, and Y. Lin. A polynomial dual simplex algorithm for the generalized circulation problem. *Mathematical Programming*, 91(2):271–288, 2002.
- [17] D. Goldfarb, Z. Jin, and J. B. Orlin. Polynomial-time highest-gain augmenting path algorithms for the generalized circulation problem. *Mathematics of Operations Research*, 22(4):793–802, 1997.
- [18] D. Goldfarb and Y. Lin. Combinatorial interior point methods for generalized network flow problems. *Mathematical Programming*, 93(2):227–246, 2002.
- [19] D. S. Hochbaum and J. G. Shanthikumar. Convex separable optimization is not much harder than linear optimization. *Journal of the ACM (JACM)*, 37(4):843–862, 1990.
- [20] K. Jain and V. V. Vazirani. Eisenberg-gale markets: Algorithms and game-theoretic properties. *Games and Economic Behavior*, 70(1):84–106, 2010.
- [21] W. S. Jewell. Optimal flow through networks. *Operations Research*, 10:476–499, 1962.
- [22] E. Kalai. Nonsymmetric Nash solutions and replications of 2-person bargaining. *International Journal of Game Theory*, 6(3):129–133, 1977.

- [23] L. V. Kantorovich. Mathematical methods of organizing and planning production. *Publication House of the Leningrad State University*, page 68, 1939. English translation in *Management Science* 6(4):366-422, 1960.
- [24] S. Kapoor and P. M. Vaidya. Speeding up Karmarkar’s algorithm for multicommodity flows. *Mathematical Programming*, 73(1):111–127, 1996.
- [25] A. V. Karzanov and S. T. McCormick. Polynomial methods for separable convex optimization in unimodular linear spaces with applications. *SIAM J. Comput.*, 26(4):1245–1275, 1997.
- [26] M. Minoux. Solving integer minimum cost flows with separable convex cost objective polynomially. *Mathematical Programming Study*, 25:237, 1985.
- [27] N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press New York, NY, USA, 2007.
- [28] K. Onaga. Dynamic programming of optimum flows in lossy communication nets. *IEEE Transactions on Circuit Theory*, 13:308–327, 1966.
- [29] K. Onaga. Optimum flows in general communication networks. *Journal of the Franklin Institute*, 283(4):308–327, 1967.
- [30] J. B. Orlin. A faster strongly polynomial minimum cost flow algorithm. *Operations Research*, 41(2):338–350, 1993.
- [31] J. B. Orlin. Improved algorithms for computing fisher’s market clearing prices. In *Proceedings of the 42nd ACM Symposium on Theory of Computing (STOC)*, pages 291–300. ACM, 2010.
- [32] T. Radzik. Approximate generalized circulation. *Technical Report93-2, Cornell Computational Optimization Project, Cornell University*, 1993.
- [33] T. Radzik. Improving time bounds on maximum generalised flow computations by contracting the network. *Theoretical Computer Science*, 312(1):75–97, 2004.
- [34] M. Restrepo and D. P. Williamson. A simple GAP-canceling algorithm for the generalized maximum flow problem. *Mathematical Programming*, 118(1):47–74, 2009.
- [35] M. Shigeno. A survey of combinatorial maximum flow algorithms on a network with gains. *Journal of the Operations Research Society of Japan*, 47:244–264, 2004.
- [36] M. Shigeno. Maximum network flows with concave gains. *Mathematical Programming*, 107(3):439–459, 2006.
- [37] V. I. Shmyrev. An algorithm for finding equilibrium in the linear exchange model with fixed budgets. *Journal of Applied and Industrial Mathematics*, 3(4):505–518, 2009.
- [38] É. Tardos and K. D. Wayne. Simple maximum flow algorithms in lossy networks. In *Proceedings of IPCO, Lecture Notes in Computer Science*, volume 1412, pages 310–324, 1998.
- [39] K. Truemper. On max flows with gains and pure min-cost flows. *SIAM Journal on Applied Mathematics*, 32(2):450–456, 1977.
- [40] K. Truemper. Optimal flows in nonlinear gain networks. *Networks*, 8(1):17–36, 1978.

- [41] V. V. Vazirani. Spending constraint utilities with applications to the Adwords market. *Mathematics of Operations Research*, 35(2):458–478, 2010.
- [42] V. V. Vazirani. The notion of a rational convex program, and an algorithm for the Arrow-Debreu Nash bargaining game. *Journal of ACM (JACM)*, 59(2), 2012.
- [43] L. A. Végh. Strongly polynomial algorithm for a class of minimum-cost flow problems with separable convex objectives. In *Proceedings of STOC*, pages 27–40. ACM, 2012.
- [44] K. D. Wayne. A polynomial combinatorial algorithm for generalized minimum cost flow. *Mathematics of Operations Research*, pages 445–459, 2002.