

Negative-Weight Shortest Paths and Unit Capacity Minimum Cost Flow in $\tilde{O}(m^{10/7} \log W)$ Time

Michael B. Cohen	Aleksander Mądry	Piotr Sankowski	Adrian Vladu
MIT	MIT	University of Warsaw	MIT
micohen@mit.edu	madry@mit.edu	sank@mimuw.edu.pl	avladu@mit.edu

July 13, 2016

Abstract

In this paper, we study a set of combinatorial optimization problems on weighted graphs: the shortest path problem with negative weights, the weighted perfect bipartite matching problem, the unit-capacity minimum-cost maximum flow problem and the weighted perfect bipartite b -matching problem under the assumption that $\|b\|_1 = O(m)$. We show that each one of these four problems can be solved in $\tilde{O}(m^{10/7} \log W)$ time, where W is the absolute maximum weight of an edge in the graph, which gives the first in over 25 years polynomial improvement in their sparse-graph time complexity.

At a high level, our algorithms build on the interior-point method-based framework developed by Mądry (FOCS 2013) for solving unit-capacity maximum flow problem. We develop a refined way to analyze this framework, as well as provide new variants of the underlying preconditioning and perturbation techniques. Consequently, we are able to extend the whole interior-point method-based approach to make it applicable in the weighted graph regime.

1 Introduction

In 2013, Mądry [24] put forth an algorithm for the maximum flow and maximum-cardinality bipartite matching problems that improved over a long standing $O(n^{3/2})$ running time barrier for sparse graphs. Specifically, he presented an $\tilde{O}(m^{10/7})$ time algorithm for computing maximum flow in a unit capacity network – which implies an $\tilde{O}(m^{10/7})$ time algorithm for the maximum-cardinality bipartite matching problem as well. The core of his approach is a new path-following interior-point type algorithm for solving a certain “bipartite b -matching” problem. At a high level, this algorithm uses electrical flow computations to improve the maintained primal dual solution and move it along the so-called central path, i.e., a set of primal dual solutions in which every edge contributes approximately the same amount to the duality gap. As Mądry has shown, one can use this framework to establish an $O(m^{3/7})$ bound on the number of such electrical flow computations needed to compute a (near-) optimal solution to the considered problem and thus to improve upon the generic worst-case $O(\sqrt{m})$ bound that all the previous interior-point-method-based algorithms provided. The key ingredient needed to obtaining this improved bound was a technique for perturbing and preconditioning the intermediate solutions that emerge during the computations. Unfortunately, the technique [24] used in was inherently unable to cope with large capacity or *any* weights on edges. In fact, it did not provide any meaningful result for the unit-capacity minimum cost maximum flow problem even when all the edge weights were equal to 1. Consequently, it remains an open question whether a similar running time improvement can be achieved for either: (a) non-unit capacity networks; or (b) weighted variants of the unit-capacity graph problems.

1.1 Our Contribution

In this paper, we answer the second question above affirmatively by providing an $\tilde{O}(m^{10/7} \log W)$ time algorithm for the minimum cost unit-capacity maximum flow problem. In addition to the improvement for this fundamental graph problem, this result also improves several other standard problems as immediate corollaries. Namely, by well-known reductions, it implies $\tilde{O}(m^{10/7} \log W)$ time algorithms for the minimum-weight bipartite perfect matching problem, the minimum-weight bipartite b -matching problem, and, the shortest path problem for graphs with negative weights. This constitutes the first in more than 25 years polynomial improvement of sparse graph time complexity for each one of these problems.

To obtain these results we simplify and extend the framework from [24] by developing new preconditioning and perturbation techniques. These techniques provide us with much better control of the intermediate solutions that we precondition/perturb. In particular, in stark contrast to [24], the preconditioning and perturbation steps do not lead to any changes in edge costs. Also, the resulting changes in vertex demands are very minimal. Thus, there is no more need for repeated fixing of these vertex demand changes throughout the execution of the algorithm – a single demand correction step is performed only at the very end. Finally, our analysis of the resulting algorithm is much more streamlined and principled compared to the analysis performed in [24], providing us with much tighter grip of the trade-offs underlying the whole framework.

1.2 Previous Work

The minimum-cost flow, min-weight bipartite perfect matching as well as the shortest path with negative weights problems are core combinatorial optimization tasks that now have been studied for over 85 years, starting with the work of Egerváry [7] from 1931. Due to immense number of works on these topics we will not review them. Instead we will concentrate only on the ones that are relevant to the sparse graph case, as that is the regime where our results are of main importance.

	Complexity	Author
	$O(n^4)$	Shimbel (1955) [30]
	$O(Wn^2m)$	Ford (1956) [14]
*	$O(nm)$	Bellman (1958) [1], Moore (1959) [25]
	$O(n^{\frac{3}{4}}m \log W)$	Gabow (1983) [9]
	$O(\sqrt{nm} \log(nW))$	Gabow and Tarjan (1989) [10]
*	$O(\sqrt{nm} \log(W))$	Goldberg (1993) [12]
*	$\tilde{O}(Wn^\omega)$	Sankowski (2005) [27] Yuster and Zwick (2005) [35]
*	$\tilde{O}(m^{10/7} \log W)$	this paper

Table 1: The complexity results for the SSSP problem with negative weights (* indicates asymptotically the best bound for some range of parameters).

Shortest Paths with Negative Weights A list of the complexity results on single source shortest paths with negative weights is included in Table 1. Observe that the sparse case can either be solved in $O(mn)$ time [1, 25] or $\tilde{O}(\sqrt{nm} \log W)$ time [10, 12]. The only progress that we had since these papers was the reduction of the problem to fast matrix multiplication [27, 35] that is relevant only for dense graphs with small integral weights.

	Complexity	Author
	$O(Wn^2m)$	Egerváry (1931) [7]
	$O(n^4)$	Khun (1955) [20] and Munkers (1957) [26]
	$O(n^2m)$	Iri (1960) [13]
	$O(n^3)$	Dinic and Kronrod (1969) [5]
*	$O(nm + n^2 \log n)$	Edmonds and Karp (1970) [6]
	$O(n^{\frac{3}{4}}m \log W)$	Gabow (1983) [9]
*	$O(\sqrt{nm} \log(nW))$	Gabow and Tarjan (1989) [10]
*	$O(W\sqrt{nm})$	Kao, Lam, Sung and Ting (1999) [16]
*	$O(Wn^\omega)$	Sankowski (2006) [28]
*	$\tilde{O}(m^{10/7} \log W)$	this paper

Table 2: The complexity results for the minimum weight bipartite perfect matching problem (* indicates asymptotically the best bound for some range of parameters).

Min-cost Perfect Bipartite Matching The complexity survey of for the minimum weight bipartite perfect matching problem is given in Table 2. Here, the situation is very similar to the case of shortest paths. We have two results that are relevant for the sparse case considered here: $O(nm + n^2 \log n)$ time [6] and $O(\sqrt{nm} \log(nW))$ time [10]. Again the only polynomial improvement that was achieved during the last 25 years is relevant to the dense case only [28].

Complexity	Author
$O(m(m + n \log n))$	Edmonds and Karp (1972) [6]
$O(n^{5/3}m^{2/3} \log(nW))$	Goldberg and Tarjan (1987) [11]
$O(\min(\sqrt{m}, n^{2/3})m \log(nW))$	Gabow and Tarjan (1989) [10]
$\tilde{O}(m^{3/2})$	Daitch and Spielman (2008) [4]
$\tilde{O}(\sqrt{nm})$	Lee and Sidford (2014) [23]
$\tilde{O}(m^{10/7} \log W)$	this paper

Table 3: The summary of the results for the min-cost unit-capacity max-flow problem. For simplicity we only list exact algorithms that yielded polynomial improvement or new strongly polynomial bounds. For the full list of complexities we refer the reader to Chapter 12 in [29].

Minimum-cost Unit-capacity Maximum Flow Problem Due to the vast number of results for this topic we restricted ourselves to present in Table 3 only algorithms for the unit-capacity case that yielded significant improvement. We note, however, that handling general capacities for this problem is a major challenge and our restriction to unit-capacity networks oversimplifies the history of this problem. Nevertheless, for there sparse case there are two relevant complexities: a $\tilde{O}(\sqrt{nm})$ time bound of [23] that, for the sparse graph case, matches the previously best known bound $O(m^{3/2} \log(nW))$ due to [10]. We note that there was also a limited progress [8] on fast matrix multiplication based algorithms for the small vertex capacity variant of this problem.

Complexity	Author
$O(m(m + n \log n))$	Edmonds and Karp (1972) [6]
$O(m^{7/4} \log W)$	Gabow (1985) [9]
$O(m^{3/2} \log(nW))$	Gabow and Tarjan (1989) [10]
$\tilde{O}(m^{10/7} \log W)$	this paper

Table 4: The summary of the results for the min-cost perfect bipartite b -matching problem under assumption that $b(V) = O(m)$. For simplicity we only list exact algorithms that yielded polynomial improvement or new strongly polynomial bounds. For the full list of complexities we refer the reader to Chapter 21 in [29].

Minimum-cost Perfect Bipartite b -Matching Minimum-cost perfect bipartite b -matching problem bears many similarities to the minimum-cost maximum flow problem, e.g., see our reduction of min-cost flow to b -matchings in Section 3. Hence some of the complexities in Table 4 are the same as in Table 3. However, not all results seem to carry over as exemplified in the tables. b -matchings seems slightly harder than max-flow as one needs to route exact amount of flow between many sources and sinks. The results relevant for sparse case are: weakly polynomial $O(m^{3/2} \log(nW))$ time algorithm [10] or strongly polynomial $O(m(m + n \log n))$ time algorithm [6].

1.3 The Outline of Our Algorithm

As mentioned above, our focus is on development of a faster, $\tilde{O}(m^{10/7} \log W)$ -time algorithm for the minimum-cost unit-capacity maximum flow problem, since well-known reductions immediately yield $\tilde{O}(m^{10/7} \log W)$ -time algorithms for the remaining problems (see also Section 8). In broad outline, our approach to solving that flow problem follows the framework introduced by Mądry [24] and comprises three major components.

First, in Section 3, we reduce our input minimum-cost flow instance to an instance of the bipartite minimum-cost b -matching problem. The latter instance will have a special structure. In particular, it can be viewed as a minimum-cost flow problem instance that has general flow demands but no capacities.

Then, in Section 4, we put forth a basic interior-point method framework that builds on the framework of Mądry [24] for solving this kind of uncapacitated minimum-cost flow instances. This basic framework alone will not be capable of breaking the $O(\sqrt{m})$ iteration bound. Therefore, in Sections 5 and 6, we develop a careful perturbation and preconditioning technique to help us control the convergence behavior of our interior-point method framework. An important feature of this technique is that, in contrast to the technique used by [24], our perturbations do not alter the arc costs. Thus they are suitable for dealing with weighted problems. We then prove that the resulting algorithms indeed obtain the improved running time bound of $\tilde{O}(m^{10/7} \log W)$ but the (near-)optimal flow solution it outputs might have some of the flow demands changed.

Finally, in Section 7, we address this problem by developing a fast procedure that relies on classic combinatorial techniques and recovers from this perturbed (near-)optimal solution an optimal solution to our original minimum-cost flow instance.

2 Preliminaries

In this section, we introduce some basic notation and definitions that we will need later.

2.1 Minimum Cost σ -Flow

We denote by $G = (V, E, c)$ a directed graph with vertex set V , arc set E and cost function c . We denote by $m = |E|$ the number of arcs in G , and by $n = |V|$ its number of vertices. An arc e of G connects an

ordered pair (w, v) , where w is called *tail* and v is called *head*. We will be mostly working with σ -flows in G , where $\sigma \in \mathbb{R}^n$, satisfying $\sum_v \sigma_v = 0$, is the *demand vector*. A σ -flow in G is defined to be a vector $f \in \mathbb{R}^m$ (assigning values to arcs of G) that satisfies the following *flow conservation constraints*:

$$\sum_{e \in E^+(v)} f_e - \sum_{e \in E^-(v)} f_e = \sigma_v, \quad \text{for each vertex } v \in V. \quad (2.1)$$

We denote by $E^+(v)$ (resp. $E^-(v)$) the set of arcs of G that are leaving (resp. entering) vertex v . The above constraints enforce that, for every $v \in V$, the total out-flow from v minus the total in-flow (i.e. the net flow) out of v is equal to σ_v .

Furthermore, we say that a σ -flow f is *feasible* in G iff it satisfies *non-negativity and capacity constraints* (in this paper we are only concerned with unit capacities):

$$0 \leq f_e \leq 1, \quad \text{for each arc } e \in E. \quad (2.2)$$

For our interior point method the basic object is the *minimum cost σ -flow problem* consists of finding feasible σ -flow f that minimizes *the cost of the flow* $c(f) = \sum_e c_e f_e$.

2.2 Electrical Flows and Potentials

A notion that will play a fundamental role in this paper is the notion of electrical flows. Here, we just briefly review some of the key properties that we will need later. For an in-depth treatment we refer the reader to [2].

Consider an undirected graph G and a vector of resistances $r \in \mathbb{R}^m$ that assigns to each edge e its *resistance* $r_e > 0$. For a given σ -flow f in G , let us define its *energy* (with respect to r) $\mathcal{E}_r(f)$ to be

$$\mathcal{E}_r(f) := \sum_e r_e f_e^2 = f^\top R f, \quad (2.3)$$

where $R = \text{diag}(r)$ is an $m \times m$ diagonal matrix with $R_{e,e} = r_e$, for each edge e . In order to simplify notation, we drop the subscript or the parameter whenever it is clear from the context.

For a given undirected graph G , a demand vector σ , and a vector of resistances r , we define an *electrical σ -flow* in G (that is *determined* by the resistances r) to be the σ -flow that minimizes the energy $\mathcal{E}_r(f)$ among all σ -flows in G . As energy is a strictly convex function, one can easily see that such a flow is unique. Also, we emphasize that we do *not* require here that this flow is feasible with respect to the (unit) capacities of G (cf. (2.2)). Furthermore, whenever we consider electrical flows in the context of a directed graph G , we will mean an electrical flow – as defined above – in the (undirected) projection \bar{G} of G .

One of very useful properties of electrical flows is that it can be characterized in terms of vertex potentials inducing it. Namely, one can show that a σ -flow f in G is an electrical σ -flow determined by resistances r iff there exist *vertex potentials* ϕ_v (that we collect into a vector $\phi \in \mathbb{R}^n$) such that, for any edge $e = (u, v)$ in G that is oriented from u to v ,

$$f_e = \frac{\phi_v - \phi_u}{r_e}. \quad (2.4)$$

In other words, a σ -flow f is an electrical σ -flow iff it is *induced* via (2.4) by some vertex potentials ϕ . (Note that orientation of edges matters in this definition.)

Using vertex potentials, we are able to express the energy $\mathcal{E}_r(f)$ (see (2.3)) of an electrical σ -flow f in terms of the potentials ϕ inducing it as

$$\mathcal{E}_r(f) = \sum_{e=(u,v)} \frac{(\phi_v - \phi_u)^2}{r_e}. \quad (2.5)$$

2.3 Laplacian Solvers

A very important algorithmic property of electrical flows is that one can compute very good approximations of them in nearly-linear time. Below, we briefly describe the tools enabling that.

To this end, let us recall that electrical σ -flow is the (unique) σ -flow induced by vertex potentials via (2.4). So, finding such a flow boils down to computing the corresponding vertex potentials ϕ . It turns out that computing these potentials can be cast as a task of solving certain type of linear system called *Laplacian* systems. To see that, let us define the *edge-vertex incidence matrix* B being an $n \times m$ matrix with rows indexed by vertices and columns indexed by edges such that

$$B_{v,e} = \begin{cases} 1 & \text{if } e \in E^+(v), \\ -1 & \text{if } e \in E^-(v), \\ 0 & \text{otherwise.} \end{cases}$$

Now, we can compactly express the flow conservation constraints (2.1) of a σ -flow f (that we view as a vector in \mathbb{R}^m) as

$$Bf = \sigma.$$

On the other hand, if ϕ are some vertex potentials, the corresponding flow f induced by ϕ via (2.4) (with respect to resistances r) can be written as

$$f = R^{-1}B^T\phi,$$

where again R is a diagonal $m \times m$ matrix with $R_{e,e} := r_e$, for each edge e .

Putting the two above equations together, we get that the vertex potentials ϕ that induce the electrical σ -flow determined by resistances r are given by a solution to the following linear system

$$BR^{-1}B^T\phi = L\phi = \sigma, \tag{2.6}$$

where $L := BR^{-1}B^T$ is the (weighted) *Laplacian* L of G (with respect to the resistances r). One can easily check that L is a symmetric $n \times n$ matrix indexed by vertices of G with entries given by

$$L_{u,v} = \begin{cases} \sum_{e \in E(v)} 1/r_e & \text{if } u = v, \\ -1/r_e & \text{if } e = (u, v) \in E, \text{ and} \\ 0 & \text{otherwise.} \end{cases} \tag{2.7}$$

One can see that the Laplacian L is not invertible, but – as long as, the underlying graph is connected – its null-space is one-dimensional and spanned by the all-ones vector. As we require our demand vectors σ to have its entries sum up to zero (otherwise, no σ -flow can exist), this means that they are always orthogonal to that null-space. Therefore, the linear system (2.6) has always a solution ϕ and one of these solutions¹ is given by

$$\phi = L^+\sigma,$$

where L^+ is the Moore-Penrose pseudo-inverse of L .

Now, from the algorithmic point of view, the crucial property of the Laplacian L is that it is symmetric and *diagonally dominant*, i.e., for any $v \in V$, $\sum_{u \neq v} |L_{u,v}| \leq L_{v,v}$. This enables us to use fast approximate solvers for symmetric and diagonally dominant linear systems to compute an approximate electrical σ -flow. Namely, there is a long line of work [33, 18, 19, 17, 3, 21, 22] that builds on an earlier work of Vaidya [34] and Spielman and Teng [32], that designed an SDD linear system solver that implies the following theorem.

¹Note that the linear system (2.6) will have many solutions, but any two of them are equivalent up to a translation. So, as the formula (2.4) is translation-invariant, each of these solutions will yield the same unique electrical σ -flow.

Theorem 1. For any $\epsilon > 0$, any graph G with n vertices and m edges, any demand vector σ , and any resistances r , one can compute in $\tilde{O}(m \log m \log \epsilon^{-1})$ time vertex potentials $\tilde{\phi}$ such that $\|\tilde{\phi} - \phi^*\|_L \leq \epsilon \|\phi^*\|_L$, where L is the Laplacian of G , ϕ^* are potentials inducing the electrical σ -flow determined by resistances r , and $\|\phi\|_L := \sqrt{\phi^\top L \phi}$.

To understand the type of approximation offered by the above theorem, observe that $\|\phi\|_L^2 = \phi^\top L \phi$ is just the energy of the flow induced by vertex potentials ϕ . Therefore, $\|\tilde{\phi} - \phi^*\|_L$ is the energy of the electrical flow \tilde{f} that “corrects” the vertex demands of the electrical $\tilde{\sigma}$ -flow induced by potentials $\tilde{\phi}$, to the ones that are dictated by σ . So, in other words, the above theorem tells us that we can quickly find an electrical $\tilde{\sigma}$ -flow \tilde{f} in G such that $\tilde{\sigma}$ is a slightly perturbed version of σ and \tilde{f} can be corrected to the electrical σ -flow f^* that we are seeking, by adding to it some electrical flow \tilde{f} whose energy is at most ϵ fraction of the energy of the flow f^* . (Note that electrical flows are linear, so we indeed have that $f^* = \tilde{f} + \tilde{f}$.) As we will see, this kind of approximation is completely sufficient for our purposes.

2.4 Bipartite b-Matchings

For a given weighted bipartite graph $G = (V, E)$ with $V = P \cup Q$ – where P and Q are the two sets of bipartition – as well as, a demand vector $b \in \mathbb{R}_+^V$, a perfect b -matching is a vector $x \in \mathbb{R}_+^E$ such that $\sum_{e \in E(v)} x_e = b_v$ for all $v \in V$. A perfect b -matching is a generalization of perfect matching; in the particular case where all b 's equal 1, integer b -matchings (or 1-matchings) are exactly perfect matchings. x is called b -matching if the equality in the above equation is satisfied with inequality, i.e., we have just $\sum_{e \in E(v)} x_e \leq b_v$ for all $v \in V$. For perfect b -matchings we usually require that $b(P) = b(Q)$ as otherwise they trivially do not exist.

A weighted perfect bipartite b -matching problem is a problem in which, given a weighted bipartite graph $G = (V, E, w)$ our goal is to either return a perfect b -matching in G that has minimum weight, or conclude that there is no perfect b -matching in G . We say that a b -matching x in a graph G is h -near if the size of x is within an additive factor of h of the size of a perfect b -matching. The dual problem to the weighted perfect bipartite b -matching is a b -vertex packing problem where we want to find a vector $y \in \mathbb{R}^V$ satisfying the following LP

$$\begin{aligned} \max \quad & \sum_{v \in V} y_v b_v, \\ & y_u + y_v \leq w_{uv} \quad \forall uv \in E. \end{aligned}$$

Also, we define the size of a b -matching x to be $\|x\|_1 / 2$. A h -near b -matching is a b -matching with size at least $\|x\|_1 / 2 - h$. Finally, we observe that a bipartite b -matching instance can be reinterpreted as a unit capacitated σ -flow instance just by setting

$$\sigma_v = \begin{cases} b_v & \text{if } v \in P, \\ -b_v & \text{otherwise.} \end{cases}$$

and leaving costs unchanged. We require this alternative view, since our interior-point algorithm will work with σ -flows, but the rounding that will need to be performed at the end is done in the b -matching view.

3 Reducing Minimum-Cost Flow to Bipartite b -matching

In this section we show how to convert an instance of unit capacity min-cost flow into an instance of min-cost b -matching. This is done via a simple combinatorial reduction similar to the one in [24]. As noted in the preliminaries, bipartite b -matchings can be reinterpreted as σ -flows. This alternative view will be useful for

us, as our interior point method will work with σ -flows, whereas it is easier to repair a near-optimal solution in the b -matching view.

We first show a straightforward reduction from min-cost flow to b -matching. One desirable feature of this reduction is that the obtained b -matching instance does not contain upper capacities on arcs, since these are going to be implicitly encoded by properly setting the demands. The part of lemma that refers to half-integral flow and half-integral matching will be essential for the initialization step.

Lemma 2. *Given a directed graph $G = (V, E, c)$ and a demand vector σ , one can construct in linear time a bipartite graph $G' = (V', E', c')$, $V' = P \cup Q$ along with a demand vector b' such that given a minimum cost b' -matching in G' , one can reconstruct a flow f that routes demand σ in G with minimum cost. Moreover, if flow $f = \frac{1}{2} \cdot \vec{1}$ is feasible in G then $x = \frac{1}{2} \cdot \vec{1}$ is a feasible fractional b -matching in G' .*

Proof. Let $P = V$ and $Q = E$. For each arc $(u, v) \in E$, let e_{uv} be the corresponding “edge” vertex in Q . Create arcs (u, e_{uv}) with cost c_{uv} , and (v, e_{uv}) with cost 0. For each $e_{uv} \in Q$ set demand $b'(e_{uv}) = 1$. For each $v \in P$, set demand $b'(v) = \sigma(v) + \deg_{in}^G(v)$. Let us now argue that the solution to b' -matching encodes a valid flow in G . Observe that in the b' -matching instance each vertex e_{uv} can be in two states: it is either matched to u or to v . When e_{uv} is matched to u we “read” that there is one unit of flow on arc (u, v) , whereas when e_{uv} is matched to v we “read” that there is no flow on arc (u, v) . With this interpretation flow conservation constraints (2.1) are satisfied.

Now assume that the flow $f = \frac{1}{2} \cdot \vec{1}$ is feasible in G and consider the b -matching $x = \frac{1}{2} \cdot \vec{1}$. We observe that feasibility of f implies that x is feasible for each vertex in P , as each such vertex has the same number of incident edges and the same demand as the corresponding vertex in G . On the other hand, by construction vertices in Q have demand -1 and two incoming edges, what settles feasibility of x for them. \square

The effect of this reduction on a single arc is presented in the figure below. Note that vertices in P correspond to *vertices* from the original graph, whereas vertices in Q correspond to *arcs* from the original graph. Essentially, every arc (u, v) in G adds a demand pair of $(1, 1)$ on the pair of vertices (v, e_{uv}) in the b -matching instance. The amount of flow routed on (v, e_{uv}) corresponds to the residual capacity of the arc (u, v) in the original graph.

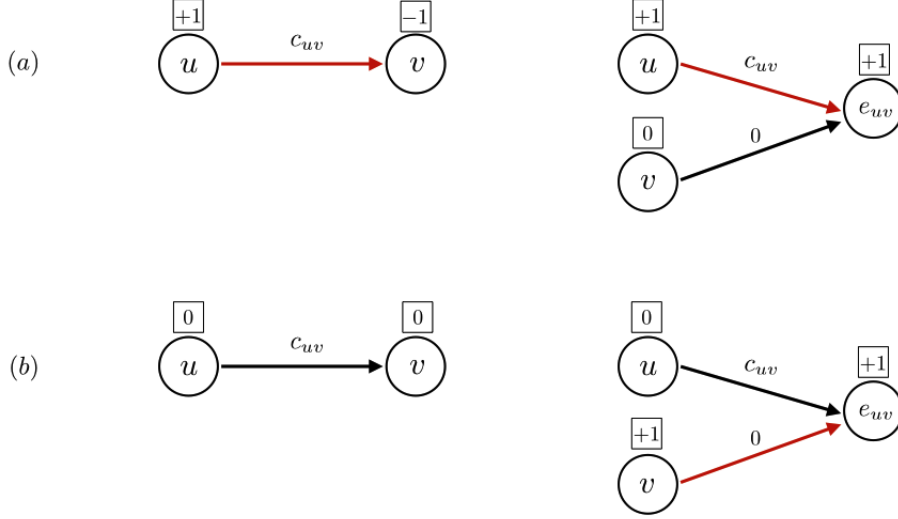


Figure 3.1: The left side contains two σ -flow instances, the right side contains their corresponding reductions to b -matching. In (a) we can see that the unit flow from u to v gets routed from u to e_{uv} in the reduction. In (b), since there is no flow from u to v , the new demand added by the reduction gets routed directly from v to e_{uv} .

While this reduction rephrases the problem into a form that is amenable to our interior-point framework, the remaining issue is that we have to be able to start the algorithm with a feasible solution where all the flows on arcs are similar in magnitude (this enables us to enforce the centrality property defined in 4.1). Ideally, we should be able to obtain a feasible instance simply by placing half a unit of flow on every arc of the b -matching instance. While doing so clearly satisfies the demands of vertices in Q , demand violations might occur on vertices in P .

We can easily fix this problem by adding one extra vertex in the original graph, along with a set of arcs with costs chosen sufficiently high, that the optimal solution will never consider them. The goal is to add the extra arcs such that flowing $\frac{1}{2}$ on every arc satisfies the demand. This instance can then be converted to a b -matching instance with the same property. The reduction is described by the following lemma:

Lemma 3. *Given a directed graph $G = (V, E, c)$ and an integral demand vector σ , one can construct in linear time a graph $G' = (V', E', c')$ along with an integral demand vector σ' such that the demand is satisfied by placing $\frac{1}{2}$ units of flow on every arc. Furthermore, given a solution to the min-cost σ' -flow satisfying demand σ' in G' , one can reconstruct in linear time a solution f that routes demand σ in G with minimum cost.*

Proof. Create one extra vertex v_{aux} with demand 0. Then, for all $v \in V$, let $t(v) = \sigma(v) + \frac{1}{2} \cdot \deg_{\text{in}}^G(v)/2 - \frac{1}{2} \cdot \deg_{\text{out}}^G(v)$ be the residual demand corresponding to the flow that has value $\frac{1}{2}$ everywhere.

Fix this residual by creating $|2t(v)|$ parallel arcs (v_{aux}, v) with costs $\|c\|_1$, for each vertex with $t(v) < 0$, respectively $|2t(v)|$ parallel arcs (v, v_{aux}) with costs $\|c\|_1$ for each vertex v with $t(v) > 0$. This enforces our condition to be satisfied for all vertices in V .

Also note that v_{aux} has an equal number of arcs entering and leaving it, since the sum of residuals at vertices in V equals the sum of degree imbalances, plus the sum of demands, both of which are equal to 0. More precisely, $\deg_{\text{in}}(v_{\text{aux}}) = \sum_{v:t(v)>0} 2t(v)$, and $\deg_{\text{out}}(v_{\text{aux}}) = \sum_{v:t(v)<0} -2t(v)$; since $0 = \sum_v t(v) = \sum_{v:t(v)>0} t(v) - \sum_{v:t(v)<0} (-t(v)) = \deg_{\text{in}}^{G'}(v_{\text{aux}})/2 - \deg_{\text{out}}^{G'}(v_{\text{aux}})/2$, the vertex v_{aux} is balanced; hence $t(v_{\text{aux}}) = 0$, and the residual demand is 0 at all vertices in the graph.

Finally, observe that a flow in G satisfying σ is a valid flow in G' for σ' that does not use any edge incident to v_{aux} . This flow has cost smaller than $\|c\|_1$. Hence, a min-cost σ' -flow in G' , if it has cost smaller than $\|c\|_1$, gives a min-cost σ -flow in G . \square

A pictorial description of the reduction is presented in the figure below.




Figure 3.2: Two examples of balancing a vertex using extra arcs connected to v_{aux} . In the first case, adding two arcs to v_{aux} makes the residual at v equal to 0, when routing $1/2$ on every arc, since the in-degree of v is equal to its out-degree, and there is no demand on v . In the second case, we add four arcs from v_{aux} to v in order for the net flow at v (when routing $1/2$ on every arc) to match the demand -3 .

We can use these two reductions together, first making the flow $f = \frac{1}{2} \cdot \bar{1}$ feasible, then converting the instance to a b -matching instance that can be directly changed back to σ -flow instance. As a result we obtain an instance where constructing a feasible starting solution that is required for our interior point method is straightforward.

4 Our Interior-Point Method Framework

In this section we describe our interior point method framework for solving our instance of the uncapacitated minimum-cost σ -flow problem that results from casting the bipartite b -matching problem instance we produced in Section 3 as an instance of the minimum-cost σ -flow problem. Our basic setup is largely following the framework used by Mądry [24]. The crucial differences emerge only later on, in Sections 5 and 6, when we refine it to obtain our desired running time improvements.

4.1 Primal and Dual Solutions and Centrality

Our framework will be inherently primal dual. That is, we will always maintain a feasible primal dual solution (f, y) . Here, the primal solution f is simply a σ -flow, i.e., a flow with demands σ and its feasibility condition is that $f_e \geq 0$ for each arc e , i.e., the flow never flows against the orientation of the arc. The dual solution y , on the other hand, corresponds to embedding of the vertices of our graph into the line, i.e., y assigns a real value y_v to each vertex v . The dual feasibility condition is that, for each arc $e = (u, v)$, we have that its *slack* $s_e := c_e + y_u - y_v$ is non-negative, i.e., $s_e \geq 0$, for each arc e . Intuitively, this means that no arc e is stretched in the embedding defined by y more than its actual cost c_e . (Note that here we measure stretch only in the direction of the arc orientation.)

Observe that the dual solution y is define uniquely (up to a translation) by the slack vector s and the arc costs c . So, we will sometime find it more convenient to represent our dual solution in terms of the slack

vector s instead of the embedding vector y . From this perspective, the dual feasibility condition is simply non-negativity of all the slacks s . (In fact, we will use y and s representation interchangeably, depending on which one of them is more convenient in given context.)

Duality Gap. A very convenient feature of working with primal dual solutions is that they provide a very natural way of measuring optimality: the duality gap, i.e., the difference between the upper bound on the cost of the optimal solution provided by current primal feasible solution f and the lower bound on the optimal cost provided by our current dual solution s . It turns out that the duality gap of a primal dual solution (f, s) is exactly

$$f^\top s = \sum_e f_e s_e = \sum_e \mu_e,$$

where $\mu_e := f_e s_e$ is the contribution of the arc e to that duality gap. Observe that by our primal dual feasibility condition we have always that $\mu_e \geq 0$, for each arc e .

Centrality. In the light of the above, the duality gap $\|\mu\|_1 = \sum_e \mu_e$ provides a natural measure of progress for our algorithm. In particular, our goal is simply to find a primal dual feasible solution (f, s) with $\|\mu\|_1$ sufficiently close to 0, and, from this perspective, any way of updating that primal dual solution that leads to reduction of the duality gap should be desirable.

However, for reasons that will become clear later, we will insist on reducing the duality gap $\|\mu\|_1$ in a more restricted manner. To define this, let us first associate with each arc e a value ν_e that we will refer to as the *measure* of e . We will always make sure that $\nu_e \geq 1$, for each arc e , and that the total sum of all the measures is not too large. Specifically, we want to maintain the following invariant that ensures that the average measure of an arc is at most 3.

Invariant 1. We always have that $\|\nu\|_1 = \sum_e \nu_e \leq 3m$.

As it turns out, once we introduced the notion of arc measure, it will be much more convenient to introduce analogues of the standard ℓ_p -norms that are reweighted by the measures of arcs. To this end, for a given vector $x \in \mathbb{R}^m$, and measure vector ν , let us define

$$\|x\|_{\nu,p} := \left(\sum_e \nu_e |x_e|^p \right)^{\frac{1}{p}} \quad (4.1)$$

We will sometimes extend this notation to refer to the weighted norm of a subset of indices in the support. More specifically, given $S \subseteq \{1, \dots, m\}$, we will call

$$\|x_S\|_{\nu,p} := \left(\sum_{e \in S} \nu_e |x_e|^p \right)^{\frac{1}{p}} \quad (4.2)$$

Also, we will extend our notation and use (f, s, ν) to denote a primal dual solution (f, s) with its corresponding measure vector ν .

Now, we can make precise the additional constraint on our primal dual solutions (f, s, ν) that we want to maintain. Namely, we want each solution (f, s, ν) to be $\widehat{\mu}$ -centered (or, simply, *centered*), i.e., we want that, for each arc e ,

$$\mu_e = \nu_e \widehat{\mu}, \quad (4.3)$$

where $\hat{\mu}$ is a normalizing value we will refer to as *average duality gap*. Intuitively, centrality means that each arc's contribution to the duality gap is exactly proportional to its measure.²

Note that the above notions enable us to express the duality gap of a solution (f, y, ν) as exactly $\sum_e \nu_e \hat{\mu}$, which by Invariant 1 is at most $3m\hat{\mu}$. That is, we have that

$$\sum_e \mu_e = \sum_e \nu_e \hat{\mu} \leq 3m\hat{\mu} \quad (4.4)$$

Consequently, we can view $\hat{\mu}$ as a measure of our progress - driving it to be smaller translates directly into making the duality gap smaller too.

4.2 Making Progress with Electrical Flow Computations

Once we setup basic definitions, we are ready to describe how we initialize our framework and then how we can use electrical flow computations to gradually improve the quality, i.e., the average duality gap $\hat{\mu}$ of our $\hat{\mu}$ -centered primal dual solution (f, y, ν) .

Initialization. As we want our solutions (f, y, ν) to be always centered, initialization of our framework, i.e., finding the initial centered primal dual solution, might be difficult. Fortunately, one of the important properties of the reduction we performed in Section 3 is that a centered primal dual feasible solution of the resulting uncapacitated minimum-cost σ -flow instance can be specified explicitly.

Lemma 4. *Given a bipartite graph $G = (V, E, c)$ along with an integral demand vector σ and a subset of vertices P such that that for all $v \in P$ we have $\sigma_v = \deg(v)/2$, whereas for all $v \notin P$ we have $\deg(v) = 2$, one can construct in linear time a feasible primal-dual set of variables (f, s) that satisfy the centrality bound for $\hat{\mu} = \|c\|_\infty$ and $\|\nu\|_1 \leq m$.*

Proof. Since $\sigma_v = \deg(v)/2$ for all $v \in P$, while all $w \notin P$ have degree precisely 2, we can set $f = \frac{1}{2} \cdot \vec{1}$ and have all the demands satisfied exactly. Moreover, we set the dual variables $y_v = \|c\|_\infty$ for all $v \in P$, and $y_w = 0$ for all $w \notin P$. This way the slacks $s_{vw} = c_{vw} - y_w + y_v$ are all within the range $[\|c\|_\infty, 2\|c\|_\infty]$. We set $\nu_e = \frac{s_e}{2\|c\|_\infty}$ and $\hat{\mu} = \|c\|_\infty$ so that

$$\mu_e = f_e s_e = \frac{1}{2} s_e = \frac{s_e}{2\|c\|_\infty} \|c\|_\infty = \nu_e \hat{\mu},$$

and

$$\|\nu\|_1 = \sum_e \nu_e = \sum_e \frac{s_e}{2\|c\|_\infty} \leq \sum_e 1 = m.$$

□

Taking an Improvement Step. Let us fix some $\hat{\mu}$ -centered primal dual solution (f, y, ν) and let us define resistances r to be equal to

$$r_e := \frac{1}{\hat{\mu}} \cdot \frac{s_e}{f_e} = \frac{1}{\hat{\mu}} \cdot \frac{\mu_e}{f_e^2} = \frac{\nu_e}{f_e^2}, \quad (4.5)$$

for each arc e . (Note that f has to always be positive due to centrality condition, and thus these resistances are well-defined.)

²The framework in [24] works with a slightly relaxed notion of centrality. However, we deviate from that here.

The fundamental object that will drive our improvements of the quality of our current primal dual solution will be the electrical σ -flow \hat{f} determined by the above resistances r . For future reference, we will call the electrical flow \hat{f} *associated with* (f, s, ν) . The key property of that electrical flow is that it will enable us to update our primal and dual solutions *simultaneously*. That is, we can use the flow itself to update the primal solution f , and we can use the vertex potentials $\hat{\phi}$ that induced \hat{f} to update our dual solution s . Specifically, our main improvement update step, for each arc $e = (u, v)$ is:

$$\begin{aligned} f'_e &:= (1 - \delta)f_e + \delta\hat{f}_e, \\ s'_e &:= s_e - \frac{\delta}{(1 - \delta)} \left(\hat{\phi}_v - \hat{\phi}_u \right), \end{aligned}$$

where δ is a *step size* that we will choose later.

Remark 5. The step derived from the standard primal-dual interior-point method computes an electrical flow along with potentials determined by resistances $\frac{s_e}{f_e}$, which are off by precisely a factor of $\hat{\mu}$ from the resistances we consider in this paper. However, scaling all resistances by the same factor has no effect on the electrical flow or the potentials produced. Setting resistances the way we do in (4.5) has the benefit that it will enable us to relate the electrical energy with another quantity of interest without having to carry along the extra $\frac{1}{\hat{\mu}}$ factor, as we will see in Lemma 7.

Intuitively, this update step mixes the electrical flow \hat{f} with the current solution f by taking a convex combination of them. (Note that the resulting flow is guaranteed to be a σ -flow in this way.) On the other hand, the dual update corresponds to updating the line embedding of each vertex by adding an appropriately scaled vertex potential to it.

It is worth pointing out that the electrical flow \hat{f} is inherently undirected. So, it is not a priori clear if the flow f' resulting from the above update is even feasible. As a result, we will need to ensure, in particular, that the step size δ is chosen to be small enough so as to ensure that f' is still feasible. (In fact, as we will see shortly, there are some even stronger restrictions on the value of δ . So, the feasibility will be enforced implicitly.)

Congestion Vector. A notion that will be extremely useful in analyzing our improvement step and the performance of our algorithm in general is the notion of congestion vectors. Specifically, given the electrical σ -flow \hat{f} associated with our solution (f, y, ν) , let us define *congestion* ρ_e of an arc e as

$$\rho_e := \frac{|\hat{f}_e|}{f_e} \tag{4.6}$$

Now, observe that we can express the duality contribution μ'_e of an arc e in the new solution (f', s') as

$$\begin{aligned} \mu'_e &= f'_e s'_e = \left((1 - \delta)f_e + \delta\hat{f}_e \right) \left(s_e - \frac{\delta}{(1 - \delta)} \left(\hat{\phi}_v - \hat{\phi}_u \right) \right) \\ &= (1 - \delta)f_e s_e - \delta f_e \left(\hat{\phi}_v - \hat{\phi}_u \right) + \delta \hat{f}_e s_e - \frac{\delta^2}{1 - \delta} \hat{f}_e \left(\hat{\phi}_v - \hat{\phi}_u \right) \\ &= (1 - \delta)f_e s_e - \delta f_e \cdot \hat{f}_e \frac{s_e}{f_e} + \delta \hat{f}_e s_e - \frac{\delta^2}{1 - \delta} \hat{f}_e \cdot \hat{f}_e \frac{s_e}{f_e} \\ &= (1 - \delta)\mu_e - \frac{\delta^2}{1 - \delta} \mu_e \rho_e^2 \end{aligned}$$

So, if we ignore the second-order term in δ , the duality gap contribution of each arc e goes down at the same rate. In this way, not only the duality gap gets reduced by a factor of $(1 - \delta)$ but also the centrality of the solution would be perfectly preserved.

However, we cannot really ignore the second-order term and this term will make our solution lose centrality. Fortunately, one can show that as long as the total degradation of centrality condition is not too large one can easily correct it with a small number of electrical flow computations. Specifically, for the correction to be possible, we need to have that the total ℓ_2^2 -norm of the degradations (measured with respect to measure ν and normalized by the duality gap contributions μ) has to be a small constant. That is, we need that

$$\sum_e \nu_e \left(\frac{\mu'_e}{(1 - \delta)\mu_e} - 1 \right)^2 = \sum_e \nu_e \left(\frac{\delta^2 \mu_e \rho_e^2}{(1 - \delta)^2 \mu_e} \right)^2 \frac{\delta^4}{(1 - \delta)^4} \sum_e \nu_e \rho_e^4 = \frac{\delta^4}{(1 - \delta)^4} \|\rho\|_{\nu,4}^4 \leq \frac{1}{256}, \quad (4.7)$$

which implies that it is sufficient to have

$$\delta \leq \frac{1}{8 \cdot \|\rho\|_{\nu,4}},$$

i.e., that the step size δ should be bounded by the ℓ_4 norm of the congestion vector ρ . The following theorem makes these requirements, as well as the result of the full improvement step, precise. Its complete proof can be found in Appendix C.

Theorem 6. *Let (f, s, ν) be a $\hat{\mu}$ -centered solution and let ρ the congestion vector of the electrical flow \hat{f} associated with that solution. For any $\delta > 0$ such that*

$$\delta \leq \min \left\{ \frac{1}{8 \cdot \|\rho\|_{\nu,4}}, \frac{1}{8} \right\},$$

we can compute in $\tilde{O}(m)$ time a $\hat{\mu}'$ -centered solution (f', s', ν') , such that $\nu' = \nu$, $\hat{\mu}' \leq (1 - \delta)\hat{\mu}$, and, for each arc e ,

$$r'_e = \frac{1}{\hat{\mu}'} \cdot \frac{s'_e}{f'_e} \geq (1 + 4 \cdot \delta \rho_e + \kappa_e)^{-1} r_e,$$

where κ is a vector with $\|\kappa\|_{\nu,2} \leq 1$.

4.3 A Simple $O(\sqrt{m} \log W)$ -iteration Bound

Once the ℓ_4 norm bound provided in Theorem 6 is established we are already able to prove in a simple way that our algorithm needs at most $O(\sqrt{m} \log W)$ iterations to compute the optimal solution, making its total running time be at most $O(m^{3/2} \log W)$. To achieve that, we just need to argue that we always have that

$$\|\rho\|_{\nu,4} \leq O(\sqrt{m}). \quad (4.8)$$

Once this is established, by Theorem 6, we know that we can always take $\delta = \Omega(m^{-1/2})$ and thus make $\hat{\mu}$ decrease by a factor of $(1 - \delta)$ in each iteration. So, after $O(\sqrt{m} \log W)$ iterations, $\hat{\mu}$ and thus the duality gap becomes small enough that a simple rounding (see Section 7) will recover the optimal solution.

Now, to argue that 4.8 indeed holds, we first notice that we can always upper bound ℓ_4 norm $\|\rho\|_{\nu,4}$ by the ℓ_2 norm $\|\rho\|_{\nu,2}$ and then bound the latter norm instead. Next, as it turns out, we can tie the energy $\mathcal{E}(\hat{f})$ of the electrical flow \hat{f} associated with a given solution (f, s, ν) to the corresponding ℓ_2 norm $\|\rho\|_{\nu,2}$. Specifically, we have the following lemma.

Lemma 7. For any centered solution (f, s, ν) , we have that

$$\|\rho\|_{\nu,2}^2 = \mathcal{E}(\hat{f}),$$

where \hat{f} is the electrical flow associated with that solution and ρ is its congestion vector.

Proof. Observe that by definition 4.6 and 4.5, we have that

$$\|\rho\|_{\nu,2}^2 = \sum_e \nu_e \rho_e^2 = \sum_e \nu_e \left(\frac{\hat{f}_e}{f_e} \right)^2 = \sum_e r_e \hat{f}_e^2 = \mathcal{E}(\hat{f})$$

Note that we used (4.5) to write $r_e = \frac{\nu_e}{f_e^2}$, which assumes centrality. □

Due to this identity, we can view the ℓ_2 norm $\|\rho\|_{\nu,2}^2$ as energy. Finally, we can use the bound from Invariant 1 to show that the energy $\mathcal{E}(\hat{f})$ and thus the norm $\|\rho\|_{\nu,2}$ can be appropriately bounded as well.

Lemma 8. For a centered solution, $\|\rho\|_{\nu,2}^2 \leq \sum_e \nu_e = \|\nu\|_1$.

Proof. By Lemma 7 and (4.5), we have that

$$\|\rho\|_{\nu,2}^2 = \mathcal{E}(\hat{f}) \leq \mathcal{E}(f) = \sum_e r_e f_e^2 = \sum_e \nu_e \left(\frac{f_e}{f_e} \right)^2 = \sum_e \nu_e = \|\nu\|_1,$$

where the inequality follows as f is a σ -flow and, by the virtue of being an electrical σ -flow, \hat{f} has to have minimum among all the σ -flows. □

Now, since by Invariant 1, $\|\nu\|_1 \leq 3m$, we can conclude that

$$\|\rho\|_{\nu,4}^2 \leq \|\rho\|_{\nu,2}^2 \leq \|\nu\|_1 \leq 3m$$

and the bound 4.8 follows.

With this upper bound on the $\|\rho\|_{\nu,4}$ we can immediately derive a bound on the running time required to obtain an exact solution. This is summarized in the following theorem.

Theorem 9. We can produce an exact solution to the unit-capacitated minimum cost σ -flow problem in $\tilde{O}(m^{3/2} \log W)$ time.

Proof. Given an instance of the unit-capacitated minimum cost σ -flow, we can apply the reduction from Section 3 in linear time. Then, Lemma 4 establishes the initial centering with $\hat{\mu} = \|c\|_\infty \leq W$. We previously saw that $\|\rho\|_{\nu,4} = O(\sqrt{m})$. Therefore, according to Theorem 6, we can set $\delta = 1/O(\sqrt{m})$, and reduce $\hat{\mu}$ by a factor of $1 - \frac{1}{O(\sqrt{m})}$ with every interior-point iteration. Therefore, in $O(m^{1/2} (\log m + \log W))$ iterations we reduce $\hat{\mu}$ to $O(m^{-3})$. Using the fact that $\|\nu\|_1 \leq 3m$, the duality gap of this solution will be $\sum_e \nu_e \hat{\mu} \leq m^{-2}$. Note that each iteration requires $\tilde{O}(1)$ electrical flow computations, and each of them can be implemented in near-linear time, according to Theorem 1.

Therefore in $\tilde{O}(m^{3/2} \log W)$ time, we obtain a feasible primal-dual solution with duality gap less than m^{-2} . This can easily be converted to an integral solution in nearly-linear time using the method described in Section 7. Hence the total running time of the algorithm is $\tilde{O}(m^{3/2} \log W)$. □

5 Taking Longer Steps

In the previous section, we established that the amount of progress we can make in one iteration of our framework is limited by the ℓ_4 norm of the congestion, $\|\rho\|_{\nu,4}$ - see Theorem 6. We then showed (cf. (4.8)) that this norm is always upper bounded by $O(\sqrt{m})$, which resulted in our overall $\tilde{O}(m^{3/2} \log W)$ time bound.

Unfortunately, a priori, this upper bound is tight, i.e., it indeed can happen that $\|\rho\|_{\nu,4} = \|\rho\|_{\nu,2} = \Omega(\sqrt{m})$. In fact, this is exactly the reason why all classic analyses of interior-point algorithms are able to obtain only an $O(\sqrt{m})$ iteration bound.

To circumvent this problem, [24] introduced a perturbation technique into the framework. These perturbations target arcs that contribute a significant fraction of the norm $\|\rho\|_{\nu,4}$, by increasing their resistance (and thus the corresponding energy), in order to discourage the emergence of such high contributing arcs in the future. A careful analysis done in [24] shows that such perturbations indeed ensure that there are sufficiently many iterations with relatively small $\|\rho\|_{\nu,4}$ norm to guarantee ability to take longer steps, and thus converge faster. Unfortunately the changes to the underlying optimization problem that these perturbations introduced, although sufficiently mild to enable obtaining a result for unit-capacity maximum flow, were too severe to enable solving any of the weighted variants that we aim to tackle here.

Our approach will also follow the same general outline. The crucial difference though is that we use a different perturbation technique, along with a somewhat simpler analysis. This technique still achieves the desired goal of increasing the resistance of the perturbed arcs. However, in big contrast to the technique used by [24], our technique does not affect the costs of those arcs – it affects only their measure. Also, as an added benefit, our perturbation treatment simplifies the analysis significantly.

We describe our preconditioning technique in Section 5.1. Also, in the table below we present a general outline of our algorithm. (This algorithm will be later modified further to include a preconditioning step.) Observe that this algorithm uses a stronger, ℓ_3 norm criterion for whether to make a perturbation or a progress step, instead of the ℓ_4 norm criterion that Theorem 6 suggests. As we will see, the reason for that is that maintaining such an ℓ_3 norm condition will enable us to have a sufficiently tight control over the change after each progress step of our potential function: the energy of electrical flows \hat{f} associated with our primal dual solutions.

Our goal is to obtain an $\tilde{O}(m^{1/2-\eta})$ bound on the overall number of iterations, where we fix η to be

$$\eta = \frac{1}{14}$$

Algorithm 1 Perturbed interior-point method (parameters: $c_\rho = 400\sqrt{3} \cdot \log^{1/3} W$, $c_T = 3c_\rho \log W$)

1. initialize primal and dual variables (f, y) (as shown in Section 4);
 2. repeat $c_T \cdot m^{1/2-\eta}$ times
 3. while $\|\rho\|_{\nu,3} > c_\rho \cdot m^{1/2-\eta}$
 4. perform perturbation (as shown in Section 5.1)
 5. perform progress steps (as shown in Section 4.2)
-

5.1 Our Perturbation Technique

Let us start by describing our perturbation technique, which heavily uses the structure of the b -matching instance obtained after applying the reduction from Section 3. We first show how a perturbation is applied to an arc, then we define the set of arcs that get perturbed in each iteration. As we will see, whenever we perturb a particular arc (to increase its resistance) we always make sure to perturb its “partner” arc, i.e., the unique arc sharing a common vertex from the set Q of the bipartition, as well.

Definition 10. Given an arc $e = (u, v)$, with $u \in P$, $v \in Q$, the *partner arc* of e is the unique arc $\bar{e} = (\bar{u}, v)$, $\bar{u} \in P$ sharing vertex v with e .

5.1.1 Perturbing an Arc.

Let $e = (u, v)$ be an arc with cost c_{uv} and vertex potentials y_u , respectively y_v , and slack $s_{uv} = c_{uv} + y_u - y_v$. Note that due to the structure of our b -matching instance (see Section 3), vertex v is of degree 2. Let $\bar{e} = (\bar{u}, v)$ be the partner arc that shares with e this vertex v . We first modify our dual solution by setting $y_v \leftarrow y_v - s_{uv}$. This effectively doubles the resistance of e , defined as in (4.5), which is our desired effect.

Unfortunately, this update breaks centrality of *both* arc e and its partner arc \bar{e} . To counteract that, we first double the measure ν_e of e - this immediately restores the centrality of that arc. Now, it remains to fix the centrality of the partner arc $\bar{e} = (\bar{u}, v)$. Specifically, we need to deal with the fact that the slack $s_{\bar{e}}$ of that partner arc gets increased by s_e . To fix this problem, recall that the centrality condition for \bar{e} guaranteed that $s_{\bar{e}}f_{\bar{e}} = \nu_{\bar{e}}\hat{\mu}$. So, we need to set the new measure $\nu'_{\bar{e}}$ such that $(s_{\bar{e}} + s_e)f_{\bar{e}} = \nu'_{\bar{e}}\hat{\mu}$. Therefore we just set the new measure to be

$$\nu'_{\bar{e}} = \frac{(s_{\bar{e}} + s_e)f_{\bar{e}}}{\hat{\mu}} = \nu_{\bar{e}} + \frac{s_e f_{\bar{e}}}{\hat{\mu}} = \nu_{\bar{e}} + \nu_e \cdot \frac{f_{\bar{e}}}{f_e}$$

Consequently, the total change in measure of that arc is

$$\nu_e \left(1 + \frac{f_{\bar{e}}}{f_e} \right) \leq \nu_e \left(1 + \frac{1}{f_e} \right) \quad (5.1)$$

as in our instance we have that $f_{\bar{e}} \leq 1$, since the arcs are unit capacitated, and f is always feasible.

We remark that we may want to perturb both an arc $e = (u, v)$ and its partner $\bar{e} = (\bar{u}, v)$. In this case, we can perturb the arcs simultaneously by setting $y_v \leftarrow y_v - s_{uv} - s_{\bar{u}v}$, and updating the measures: $\nu_{\bar{e}} \leftarrow 2\nu_{\bar{e}} + \nu_e \cdot \frac{f_{\bar{e}}}{f_e}$, $\nu_e \leftarrow 2\nu_e + \nu_{\bar{e}} \cdot \frac{f_e}{f_{\bar{e}}}$. This maintains centrality, and the bound from (5.1) still holds.

So, to summarize, one effect of the above operation is that it made the resistance of the perturbed arc e double. As we will see, similarly as it was the case in [24], this will enable us to ensure that the total number of perturbation steps is not too large. Also, note that the above operation does not change any vertex demands or costs. It only affects the dual solution and the arcs' measure. Therefore, the only undesirable long term effect of it is the measure increase, since it might lead to violation of Invariant 1. ³

5.1.2 Which Arcs to Perturb?

As we have just seen, while perturbing an arc doubles its resistance, this operation has the side effect of increasing the total measure. To control the latter, undesirable effect, we show that every time we need to perturb the problem, we can actually select a subset of arcs with the property that perturbing all of them

³In fact, if we were to formulate our problem as a primal interior-point method, one could think of these perturbations on arcs and their partners as acting on both the lower and the upper barriers. In that formulation, the barrier would be of the form $-\sum_e \nu_e \log f_e + \nu_{\bar{e}} \log(1 - f_e)$. The reduction from Section 3 essentially eliminates the upper barrier, in order to be make our problem amenable to a primal-dual approach, which we preferred to use here.

increases the energy by a lot while keeping the measure increase bounded. Ultimately, the entire goal of the analysis will be to show that:

1. We do not need to perturb the problem more than $\tilde{O}(m^{1/2-\eta})$ times
2. The total increase in measure caused by the perturbations is at most $2m$, thus maintaining Invariant 1.

Below we define the subset of arcs that will get perturbed. Intuitively, we only want to perturb arcs e with high congestion ρ_e . Furthermore, we choose to perturb them only if the amount of flow they carry is not very small. This extra restriction enforces an upper bound on the amount by which the measure of the perturbed arc increases, as per equation 5.1.

As we will see in Corollary 16, perturbing edge e will increase total energy by at least a quantity that is proportional to the amount of contribution to total energy of that edge; therefore the total measure increase will be upper bounded by a quantity proportional the total energy increase due to perturbations.

We will soon see that the total energy increase across iterations is, as a matter of fact, bounded by $O(c_T c_\rho^2 \cdot m^{3/2-3\eta})$, where c_T and c_ρ are some appropriately chosen constants, which immediately yields the desired bound on the total measure increase.

Definition 11. An arc e is *perturbable* if $\frac{1}{f_e} \leq m^{1/2-3\eta}$ and $\rho_e \geq \sqrt{40c_T c_\rho^2} \cdot m^{1/2-3\eta}$. An arc that is not perturbable is called *unperturbable*. Denote by \mathcal{S} the set of all perturbable arcs.

A useful property of perturbable arcs the way they are chosen enforces a small increase in measure compared to that in energy during each perturbation. We will make this property precise below, and it is what we will be using for the remainder of the section.

Corollary 12. A perturbable arc satisfies $1 + \frac{1}{f_e} \leq C \rho_e^2$, where $C = \frac{1}{20c_T c_\rho^2} \cdot m^{-1/2+3\eta}$.

5.2 Runtime Analysis

The analysis of our algorithm is based on two major parts.

The first part shows that throughout the execution of the algorithm, the total energy increase caused by perturbations can not be too large. This will automatically imply that total measure increase will be bounded by $2m$, and therefore Invariant 1 is preserved. The key idea is that, since they are applied only when the ℓ_3 norm of the congestion is “small” (i.e. $c_\rho \cdot m^{1/2-\eta}$), progress steps do not decrease the energy by a lot (i.e. $O(c_\rho^2 \cdot m^{1-2\eta})$), as we will soon see). However, since total measure, and hence energy, was $O(m)$ to begin with, perturbations could not have increased energy by more than progress steps have decreased it overall. Over the $O(c_T \cdot m^{1/2-\eta})$ iterations, progress steps decrease energy by at most $O(c_T c_\rho^2 \cdot m^{3/2-3\eta})$; therefore this is also a bound on the total increase in energy.

The second part use the invariant that perturbable edges consume most of the energy in the graph, in order to argue that the number of perturbations is small. While a priori we only had a bound on the time required for progress steps, with no guarantee on how many iterations the algorithm spends performing perturbations, this argument provides a bound on the number of perturbations, and hence on the running time of the algorithm. Showing that, every time we perform a perturbation, energy increases by at least $\Omega(c_\rho \cdot m^{1-2\eta})$ implies, together with the bound proven in the first part, that throughout the execution of the algorithm we perform only $O(c_T c_\rho \cdot m^{1/2-\eta})$ perturbations. This bounds the running time by $\tilde{O}(c_T c_\rho \cdot m^{3/2-\eta})$, since perturbing the problem takes only $\tilde{O}(m)$ time.

The invariant that the second part relies on is motivated by the fact that, whenever we have to perturb the problem, the ℓ_3 norm of the congestion vector is large, so the energy of the system is also large (at least $c_\rho^2 \cdot m^{1-2\eta}$). Since perturbable arcs are highly congested, we expect them to contribute most of the energy; so perturbing those should increase the energy of the system by a quantity proportional to their

current contribution to energy. Maintaining this invariant requires a finer control over how the electrical flows behave, and will be guaranteed via a modification of the algorithm, which will be carefully analyzed in Section 6. However, the future modification will not affect any of the analysis described in this section. Therefore this section will be concerned only with proving the runtime guarantee, assuming validity of the invariant.

5.2.1 Bounding the Total Increase in Measure and Energy

We formalize the intuition described at the beginning of the section. First, we show that Theorem 6 provides a bound on how much energy can decrease during one progress step. This relies on the following lemma, which allows us to lower bound the energy of an electrical flow.

Lemma 13. *Let \mathcal{E}_r be the energy of the electrical flow in a graph with demands σ and resistances r . Then*

$$\mathcal{E}_r = \max_{\phi} \left(2\sigma^\top \phi - \sum_{e=(u,v)} \frac{(\phi_u - \phi_v)^2}{r_e} \right) \quad (5.2)$$

Proof. The result can be derived by letting L be the Laplacian corresponding to the graph with resistances r , and rewriting the above maximization problem as $\max_{\phi} 2\sigma^\top \phi - \phi^\top L \phi$. By first order optimality conditions we get that the maximizer satisfies $L\phi = \sigma$, hence $\phi = L^+ \sigma$. Plugging in makes the expression equal to $\sigma^\top L^+ \sigma$, which is precisely the energy \mathcal{E}_r . \square

In our context, this lemma enables us to provide a more convenient formula for lower bounding the new value of energy after resistances change.

Lemma 14. *Let \mathcal{E}_r be the energy of the electrical flow corresponding to a centered instance with resistances r , and let $\mathcal{E}_{r'}$ be the energy of the electrical flow corresponding to the new centered instance with resistances r' , obtained after applying one progress step or one perturbation. Then the change in energy can be lower bounded by:*

$$\mathcal{E}_{r'} - \mathcal{E}_r \geq \sum_{e=(u,v)} \nu_e \rho_e^2 (1 - r_e/r_{e'}) \quad (5.3)$$

Proof. Let $\hat{\phi}$ be the potentials that maximize the expression from (5.2) for resistances \mathcal{E}_r . Therefore we have

$$\mathcal{E}_r = 2\sigma^\top \hat{\phi} - \sum_{e=(u,v)} \frac{(\hat{\phi}_u - \hat{\phi}_v)^2}{r_e}$$

Using the same set of potentials in order to certify a lower bound on the new energy, we obtain:

$$\begin{aligned}
\mathcal{E}_{r'} &\geq 2\sigma^\top \hat{\phi} - \sum_{e=(u,v)} \frac{(\hat{\phi}_u - \hat{\phi}_v)^2}{r'_e} \\
&= 2\sigma^\top \hat{\phi} - \sum_{e=(u,v)} \frac{(\hat{\phi}_u - \hat{\phi}_v)^2}{r_e} + \sum_{e=(u,v)} \frac{(\hat{\phi}_u - \hat{\phi}_v)^2}{r_e} \left(1 - \frac{r_e}{r'_e}\right) \\
&= \mathcal{E}_r + \sum_{e=(u,v)} \frac{(\hat{\phi}_u - \hat{\phi}_v)^2}{r_e} \left(1 - \frac{r_e}{r'_e}\right) \\
&= \mathcal{E}_r + \sum_{e=(u,v)} \nu_e \rho_e^2 \left(1 - \frac{r_e}{r'_e}\right)
\end{aligned}$$

For the last identity we used the fact that if \hat{f} is the electrical flow corresponding to potentials $\hat{\phi}$, then $\frac{(\hat{\phi}_u - \hat{\phi}_v)^2}{r_e} = r_e \hat{f}_e^2 = \frac{\nu_e}{\hat{f}_e^2} \hat{f}_e^2 = \nu_e \rho_e^2$. \square

A first application of this lemma is that it enables us to lower bound the increase in energy when perturbing arcs.

Lemma 15. *After perturbing arcs in \mathcal{S} , energy increases by at least $\frac{1}{2} \|\rho_{\mathcal{S}}\|_{\nu,2}^2$.*

Proof. According to the effects of the perturbation described in Section 5.1.1, all resistances of arcs in \mathcal{S} get doubled, while the others can only increase. Therefore, applying Lemma 14, we obtain a lower bound on the energy increase:

$$\begin{aligned}
\mathcal{E}_{r'} - \mathcal{E}_r &\geq \sum_{e \in \mathcal{S}} \nu_e \rho_e^2 \left(1 - \frac{r_e}{2r_e}\right) \\
&= \frac{1}{2} \sum_{e \in \mathcal{S}} \nu_e \rho_e^2 \\
&= \frac{1}{2} \|\rho_{\mathcal{S}}\|_{\nu,2}^2
\end{aligned}$$

\square

An immediate corollary is that the increase in energy during a perturbation upper bounds the increase in measure.

Corollary 16. *If a perturbation increases energy by Δ , then the total measure increases by at most $2C \cdot \Delta$.*

Proof. By definition, the arcs we perturb satisfy $1 + \frac{1}{f_e} \leq C\rho_e^2$. According to (5.1), the measure increase cause by perturbing an arc e is at most $\nu_e \left(1 + \frac{1}{f_e}\right)$. Therefore, perturbing all the arcs in $\hat{\mathcal{S}}$, increases measure by at most $\sum_{e \in \hat{\mathcal{S}}} \nu_e \cdot C\rho_e^2 = C \cdot \|\rho_{\mathcal{S}}\|_{\nu,2}^2$. But Lemma 15 shows that $\Delta \geq \frac{1}{2} \|\rho_{\mathcal{S}}\|_{\nu,2}^2$. Combining these two bounds yields the result. \square

While Lemma 15 tells us that perturbations increase energy, we can show that progress steps do not decrease it by too much, using another application of Lemma 14.

Lemma 17. *Let \mathcal{E}_r be the energy of an electrical flow corresponding to a centered solution with congestion vector ρ , and let $\mathcal{E}_{r'}$ be the new energy after applying one progress step. Then $\mathcal{E}_{r'} \geq \mathcal{E}_r - 5 \cdot \|\rho\|_{\nu,3}^2$.*

Proof. Combining Lemma 14 with Theorem 6 we obtain:

$$\begin{aligned}
\mathcal{E}_{r'} - \mathcal{E}_r &\geq \sum_{e=(u,v)} \nu_e \rho_e^2 \left(1 - \frac{r_e}{r'_e}\right) \\
&\geq \sum_{e=(u,v)} \nu_e \rho_e^2 (-4\delta \rho_e - \kappa_e) \\
&= -4\delta \cdot \sum_{e=(u,v)} \nu_e \rho_e^3 - \sum_{e=(u,v)} \nu_e \rho_e^2 \kappa_e \\
&= -4\delta \cdot \|\rho\|_{\nu,3}^3 - \sum_{e=(u,v)} \sqrt{\nu_e} \rho_e^2 \cdot \sqrt{\nu_e} \kappa_e \\
&\geq -4 \frac{\|\rho\|_{\nu,3}^3}{\|\rho\|_{\nu,4}} - \sqrt{\left(\sum_{e=(u,v)} \nu_e \rho_e^4\right) \left(\sum_{e=(u,v)} \nu_e \kappa_e^2\right)} \quad (\text{by Cauchy-Schwarz}) \\
&\geq -4 \|\rho\|_{\nu,3}^2 - \sqrt{\|\rho\|_{\nu,4}^4 \cdot \|\kappa\|_{\nu,2}^2} \\
&= -4 \|\rho\|_{\nu,3}^2 - \|\rho\|_{\nu,4}^2 \quad (\text{using } \|\kappa\|_{\nu,2}^2 \leq 1) \\
&\geq -5 \|\rho\|_{\nu,3}^2
\end{aligned}$$

With this tool in hand we can now upper bound the total energy increase caused by perturbations. \square

Lemma 18. *The total energy increase due to perturbations is at most $16c_T c_\rho^2 \cdot m^{3/2-3\eta}$. Furthermore, the total measure always satisfies $\|\nu\|_1 \leq 3m$, i.e. Invariant 1 is preserved.⁴*

Proof. We start by introducing some notation. Let \mathcal{E}^t and ν^t be the energy, respectively the vectore of mesures at the end of the t^{th} iteration. Also, let Δ^t be the total amount of energy increases during that iteration.⁵

Note that, since we only perform progress steps when $\|\rho\|_{\nu,3} \leq c_\rho m^{1/2-\eta}$, one progress step decreases energy by at most $5 \cdot c_\rho^2 m^{1-2\eta}$, according to Lemma 17. Therefore the amount by which energy increases during an iteration can be bounded by

$$\Delta^t \leq \mathcal{E}^t - \mathcal{E}^{t-1} + 5c_\rho^2 \cdot m^{1-2\eta} \quad (5.4)$$

At any point, the energy is capped by the total measure (Lemma 8). Therefore

$$\mathcal{E}^t \leq \|\nu^t\|_1 \quad (5.5)$$

Also, using Corollary 16 we get that every increase in energy by Δ^t increases the total measure by at most $2C \cdot \Delta^t$. Hence

$$\|\nu^t\|_1 \leq \|\nu^{t-1}\|_1 + 2C \cdot \Delta^t \quad (5.6)$$

⁴While the constants provided here are worse than those seen in the proof, we will use these loose bounds in order to accomodate some future changes in the algorithm.

⁵Remember that energy can decrease during progress steps, as per Lemma 17; Δ^t measures the total amount of all increases, without accounting for the lost energy due to progress steps.

Using (5.4) and summing over all $T = c_T \cdot m^{1/2-\eta}$ iterations of the algorithm we obtain:

$$\begin{aligned}
\sum_{t=1}^T \Delta^t &\leq \sum_{t=1}^T (\mathcal{E}^t - \mathcal{E}^{t-1} + 5c_\rho^2 \cdot m^{1-2\eta}) \\
&\leq \mathcal{E}^T + T \cdot 5c_\rho^2 \cdot m^{1-2\eta} \\
&\leq \|\nu^T\|_1 + T \cdot 5c_\rho^2 \cdot m^{1-2\eta} \\
&\leq \left(\|\nu^0\|_1 + \sum_{t=1}^T 2C \cdot \Delta^t \right) + T \cdot 5c_\rho^2 \cdot m^{1-2\eta} \\
&= m + 2C \cdot \sum_{t=1}^T \Delta^t + T \cdot 5c_\rho^2 \cdot m^{1-2\eta}
\end{aligned} \tag{5.7}$$

where for the last two inequalities we applied (5.5) and (5.6).

Hence we obtain

$$\sum_{t=1}^T \Delta^t \leq \frac{1}{1-2C} (m + T \cdot 5c_\rho^2 \cdot m^{1-2\eta}) \leq 2 \cdot c_T m^{1/2-\eta} \cdot 5c_\rho^2 m^{1-2\eta} = 10c_T c_\rho^2 \cdot m^{3/2-3\eta} \tag{5.8}$$

and the measure increase is upper bounded by

$$2C \cdot \left(\sum_{t=1}^T \Delta^t \right) = 2C \cdot 10c_T c_\rho^2 \cdot m^{3/2-3\eta} = m \tag{5.9}$$

So Invariant 1 is satisfied. \square

5.2.2 Bounding the Number of Perturbations

We have just seen that the energy increase suffered due to perturbations is $O(c_T c_\rho^2 \cdot m^{3/2-3\eta})$, which should intuitively enable us to bound the number of perturbations, and thus wrap up the analysis of the algorithm. The reason is that whenever we have to perturb the problem, the ℓ_3 norm of the congestion vector is large (i.e. $\|\rho\|_{\nu,3} \geq c_\rho \cdot m^{1/2-\eta}$), so the energy of the system is large: $\mathcal{E} = \|\rho\|_{\nu,2}^2 \geq \|\rho\|_{\nu,3}^2 \geq c_\rho^2 \cdot m^{1-2\eta}$. Since perturbable arcs are highly congested (see Definition 11), we expect them to contribute most of the energy. This feature of perturbable args is highlighted by the following invariant:

Invariant 2. Whenever we perform a perturbation, $\|\rho_S\|_{\nu,2}^2 \geq \frac{1}{2} c_\rho \cdot m^{1-2\eta}$.

This guarantees that every perturbation increases energy by at least $\Omega(m^{1-2\eta})$, which automatically implies that the number of perturbations is bounded by $O(c_T c_\rho^2 \cdot m^{1/2-\eta})$. Indeed, as seen in Lemma 15, with every perturbation energy increases by $\frac{1}{2} \|\rho_S\|_{\nu,2}^2$. Therefore, assuming Invariant 2, we get that each perturbation increases energy by at least $\frac{1}{4} c_\rho \cdot m^{1-2\eta}$. Since we know from Lemma 18 that total energy increase is bounded by $16c_T c_\rho^2 \cdot m^{3/2-3\eta}$, we get that the number of perturbations performed during the execution of the algorithm is at most $64c_T c_\rho \cdot m^{1/2-\eta}$.

Enforcing the validity of this invariant will be done in Section 6, where we introduce a preconditioning technique which enables us to gain more control over the behavior of electrical flows.

Hence we have proved the following Lemma:

Lemma 19. *Assuming Invariant 2 is valid, the number of perturbations is at most $64c_T c_\rho \cdot m^{1/2-\eta}$.*

This immediately concludes the running time analysis. Indeed, both progress steps and perturbations can be implemented in $\tilde{O}(m)$ time by computing electrical flows using a fast Laplacian solver (see Theorem 1). The number of progress steps is precisely $c_T \cdot m^{1/2-\eta}$, since this is hard coded in the description of the algorithm. Also, the number of perturbations is $O(c_T c_\rho \cdot m^{1/2-\eta})$, according to Lemma 19. Therefore the total running time is $\tilde{O}(c_T c_\rho \cdot m^{3/2-\eta})$.

Theorem 20. *Assuming Invariant 2 is valid, we can produce an exact solution to the unit-capacitated minimum cost σ -flow problem in $\tilde{O}(m^{10/7} \log^{4/3} W)$ time.*

Proof. The proof is similar to the one for Theorem 9.

The algorithm performs a progress step only when $\|\rho\|_{\nu,4} \leq \|\rho\|_{\nu,3} \leq m^{1/2-\eta}$, therefore $\hat{\mu}$ decreases by a factor of $1 - \frac{1}{c_\rho \cdot m^{1/2-\eta}}$ with every iteration. Therefore, in $c_\rho m^{1/2-\eta} (2 \log m + \log \tilde{W}) \leq c_T \cdot m^{1/2-\eta}$ iterations we reduce $\hat{\mu}$ to $O(m^{-3})$, and by Invariant 1 the duality gap of this solution will be $\sum_e \nu_e \hat{\mu} \leq m^{-2}$. Each of the $c_T \cdot m^{1/2-\eta}$ progress steps requires $\tilde{O}(1)$ electrical flow computations, and each of them can be implemented in near-linear time, according to Theorem 1. Furthermore, assuming Invariant 2, we have that the number of perturbations is at most $64c_T c_\rho \cdot m^{1/2-\eta}$, by Lemma 19. Similarly, each perturbation can be implemented in nearly-linear time. Therefore the total running time required for obtaining a duality gap of m^{-2} is $\tilde{O}(c_T m^{3/2-\eta} + c_T c_\rho \cdot m^{3/2-\eta}) = \tilde{O}(m^{10/7} \log^{4/3} W)$. Then, using the repairing algorithm from Section 7, we can round the solution to an optimal one in nearly-linear time. So the total time is $\tilde{O}(m^{10/7} \log^{4/3} W)$. \square

One should note that the $\tilde{O}(m^{10/7} \log^{4/3} W)$ running time can be reduced to $\tilde{O}(m^{10/7} \log W)$ by employing the scaling technique of [9]. Thus, we can reduce our problem to solving $O(\log W)$ instances of our problem where the costs are polynomially bounded. This enables us to change the poly $\log W$ factors from the running time to poly $\log n$ (and thus have them absorbed by the \tilde{O} notation) at the cost of paying only an extra factor of $\log W$.

However, ensuring that Invariant 2 always holds is a bit more subtle. Obtaining a provable guarantee will actually be done by adding a preconditioner, which is carefully analyzed in Section 6.

6 Preconditioning the Graph

Our analysis from the previous section was crucially relying on the assumption that perturbable arcs contribute most of the energy. Unfortunately, this assumption is not always valid. To cope with this problem, we develop a modification of our algorithm that ensures that this assumption holds after all. Roughly speaking, we achieve that by an appropriate preconditioning of our graph. This preconditioning is based on augmenting the graph with additional, auxiliary edges which make the computed electrical flows better behaved. These edges *should not* be thought of as being part of the graph we are computing our σ -flows on. Their sole effect is to guide the electrical flow computation in order to obtain a better electrical flow in the original graph at the cost of slightly changing the demand we are routing.

These edges achieve the optimal trade-off between providing good connectivity in the augmented graph (which lowers the congestion of arcs with low residual capacity) and preventing too much flow from going through them (because of their sufficiently high resistance).

One difficulty posed by this scheme is that we need to control how much the routed demand gets modified. This is easily controlled by setting the resistances of the auxiliary edges to be sufficiently high; in contrast, the magnitude of these resistances needs to be traded against the effect they have on the computed electrical flow. At the end, we fix the demand using a combinatorial procedure (see Section 7) whose running time is

proportional to the ℓ_1 difference between the initial demand and the one routed by the algorithm. Therefore we need to simultaneously ensure that preconditioner edges have sufficiently high resistance such that the change in demand is not significant, and guarantee that the graph has good enough connectivity for Invariant 2 to hold. This trade-off will ultimately determine the choice of the parameter $\eta = 1/14$.

6.1 Using Auxiliary Edges for Electrical Flow Computations

In order to properly describe preconditioning, we need to partition the iterations of the algorithm into *phases* (each of them consisting of a sequence of $m^{2\eta}$ iterations), and show that a finer version of Lemma 18 holds, for each of these phases. The reason is that the resistances on the auxiliary edges need to depend on the set of measures in the graph. But measures increase over time, so the resistances need to be updated accordingly. One should be careful about this aspect, since changing the resistances of auxiliary edges during every iteration of the algorithm would destroy the potential based argument we described in Section 5. Therefore, instead of adjusting the resistances every iteration, we do this only at the beginning of a phase. Over the course of a phase, measures can not change too much, so the preconditioning argument will still be valid.

Definition 21. Partition the $c_T \cdot m^{1/2-\eta}$ iterations of the algorithm into consecutive blocks of size $m^{2\eta}$. Such a block is called a phase. Hence the algorithm consists of $c_T \cdot m^{1/2-3\eta}$ phases.

Preconditioning consists of adding an extra vertex v_0 along with undirected edges (v_0, v) for each vertex $v \in P$ (recall that vertices in P correspond to vertices from only one side of the bipartition in the b -matching instance). We will call these newly added edges *auxiliary edges*. Each of these auxiliary edges will have resistance set to

$$r_{v_0v} = \frac{m^{1+2\eta}}{a(v)} \quad (6.1)$$

where we define

$$a(v) = \sum_{u \in Q: e=(v,u) \in E} \nu_e + \nu_{\bar{e}} \quad (6.2)$$

Recall that by \bar{e} we denote the partner arc of e (introduced in Definition 10), and that the quantities above are defined with respect to the measures existing at the beginning of the phase.

Also, remember that these auxiliary edges exist only in order to provide a mildly different demand for which electrical flows are better behaved. Once we are done perturbing, we perform a progress step on the graph without auxiliary edges, but with the modified demand (i.e. the one that gets routed on the original graph, after removing auxiliary edges).

The inclusion of auxiliary edges in the electrical flow computations requires the contribution of these edges to the energy of the system to be included in the potential based argument from Theorem 18, when analyzing the preconditioned steps.⁶

Remark 22. Even though we include additional edges for electrical flow computations, the energy bound from Lemma 8 still holds (since including additional edges can only decrease energy).

This motivates partitioning the iterations into phases, since changing the resistances of auxiliary edges too often could potentially make the energy vary wildly.

The following lemma shows that we can individually bound the energy and measure increase over any single phase. What is crucial about the new proof is that it does not require any control over how energy changes between iterations belonging to different phases. Therefore, resetting the resistances of auxiliary edges at the beginning of a phase will have no effect on the result described in Lemma 18.

⁶The ℓ_3 norm of the congestion vector will still be measured only with respect to the arcs in the original graph.

The precise statement concerning energy and measure increase during a phase is summarized in the following lemma, whose proof we defer to Appendix D.

Lemma 23. *During a single phase, the total energy increase due to perturbations is at most $16c_\rho^2 \cdot m$. Furthermore, the measure increase during a single phase is at most $\frac{2}{c_T} \cdot m^{1/2+3\eta}$. Also, the total measure always satisfies $\|\nu\|_1 \leq 3m$, i.e. Invariant 1 is preserved.*

It immediately follows that this is simply a refinement of Lemma 18:

Corollary 24. *The total energy increase due to perturbations is at most $16c_T c_\rho^2 \cdot m^{3/2-3\eta}$, and the result described in Lemma 18 is still valid.*

The new version of the algorithm which includes the effect of the auxiliary edges is described below.

Algorithm 2 Perturbed interior-point method with preconditioning edges (parameters: $c_\rho = 400\sqrt{3} \cdot \log^{1/3} \tilde{W}$, $c_T = 3c_\rho \log \tilde{W}$)

1. initialize primal and dual variables (f, y) (as shown in Section 4)
 2. repeat for $c_T \cdot m^{1/2-3\eta}$ phases
 3. reset auxiliary edge resistances (as described in (6.1))
 4. repeat $m^{2\eta}$ times
 5. while $\|\rho\|_{\nu,3} > c_\rho \cdot m^{1/2-\eta}$
 6. perform perturbation (as shown in Section 5.1)
 7. perform progress steps on the original graph (as shown in Section 4.2)
-

Before proving that this version of the algorithm forces Invariant 2 to stay valid, we first bound the change in demand caused by the auxiliary edges. We first show that, due to Invariant 1, the total amount of flow that gets routed electrically through auxiliary edges is small.

Proposition 25. *Let \mathcal{P} be the set of auxiliary edges. Then the total amount of electrical flow on these edges during any progress step satisfies $\|\hat{f}_{\mathcal{P}}\|_1 \leq 5 \cdot m^{1/2-\eta}$.*

Proof. From Invariant 1 and Lemma 8 we have that the total energy satisfies $\mathcal{E} \leq 3m$. The energy contributed by the auxiliary edges is

$$\sum_{(v_0,v) \in \mathcal{P}} r_{v_0v} \cdot f_{v_0v}^2 = \sum_{e \in \mathcal{P}} \frac{m^{1+2\eta}}{a(v)} \cdot f_{v_0v}^2 \leq 3m \quad (6.3)$$

Applying Cauchy-Schwarz we obtain a bound for the ℓ_1 norm of $f_{\mathcal{P}}$:

$$\begin{aligned}
\|f_{\mathcal{P}}\|_1 &= \sum_{(v_0, v) \in \mathcal{P}} \sqrt{a(v)} \cdot \frac{|f_{v_0 v}|}{\sqrt{a(v)}} \\
&\leq \sqrt{\left(\sum_{(v_0, v) \in \mathcal{P}} a(v) \right) \left(\sum_{(v_0, v) \in \mathcal{P}} \frac{f_{v_0 v}^2}{a(v)} \right)} \\
&\leq \sqrt{2 \|\nu\|_1 \cdot \frac{3m}{m^{1+2\eta}}} \\
&\leq \sqrt{2 \cdot 3m \cdot \frac{3m}{m^{1+2\eta}}} \\
&\leq 5 \cdot m^{1/2-\eta}
\end{aligned}$$

We used the fact that summing over all $a(v)$'s we obtain precisely twice the total measure, since each measure on an arc gets counted exactly twice. Then we used Invariant 1, and (6.3). \square

This proposition shows that the demand routed within a progress step is off by at most $5m^{1/2-\eta}$ from the demand routed by the iterate f at that point. Using this fact, we can show that the flow obtained in the end routes a demand that is off by at most $c_T \cdot m^{1/2-\eta}$ from the original demand.

Lemma 26. *Consider the last flow iterate f^T , and let σ^T be the demand routed by this flow. Then the difference between σ^T and the original demand σ satisfies $\|\sigma^T - \sigma\|_1 \leq c_T \cdot m^{1/2-\eta}$.*

Proof. We show by induction that after t iterations, the demand routed by f^t satisfies $\|\sigma^t - \sigma\|_1 \leq t$. The base case is $t = 0$ where $\sigma^0 = \sigma$, and the hypothesis holds. Now let us show that if the induction hypothesis holds after $t - 1$ iterations, then it also holds after t . By Proposition 25 we have that the progress step first produces an electrical flow \hat{f}^t which routes a demand $\hat{\sigma}^t$ on the original graph satisfying $\|\hat{\sigma}^t - \sigma^{t-1}\|_1 \leq 5m^{1/2-\eta}$. Therefore, noting that the flow gets updated by setting it to $(1 - \delta)f^t + \delta\hat{f}^t$ for $\delta \leq \frac{1}{8\|\rho\|_{\nu,4}}$ (see Section 4.2), and using the fact that progress steps are done only when $\|\rho\|_{\nu,3} \leq c_\rho \cdot m^{1/2-\eta}$ (therefore $\delta \leq \frac{1}{8 \cdot c_\rho \cdot m^{1/2-\eta}}$), we have that the demand σ^t routed by the averaged flow satisfies:

$$\begin{aligned}
\|\sigma^t - \sigma\|_1 &\leq \|\sigma^{t-1} - \sigma\|_1 + \|\sigma^{t-1} - \sigma^t\| \\
&= \|\sigma^{t-1} - \sigma\|_1 + \|\sigma^{t-1} - ((1 - \delta)\sigma^{t-1} + \delta\hat{\sigma}^t)\|_1 \\
&= \|\sigma^{t-1} - \sigma\|_1 + \delta \|\sigma^{t-1} - \hat{\sigma}^t\|_1 \\
&\leq t - 1 + \delta \cdot 5m^{1/2-\eta} \\
&\leq t - 1 + \frac{1}{8 \cdot c_\rho m^{1/2-\eta}} \cdot 5m^{1/2-\eta} \\
&\leq t - 1 + 1 \\
&= t
\end{aligned}$$

Centering the solution does not change the demand, so the newly obtained flow f^t has the same demand σ^t , which satisfies the bound above. Therefore, after $T = c_T \cdot m^{1/2-\eta}$ iterations, we have that $\|\sigma^T - \sigma\|_1 \leq c_T \cdot m^{1/2-\eta}$. \square

6.2 Proving Invariant 2

We can finally proceed with proving that, for this version of the algorithm, Invariant 2 holds. We do so by upper bounding the ℓ_3 norm of the congestions of unperturbable arcs $\|\rho_{\bar{S}}\|_{\nu,3}$. Showing that this quantity is significantly smaller than $\tilde{O}\left((c_T c_\rho^2)^{1/6} \cdot m^{1/2-\eta}\right)$ whenever we perform a perturbation automatically implies our result; this is because this lower bounds the ℓ_3 norm of congestions of perturbable arcs, and therefore also their energy.

Proposition 27. *Let \bar{S} be the set of unperturbable arcs. If, whenever we perform a perturbation, $\|\rho_{\bar{S}}\|_{\nu,3} \leq 10(c_T c_\rho^2)^{1/6} \cdot m^{1/2-\eta}$, then Invariant 2 holds.*

Proof. Recall that when we perform a perturbation, we have $\|\rho\|_{\nu,3} \geq c_\rho \cdot m^{1/2-\eta}$. If $\|\rho_{\bar{S}}\|_{\nu,3} \leq 10(c_T c_\rho^2)^{1/6} \cdot m^{1/2-\eta}$, then

$$\begin{aligned}
\|\rho_{\bar{S}}\|_{\nu,3} &\geq c_\rho \cdot m^{1/2-\eta} - 10(c_T c_\rho^2)^{1/6} \cdot m^{1/2-\eta} \\
&= c_\rho \cdot m^{1/2-\eta} - 10\left(3 \log \tilde{W} \cdot c_\rho^3\right)^{1/6} \cdot m^{1/2-\eta} \\
&= c_\rho \cdot m^{1/2-\eta} - 10\left(3 \log \tilde{W}\right)^{1/6} \cdot c_\rho^{1/2} \cdot m^{1/2-\eta} \\
&= (400\sqrt{3} \log^{1/3} \tilde{W}) m^{1/2-\eta} - 10 \cdot 3^{1/6} \cdot \log^{1/6} \tilde{W} \cdot (400\sqrt{3} \log^{1/3} \tilde{W})^{1/2} \cdot m^{1/2-\eta} \\
&= 200\sqrt{3} \log^{1/3} \tilde{W} \cdot m^{1/2-\eta} \\
&= \frac{1}{2} c_\rho \cdot m^{1/2-\eta}
\end{aligned}$$

Hence $\|\rho_{\bar{S}}\|_{\nu,2}^2 \geq \|\rho_{\bar{S}}\|_{\nu,3}^2 \geq \frac{1}{2} c_\rho \cdot m^{1-2\eta}$. □

Upper bounding the ℓ_3 norm of ρ on unperturbable edges is done by partitioning them into sets, and separately bounding their ℓ_3^3 norms. As a matter of fact, all the work we have to do concerns arcs with congestions within the range $[c_\rho^3 \cdot m^{1/2-3\eta}, \sqrt{3} \cdot m^{1/2}]$, since the energy upper bound enforced by Invariant 2 controls the maximum congestion, while those with lower congestion immediately satisfy the required ℓ_3 norm bound.

Lemma 28. $\|\rho_{\bar{S}}\|_{\nu,3} \leq 10(c_T c_\rho^2)^{1/6} \cdot m^{1/2-\eta}$.

Proving Lemma 28 needs a careful analysis based on bounding the contributions from edges at different scales. Doing so requires extending the analysis from [24] in a slightly more complicated fashion. That analysis looks at the line embedding of graph vertices given by their potentials, and separately upper bounds the energy contributed by sets of arcs according to how stretched each of them is in the embedding. One particular obstacle posed by our setup is that, while the analysis crucially relies on the existence of auxiliary arcs, in our case there are no auxiliary arcs connected to vertices in Q . This makes it difficult to prove statements about the amount of energy on partner arcs (u, v) and (\bar{u}, v) (connected to the same vertex v in Q), since the auxiliary edges only control how far apart in the embedding u and \bar{u} are. Unfortunately, they do not immediately say anything about v , which could potentially be very far from both u and \bar{u} , and thus the two partner arcs could contribute a lot of energy.

We will see that our desired bound still holds. Our proof technique relies on decomposing the electrical flow into a sum of two electrical flows $\hat{f}^{(P)}$ and $\hat{f}^{(Q)}$, one of which can be bounded directly and the other of which has no net flow through any arc in Q .

We additionally express $f = f^{(P)} + \hat{f}^{(Q)}$, where we define

$$f^{(P)} = f - \hat{f}^{(Q)} \tag{6.4}$$

The following lemma, whose proof can be found in Appendix B, states the existence and properties of such electrical flows.

Lemma 29. *Let \hat{f} be an electrical flow in the graph with auxiliary edges. There exist electrical flows $\hat{f}^{(P)}$ and $\hat{f}^{(Q)}$ such that:*

1. $\hat{f}^{(P)}$ has no demand on vertices in $Q \cup \{v_0\}$.
2. For any pair of partner edges e and \bar{e} , and writing $f^{(P)} = f - \hat{f}^{(Q)}$, we have $r_e \left(f_e^{(P)}\right)^2 + r_{\bar{e}} \left(f_{\bar{e}}^{(P)}\right)^2 \leq \nu_e + \nu_{\bar{e}}$ and $r_e \left(\hat{f}_e^{(Q)}\right)^2 + r_{\bar{e}} \left(\hat{f}_{\bar{e}}^{(Q)}\right)^2 \leq \nu_e + \nu_{\bar{e}}$.

Now, we additionally define $\rho_e^{(P)} = \frac{|\hat{f}_e^{(P)}|}{f_e}$ and $\rho_e^{(Q)} = \frac{|\hat{f}_e^{(Q)}|}{f_e}$. We split $\bar{\mathcal{S}}$ into two subsets:

$$\mathcal{A}_1 = \left\{ e \in \bar{\mathcal{S}} \mid \rho_e^{(P)} \geq \frac{1}{2} \rho_e \right\}$$

and

$$\mathcal{A}_2 = \left\{ e \in \bar{\mathcal{S}} \mid \rho_e^{(P)} < \frac{1}{2} \rho_e \right\}_2$$

for which we bound the contributions to the ℓ_3 norm separately. Note that for any edge e in \mathcal{A}_2 , $\rho_e^{(Q)} \geq \frac{1}{2} \rho_e$. This implies that

$$\begin{aligned} \|\rho_{\bar{\mathcal{S}}}\|_{\nu,3}^3 &\leq \|\rho_{\mathcal{A}_1}\|_{\nu,3}^3 + 8 \|\rho_{\mathcal{A}_2}\|_{\nu,3}^3 \\ &\leq \|\rho_{\mathcal{A}_1}\|_{\nu,3}^3 + 8 \|\rho^{(Q)}\|_{\nu,3}^3 \end{aligned}$$

First, we want to bound $\|\rho^{(Q)}\|_{\nu,3}^3$:

Lemma 30. *Suppose that for all pairs of partner edges e and \bar{e} , $\frac{\max(\nu_e, \nu_{\bar{e}})}{\min(\nu_e, \nu_{\bar{e}})} \leq x$. Then $\|\rho^{(Q)}\|_{\nu,3}^3 \leq (3\sqrt{1+x})m$.*

Proof. We look at the contribution of a single pair of partner arcs, e and \bar{e} , to $\|\rho^{(Q)}\|_{\nu,3}^3$. This is equal to

$$\begin{aligned} \nu_e \left| \rho_e^{(Q)} \right|^3 + \nu_{\bar{e}} \left| \rho_{\bar{e}}^{(Q)} \right|^3 &= r_e \left(\hat{f}_e^{(Q)} \right)^2 \rho_e^{(Q)} + r_{\bar{e}} \left(\hat{f}_{\bar{e}}^{(Q)} \right)^2 \rho_{\bar{e}}^{(Q)} \\ &= r_e \left(\hat{f}_e^{(Q)} \right)^2 \sqrt{\frac{r_e \left(\hat{f}_e^{(Q)} \right)^2}{\nu_e} + r_{\bar{e}} \left(\hat{f}_{\bar{e}}^{(Q)} \right)^2} \sqrt{\frac{r_{\bar{e}} \left(\hat{f}_{\bar{e}}^{(Q)} \right)^2}{\nu_{\bar{e}}}} \\ &\leq \sqrt{\frac{\max \left(r_e \left(\hat{f}_e^{(Q)} \right)^2, r_{\bar{e}} \left(\hat{f}_{\bar{e}}^{(Q)} \right)^2 \right)}{\min(\nu_e, \nu_{\bar{e}})}} \left(r_e \left(\hat{f}_e^{(Q)} \right)^2 + r_{\bar{e}} \left(\hat{f}_{\bar{e}}^{(Q)} \right)^2 \right) \\ &\leq \sqrt{\frac{\nu_e + \nu_{\bar{e}}}{\min(\nu_e, \nu_{\bar{e}})}} (\nu_e + \nu_{\bar{e}}) \\ &\leq \sqrt{1+x} (\nu_e + \nu_{\bar{e}}). \end{aligned}$$

Here, we applied property 2 from Lemma 29. Summing over all pairs of partner arcs and using Invariant 1, we get $\|\rho^{(Q)}\|_{\nu,3}^3 \leq (3\sqrt{1+x})m$, as desired. \square

Corollary 31. $\|\rho^{(Q)}\|_{\nu,3}^3 \leq 5m^{5/4-(3/2)\eta}$

Proof. We can apply Lemma 30 with $x = m^{1/2-3\eta}$, since by definition we never perturb arcs with $\frac{1}{f_e} > m^{1/2-3\eta}$. So $\|\rho^{(Q)}\|_{\nu,3}^3 \leq \sqrt{1+m^{1/2-3\eta}} \cdot 3m \leq 5m^{5/4-(3/2)\eta}$. \square

It remains to bound $\|\rho_{\mathcal{A}_1}\|_{\nu,3}^3$. First, note that for every edge e in \mathcal{A}_1 , if $\rho_e > \sqrt{40c_T c_\rho^2} \cdot m^{1/2-3\eta}$, we have $r_e = \frac{\nu_e}{f_e^2}$, implying that

$$\begin{aligned} r_e \left| \hat{f}_e \right| &= \frac{\nu_e}{f_e} \rho_e \\ &\geq \frac{1}{f_e} \rho_e \\ &\geq m^{1/2-3\eta} \rho_e. \end{aligned}$$

By the definition of \mathcal{A}_1 , $r_e \left| \hat{f}_e^{(P)} \right|$ is at least half of $\frac{1}{2} m^{1/2-3\eta} \rho_e$. Similarly, $r_e \left(\hat{f}_e^{(P)} \right)^2 \geq \frac{1}{4} r_e f_e^2$. Thus for any choice of a threshold $T > \sqrt{40c_T c_\rho^2} \cdot m^{1/2-3\eta}$ on the ρ , we have

$$\sum_{e \in \mathcal{A}_1, \rho_e \geq T} r_e f_e^2 \leq 4 \sum_{r_e \left| \hat{f}_e^{(P)} \right| \geq \frac{1}{2} m^{1/2-3\eta} T} r_e \left(\hat{f}_e^{(P)} \right)^2.$$

Next, we consider a ‘‘quotient’’ or ‘‘Schur complement’’ electrical network N on the vertices of $P \cup \{v_0\}$ only, replacing the pair of arcs (u, v) and (\bar{u}, v) with one edge (u, \bar{u}) with resistance $r_{uv} + r_{\bar{u}v}$. We consider this edge to have measure $\nu_{uv} + \nu_{\bar{u}v}$. Note that this is in some sense undoing the b -matching reduction from Lemma 2.

Because they have 0 net flow on all vertices in Q , we can map flows $f^{(P)}$ and $\hat{f}^{(P)}$ to flows in N , $f^{(N)}$ and $\hat{f}^{(N)}$, by setting the flow from u to \bar{u} in the new flow to that from u to v (or equivalently, by flow conservation, v to \bar{u}) in the original flow. This mapping preserves the demands on the vertices in Q , so $f^{(N)}$ still satisfies the same demands as $\hat{f}^{(N)}$. It also preserves the electrical nature of $\hat{f}^{(N)}$, since $f^{(N)}$ can be induced by the same voltages. Finally, since the resistance on $e = (u, \bar{u})$ is set to $r_{uv} + r_{\bar{u}v}$, we have $r_e \left(f_e^{(N)} \right)^2 = r_{uv} \left(f_{uv}^{(P)} \right)^2 + r_{\bar{u}v} \left(f_{\bar{u}v}^{(P)} \right)^2$, and in particular for any edge e in N we have $r_e \left(f_e^{(N)} \right)^2 \leq \nu_e$.

Now we may apply our general preconditioning result, which bounds the total energy of edges with high voltage in terms of the resistance of the preconditioning edges:

Lemma 32. *Let N be an electrical network, on a set of vertices P plus a special vertex v_0 , with each edge e not incident to v_0 assigned a measure ν_e . Let ν' be another assignment of measures such that $\nu' \leq \nu$, with the total ‘‘missing measure’’ $\sum_e \nu_e - \nu'_e = M$, and let each vertex v be connected to v_0 by an edge with resistance $\frac{R}{a(v)}$, where $a(v)$ is the sum of ν'_e over edges e incident to v . Let $f^{(N)}$ be a flow on N , with no flow on any edge incident to v_0 and with $r_e \left(f_e^{(N)} \right)^2 \leq \nu_e$ for every edge not incident to v_0 , and let $\hat{f}^{(N)}$ be the electrical flow on N satisfying the same demands as $f^{(N)}$. Then*

$$\sum_{r_e \left| \hat{f}_e^{(N)} \right| \geq V} r_e \left(\hat{f}_e^{(N)} \right)^2 \leq \frac{32R \sum_e \nu_e}{V^2} + 2M.$$

This is proved in Appendix D. Here, we set ν' to the measures from the beginning of the preconditioning phase; by Lemma 23 the missing measure $M \leq \frac{2}{c_T} \cdot m^{1/2+3\eta}$. The preconditioning edges were weighted with

$R = m^{1+2\eta}$. This implies that for $T > \sqrt{40c_Tc_\rho^2} \cdot m^{1/2-3\eta}$,

$$\sum_{e \in \mathcal{A}_1, \rho_e \geq T} r_e f_e^2 \leq 4 \cdot \left(\frac{32 \cdot m^{1+2\eta} \cdot \sum_e \nu_e}{\left(\frac{1}{2}m^{1/2-3\eta} \cdot T\right)^2} + \frac{2}{c_T} \cdot m^{1/2+3\eta} \right) \leq \frac{512 \cdot m^{8\eta} \cdot 3m}{T^2} + \frac{8}{c_T} \cdot m^{1/2+3\eta}.$$

On the other hand, for all T we trivially have $\sum_{e \in \mathcal{A}_1, \rho_e \geq T} r_e f_e^2 \leq 3m$, by Invariant 1 combined with Lemma 8. Now, we can write

$$\begin{aligned} \|\rho_{\mathcal{A}_1}\|_{\nu,3}^3 &= 3 \int_0^{\sqrt{3} \cdot m^{1/2}} \left(\sum_{e \in \mathcal{A}_3, \rho_e \geq T} r_e f_e^2 \right) dT \\ &\leq 3 \left(\int_0^{\sqrt{40c_Tc_\rho^2} \cdot m^{1/2-3\eta}} 3m dT \right) + 3 \left(\int_{m^{1/2-3\eta}}^{\sqrt{3} \cdot m^{1/2}} \left(1536 \cdot \frac{m^{1+8\eta}}{T^2} + \frac{8}{c_T} \cdot m^{1/2+3\eta} \right) dT \right) \\ &\leq 9\sqrt{40c_Tc_\rho^2} \cdot m^{3/2-3\eta} + 42m^{1+3\eta} + 3 \int_{m^{1/2-3\eta}}^\infty \left(1536 \frac{m^{1+8\eta}}{T^2} \right) dT \\ &\leq 9\sqrt{40c_Tc_\rho^2} \cdot m^{3/2-3\eta} + 42m^{1+3\eta} + 4608m^{1/2+11\eta} \\ &\leq 9\sqrt{40c_Tc_\rho^2} \cdot m^{3/2-3\eta} + 42m^{1+3\eta} + 4608m^{1/2+11\eta}. \end{aligned}$$

This determines our choice of $\eta = \frac{1}{14}$ —it ensures that $m^{1/2+11\eta} = m^{3/2-3\eta}$ in the last term, which is essentially the bound we need. Plugging that in it also ensures that $m^{1+3\eta}$ in the second term is less than $m^{3/2-3\eta}$. Finally, we had $\|\rho_{\bar{S}}\|_{\nu,3}^3 \leq \|\rho_{\mathcal{A}_1}\|_{\nu,3}^3 + 8\|\rho^{(Q)}\|_{\nu,3}^3$, with $\|\rho^{(Q)}\| \leq 5m^{5/4-(3/2)\eta}$; with $\eta = \frac{1}{14}$ we also have $m^{5/4-(3/2)\eta} \leq m^{3/2-3\eta}$. Thus we have $\|\rho_{\bar{S}}\|_{\nu,3}^3 \leq 1000\sqrt{c_Tc_\rho^2}m^{3/2-3\eta}$, or $\|\rho_{\bar{S}}\|_{\nu,3} \leq 10(c_Tc_\rho^2)^{1/6} \cdot m^{1/2-\eta}$, proving Lemma 28.

7 Repairing the Matching

In this section we assume the b -matching view on σ -flows. Hence, summarizing Theorem 20, Lemma 26 and the proof of Invariant 2 in Section 6 we obtain the following result.

Theorem 33. *Consider a ternary instance $G = (V, E, c)$ of the weighted perfect bipartite b -matching problem when $\|b\|_1 = O(m)$. In $\tilde{O}(m^{10/7} \log^{4/3} W)$ time we can either conclude that G does not have a perfect b -matching or return a primal-dual pair with duality gap at most m^{-2} to the perfect b^+ -matching problem, where $\|b^+ - b\|_1 \leq c_T m^{3/7}$.*

In this section we will be “repairing” the feasible primal-dual solution given by the above theorem. Our repair procedure will consist out of four steps. In the first step – Lemma 35, we will reduce the duality gap to 0. Next, in Lemma 37 we will round the solution to be integral. In the third step, we will repair the perturbations done to demands in Theorem 40. Finally, we will use scaling to reduce the dependence on $\log W$.

The main tool we are going to use is the directed version of the graph bipartite graph that encodes alternating paths with respect to the current b -matching. Let $G = (V_1 \cup V_2, E, c)$ be a bipartite weighted graph in which we want to find a minimum weight perfect b -matching. Given some fractional b -matching x we define $\vec{G}_x = (V_1 \cup V_2, \vec{E}_x, \vec{c}_x)$ to be a directed version of the graph G where all edges are directed from V_1 to V_2 and additionally we add edges of x that are directed from V_2 to V_1 (edges of x have two copies in both

directions). Moreover, the weights of edges in x are negated. Formally,

$$\begin{aligned}\vec{E}_x &= \{(u, v) | uv \in E, u \in V_1, v \in V_2\} \cup \{(u, v) | x_{uv} \neq 0, u \in V_2, v \in V_1\}, \\ \vec{c}_x(u, v) &= \begin{cases} c_{uv} & \text{if } u \in V_1, v \in V_2, \\ -c_{uv} & \text{if } u \in V_2, v \in V_1. \end{cases}\end{aligned}$$

We observe a path in \vec{G}_x correspond to alternating paths in G with respect to x . Let F_x denote the set of vertices whose demand is not fully satisfied, i.e., $F_x = \{v \in V | x(v) < b_v\}$. We now observe that if a path π starts in $V_1 \cap F_x$ and ends in $V_2 \cap F_x$ that it is an augmenting path with respect to x and can be used to enlarge the b -matching x . When the b -matching x is integral we can interpret it as a multiset of edges which we denote by M . In such case we will use \vec{G}_M to denote \vec{G}_x . The important property of \vec{G}_x is that given an optimal primal solution x it allows to find optimal dual solution. Let $\mathcal{D}_{\vec{G}_M}(V_1, u)$ denote the distance from V_1 to u in \vec{G}_M . The following property was first observed by Iri [13].

Lemma 34. *Consider a ternary instance $G = (V, E, c)$ of the weighted perfect bipartite b -matching problem if x is an optimal b -matching then optimal dual solution y is given as:*

$$y_v = \begin{cases} -\mathcal{D}_{\vec{G}_M}(V_1, v) & \text{if } v \in V_1, \\ \mathcal{D}_{\vec{G}_M}(V_1, v) & \text{if } v \in V_2. \end{cases}$$

Proof. First, observe that \vec{G}_x does not contain negative length cycles by optimality of x , so the above distances are well defined. Second, observe that y is feasible as for each $uv \in E$, where $u \in V_1$ and $v \in V_2$ we have $\mathcal{D}_{\vec{G}_M}(V_1, v) \leq \mathcal{D}_{\vec{G}_M}(V_1, u) + c_{uv}$, so $y_v + y_u \leq c_{uv}$. Third, when $x_{uv} \neq 0$ both edges (u, v) and (v, u) are present in \vec{G}_M . One edge implies feasibility $y_v + y_u \leq c_{uv}$, whereas in the second one all signs are reversed $-y_v - y_u \leq -c_{uv}$. Hence, the equality $y_v + y_u = c_{uv}$ and y is optimal. \square

The above observation will be useful in the following lemma where we will construct optimal primal-dual pair. We will first find an optimal primal solution to the slightly perturbed instance and then compute the corresponding optimal dual.

Lemma 35. *Consider a ternary instance $G = (V, E, c)$ of the weighted perfect bipartite b -matching problem. Given a feasible primal-dual pair (x, y) with duality gap at most m^{-2} in $O(m + n \log n)$ time we can compute an optimal primal-dual pair (x^+, y^+) to the perfect b^+ -matching problem, where $\|b^+ - b\|_1 \leq m^{-2}$.*

Proof. First, we obtain x^+ from x by rounding to 0 all edges that have value of x smaller than m^{-2} , i.e.,

$$x_{uv}^+ = \begin{cases} 0 & \text{if } x_{uv} \leq m^{-2}, \\ x_{uv} & \text{otherwise.} \end{cases}$$

Moreover, we define $b^+(u) = \sum_{uv} x(v)$. Observe that all $uv \in E$ such that $x_{uv}^+ \neq 0$ we need to have $c_{uv} - y_u - y_v \leq m^{-2}$, as otherwise the duality gap would be bigger than m^{-2} . Consider *reduced weights* (or the slack) of edges that are defined as $\tilde{c}_{uv} = c_{uv} - y_u - y_v$. Now define

$$\tilde{c}_{uv}^+ = \begin{cases} 0 & \text{if } \tilde{c}_{uv} \leq m^{-2}, \\ \tilde{c}_{uv} & \text{otherwise.} \end{cases}$$

We denote by $\tilde{\mathcal{D}}_{\vec{G}_M}^+(V_1, v)$ the distances in \vec{G}_x with respect to \tilde{c}_{uv}^+ , whereas by $\tilde{\mathcal{D}}_{\vec{G}_M}(V_1, v)$ the distances with respect to \tilde{c}_{uv} . Observe that \vec{G}_x with weights given by \tilde{c}_{uv}^+ does not contain negative weight edges, so

distances with respect to \tilde{c}_{uv}^+ can be computed in $O(m + n \log n)$ time using Dijkstra's algorithm. Moreover, we observe that

$$\mathcal{D}_{\vec{G}_M}(V_1, v) = \begin{cases} \tilde{\mathcal{D}}_{\vec{G}_M}(V_1, v) + y_v & \text{if } v \in V_1, \\ \tilde{\mathcal{D}}_{\vec{G}_M}(V_1, v) - y_v & \text{if } v \in V_2. \end{cases}$$

We have that $|\tilde{\mathcal{D}}_{\vec{G}_M}(V_1, v) - \tilde{\mathcal{D}}_{\vec{G}_M}^+(V_1, v)| \leq m^{-1}$ and we know that $\mathcal{D}_{\vec{G}_M}(V_1, v)$ is integral, so for

$$\mathcal{D}_{\vec{G}_M}^+(V_1, v) = \begin{cases} \tilde{\mathcal{D}}_{\vec{G}_M}^+(V_1, v) + y_v & \text{if } v \in V_1, \\ \tilde{\mathcal{D}}_{\vec{G}_M}^+(V_1, v) - y_v & \text{if } v \in V_2. \end{cases}$$

we have $\mathcal{D}_{\vec{G}_M}(V_1, v) = \lceil \mathcal{D}_{\vec{G}_M}^+(V_1, v) \rceil$, where $\lceil \cdot \rceil$ is the nearest integer function. \square

Now we are ready to round the above solution – for the rounding step it is essential that primal solution is optimal. The following result that was proven recently in [15] will become handy for us.

Theorem 36. [15] *Let N be a flow network with integral capacities and edge costs. Let f be a flow in N with integral total value then in $O(m \log n)$ one can compute integral flow f' with the same flow value and no worse cost. Moreover, the support of f' is a subset of the support of f .*

The above result is obtained by performing fractional cycle cancelation using dynamic link-cut trees [31].

Lemma 37. *Let x be the optimal primal and y be the optimal dual solution to the b^+ -matching problem, where $\|b^+ - b\|_1 \leq c_T \cdot m^{3/7}$. In $O(m \log n)$ time we can compute an integral $(2c_T \cdot m^{3/7} + 1)$ -near b -matching $M \subseteq E$ such that M is included in the support of x , i.e., $M \subseteq \{e \in E | x_e \neq 0\}$.*

Proof. For graph G let V_1 and V_2 be the bipartition of V . For notational convenience in the remainder of this section we use $V_1 = P$ and $V_2 = Q$. First, for each vertex $v \in V$ we define b^{\leq} to be $b_v^{\leq} = \min(b_v, b_v^+)$. Moreover, we round down $b^{\leq}(V_1)$ and $b^{\leq}(V_2)$ to the nearest integer. Observe that $\|b^+ - b\|_1 \leq 2c_T \cdot m^{3/7} + 1$. Now, for each $x(E(v)) > b_v^{\leq}$ we round down $x(E(v))$ to b_v^{\leq} , i.e., we decrease x on arbitrary edges incident to v , so that the fraction of edges incident to v becomes b_v^{\leq} . Let us denote the resulting vector by x^{\leq} . Now, let us view the b^{\leq} -matching problem as a flow problem by:

- directing edge in $e \in E$ from V_1 to V_2 – the flow on arc e is equal x_e^{\leq} ,
- adding source s and sink t ,
- connecting s to all vertices in $v \in V_1$ – the flow on arc sv is equal to $x^{\leq}(E(v))$,
- connecting all vertices in $v \in V_2$ to t – the flow on arc vt is equal to $x^{\leq}(E(v))$.

By applying Theorem 36 to the above fractional flow we obtain integral flow f' with value $\geq \|b\|_1 / 2 - 2c_T \cdot m^{3/7} - 1$. This flow induces $(2c_T \cdot m^{3/7} + 1)$ -near b -matching M in G . \square

We observe that complementary slackness conditions still hold between M and y (i.e., for each $uv \in M$ we have $c_{uv} = y_u + y_v$) because M is contained in the support of x . We will exploit this fact in the following.

In this moment we have executed two steps of our repair procedure – we have reduced the duality gap to 0 and rounded the solution to be integral. Now, we are ready to repair the perturbations done to demands. Our repair procedure is presented in Algorithm 3. It finds a minimum weight perfect b -matching in G . In the procedure we first apply Theorem 33 then we round the primal solution using Theorem 36. Next, we repeatedly find shortest path in \vec{G}_M from $V_1 \cap F_M$ to $V_2 \cap F_M$ with respect to reduced weights \tilde{c} . *Reduced*

Algorithm 3 Algorithm for computing minimum weight perfect b -matching.

1. Apply Theorem 33 to G and let x and y be the resulting primal and dual solution to the perturbed b^+ .
 2. Apply Theorem 36 to x to obtain $(2c_T \cdot m^{3/7} + 1)$ -near b -matching M .
 3. while M is not a perfect b -matching repeat
 4. (Invariant: for all edges $e \in \vec{G}_M[R_M]$ we have $\tilde{c}_e \geq 0$.)
 5. Construct \vec{G}_M using G, M, \tilde{c} .
 6. Find a shortest path π from $V_1 \cap F_M$ to $V_2 \cap F_M$ in \vec{G}_M .
 7. for all $u \in V_1 \cup V_2$ do
 8. if u is reachable from V_1 in \vec{G}_M then
 9. if $u \in V_1$ then $y_u := y_u - \mathcal{D}_{\vec{G}_M}(V_1, u)$
 10. else $y_u := y_u + \mathcal{D}_{\vec{G}_M}(V_1, u)$
 11. Enlarge M using augmenting path π .
 12. Return M
-

weights (or the slack) are defined as $\tilde{c}_{uv} = c_{uv} - y_u - y_v$. These paths are used to augment the matching. Augmentation of the matching using shortest paths guarantees that M is *extremal*, i.e., M is the minimum weight perfect deg_M -matching. In deg_M -matching the demand for vertex v is equal to $\text{deg}_M(v)$, i.e., number of edges incident to v in M .⁷ The following corollary states that this is true at the beginning of the algorithm.

Corollary 38. *The matching M constructed by Theorem 36 is extremal.*

Proof. Observe that edges of M are tight in G with respect to the corresponding dual solution y , so $\sum_{v \in V(M)} y_v b_v = c(M)$ and this proves that M is optimal deg_M -matching. \square

In order to be able to efficiently find these shortest paths in \vec{G}_M we need to make sure that during the execution of the algorithm reduced weights are nonnegative. At the beginning, by the definition of b -vertex packing, the reduced weights are nonnegative. However, when augmenting the matching using some path π we change direction of some edges on π in \vec{G}_M and flip the sign of their weights. The crucial part is to first reweigh the dual to make sure that all edges on π have reduced cost equal to 0. If this is the case reversing edges does not introduce negative weights. Let us denote by R_M the set of vertices reachable from $V_1 \cap F_M$ in \vec{G}_M and by $\mathcal{D}_{\vec{G}_M}(V_1, u)$ the distance from V_1 to u in \vec{G}_M .

First, we need to prove that the invariant in the while loop of the algorithm holds.

Lemma 39. *During the execution of the while loop in Algorithm 3 for all edges $e \in \vec{G}_M[R_M]$ we have $\tilde{c}_e \geq 0$.*

Proof. First, we observe that set of reachable vertices R_M decreases during the execution of the algorithm. The only step that alters set R_M is the matching augmentation that changes the direction of some edges on π . Hence, R_M would increase when there would be an edge of π entering R_M , but this is impossible as π needs to be contained in R_M by definition.

⁷Similar definition is used in Edmonds-Karp algorithm. However, there extremal matching is defined to be maximum weight matching of given size.

Now, we need to consider the reweighing done in the algorithm. We will prove that when it is done we have that $\tilde{c}_e = 0$ for all $e \in \pi$ and $\tilde{c}_e \geq 0$ for all $e \in \vec{G}_M[R_M]$. In such case augmenting M using π will not introduce negative weights to \vec{G}_M as all weights on π will be zero.

As for the fact that $\tilde{c}_e = 0$ for all $e \in \pi$ we need to consider two cases. Let $uv \in \pi$ where $u \in V_1$ and $v \in V_2$.

- $uv \in M$ – Observe that in \vec{G}_M arc vu enters u and lies on the shortest path π , so we have $\mathcal{D}_{\vec{G}_M}(V_1, u) = \mathcal{D}_{\vec{G}_M}(V_1, v) - \tilde{c}_{uv}$, by definition of \vec{c}_M . By our reweighing rule the new reduced weight of uv is equal to $\tilde{c}(uv) + \mathcal{D}_{\vec{G}_M}(V_1, u) - \mathcal{D}_{\vec{G}_M}(V_1, v) = 0$.
- $uv \notin M$ – As π is the shortest path, for arc uv we have that $\mathcal{D}_{\vec{G}_M}(V_1, v) = \mathcal{D}_{\vec{G}_M}(V_1, u) + \tilde{c}_{uv}$ by definition of \vec{c}_M . This in turn means that the new reduced weight of uv is $\tilde{c}_{uv} + \mathcal{D}_{\vec{G}_M}(V_1, u) - \mathcal{D}_{\vec{G}_M}(V_1, v) = 0$.

Let us now consider edges $e \notin E(\pi)$ but $e \in \vec{G}_M[R_M]$. We have two cases here as well. Let $uv \in E(\pi)$ where $u \in V_1$ and $v \in V_2$.

- $uv \in M$ – We have $\mathcal{D}_{\vec{G}_M}(V_1, u) \leq \mathcal{D}_{\vec{G}_M}(V_1, v) - \tilde{c}_{uv}$ as well as $\mathcal{D}_{\vec{G}_M}(V_1, v) \leq \mathcal{D}_{\vec{G}_M}(V_1, u) + \tilde{c}_{uv}$, because edges of M are bidirected. Hence, $\mathcal{D}_{\vec{G}_M}(V_1, v) = \mathcal{D}_{\vec{G}_M}(V_1, u) + \tilde{c}_{uv}$ and the new weight of uv is $\tilde{c}_{uv} + \mathcal{D}_{\vec{G}_M}(V_1, u) - \mathcal{D}_{\vec{G}_M}(V_1, v) = 0$.
- $uv \notin M$ – By the properties of the distance function we have that $\mathcal{D}_{\vec{G}_M}(V_1, v) \leq \mathcal{D}_{\vec{G}_M}(V_1, u) + \tilde{c}_{uv}$, so the new reduced weight of uv is $\tilde{c}_{uv} + \mathcal{D}_{\vec{G}_M}(V_1, u) - \mathcal{D}_{\vec{G}_M}(V_1, v) \geq 0$.

□

We are now ready to prove the correctness of the algorithm and bound its running time.

Theorem 40. *Assuming $\|b\|_1 = O(m)$, we can find a minimum weight perfect b -matching in $\tilde{O}(m^{10/7} \log^{4/3} W)$ time.*

Proof. We apply Algorithm 3. The correctness of the algorithm follows by the fact that after each augmentation M is extremal. Hence, when M is perfect it needs to be minimum cost perfect b -matching.

The execution of Theorem 33 requires $\tilde{O}(m^{10/7} \log^{4/3} W)$ time. In order to round x to M using Theorem 36 we need $O(m \log n)$ time. Finally, as M is a $\tilde{O}(m^{3/7})$ -near b -matching we will find at most $\tilde{O}(m^{3/7})$ augmenting paths with respect to it. Finding each augmenting path takes $O(m + n \log n)$ time using Dijkstra's algorithm, so this part takes $\tilde{O}(m^{10/7})$ time. Therefore the total running time is $\tilde{O}(m^{10/7} \log^{4/3} W)$. □

Theorem 41. *Assuming $\|b\|_1 = O(m)$, we can find a maximum weight vertex b -packing in $\tilde{O}(m^{10/7} \log^{4/3} W)$ time.*

Proof. One would like to apply Theorem 33 to G' to produce the dual, but Algorithm 3 does not compute a dual solution. The problem is that y in the algorithm is defined only on the set R_M , i.e., nodes reachable in \vec{G}_M from free vertices in V_1 . There is, however, an easy fix to this. Let $v \in V_1 \setminus V(M)$ be arbitrary free vertex in V_1 . We define \vec{G}'_M to be a graph obtained from \vec{G}_M by connecting v with all vertices in V_2 with an edge of cost $n\|w\|_\infty$. Such heavy edges will never be used by a minimum cost perfect 1-matching, but now we always have $R_M = V_1 \cup V_2$. Hence, by the invariant, y computed during the execution of Algorithm 3 needs to form a maximum vertex b -packing. □

Algorithm 4 Scaling algorithm for computing maximum weight b -packing in $G = (V_1 \cup V_2, E, c)$.

1. If $c_{uv} = 0$ for all $uv \in E$ then return $y_v = 0$ for all $v \in V_1 \cup V_2$.
 2. Recursively find the optimal dual solution \bar{y} for a graph $\bar{G} = (V_1 \cup V_2, E, \bar{c})$ where $\bar{c}_{uv} = \lfloor \frac{c_{uv}}{2} \rfloor$ for all $uv \in E$.
 3. Set $y_v = 2\bar{y}_v$ for all $v \in V_1 \cup V_2$.
 4. Set $\tilde{c}_{uv} = c_{uv} - y_u - y_v$ for all $uv \in E$.
 5. Set $\tilde{c}_{uv}^n = \min(\tilde{c}_{uv}, \|b\|_1)$.
 6. Find the optimal dual solution \tilde{y}^n for the graph $\tilde{G}^n = (V_1 \cup V_2, E, \tilde{c}^n)$.
 7. Return $y + \tilde{y}^n$ as the optimal dual solution for G .
-

In order to reduce the dependency on $\log W$ in the running time, we apply the scaling technique of Gabow (see [9]). This enables us to reduce our problem in a black-box manner to solving $O(\log W)$ instances where the weights are polynomially bounded, i.e., $W \leq \|b\|_1$. So the running time will be reduced to $\tilde{O}(m^{10/7} \log W)$. This is described in the following theorem that shows how to execute scaling on the dual problem.

Theorem 42. *Assuming $\|b\|_1 = O(m)$, we can find a maximum weight vertex b -packing in $\tilde{O}(m^{10/7} \log W)$ time.*

Proof. Consider Algorithm 4 that executes cost scaling for the dual solution. It first recursively computes the dual problem for a graph with costs that are smaller by a factor of 2. Then it uses the obtained solution as a starting point for the computation of the dual solution in the original graph.

The only step of the algorithm that requires some explanation is step 5 – if we would have removed this step the algorithm would compute the dual solution as it only works with reduced weights. When this step is present we just need to argue that \tilde{y}^n is the optimal dual solution for $\tilde{G} = (V_1 \cup V_2, E, \tilde{c})$. This can be shown by arguing that the optimal primal solution in \tilde{G} never uses edges of weight higher than $\|b\|_1$. Take an optimal primal solution \tilde{M} for \tilde{G} and take an optimal primal solution \bar{M} for \bar{G} . By the scaling procedure edges in \bar{M} have weights either 0 or 1, as they were tight for \bar{y} . By optimality of \tilde{M} we know that $\tilde{c}(\tilde{M}) \leq \tilde{c}(\bar{M}) \leq \|b\|_1$, so no single edge of \tilde{M} can have weight higher than $\|b\|_1$. \square

The above theorem can be easily extended to computing the minimum weight perfect b -matching. We only need to take the optimal dual solution and restrict our attention to edges which are tight with respect to it. By finding any perfect b -matching in this unweighted tight graph we obtain the minimum weight perfect b -matching in the original graph.

Corollary 43. *Assuming $\|b\|_1 = O(m)$, we can find a minimum weight perfect b -matching in $\tilde{O}(m^{10/7} \log W)$ time.*

8 Shortest Paths with Negative Weights

We are given a directed graph $G(V, E, c)$ together with the edge weight function $c : E \rightarrow \{-W, \dots, 0, \dots, W\}$ and a source vertex s . Our goal is to compute shortest paths from s to all vertices in V . We will start by reducing this shortest paths problem to the weighted perfect 1-matching problem using the reduction that was given by Gabow [9]. The main step of this reduction is a construction of a bipartite graph $G_{12} =$

$(V_1 \cup V_2, E_{12}, c_{12})$ such that the vertex packing problem in G_{12} induces a valid potential function on G . Using this potential function we can reweigh w to remove negative weights.

We define a bipartite graph $G_{12} = (V_1 \cup V_2, E_{12}, c_{12})$ in the following way

$$\begin{aligned} V_1 &= \{v_1 | v \in V\}, \\ V_2 &= \{v_2 | v \in V\}, \\ E_{12} &= \{u_2v_1 | uv \in E\} \cup \{v_1v_2 | v \in V\}. \end{aligned}$$

The weight function w_{12} is defined as follows

$$c_{12}(u_iv_j) = \begin{cases} c_{uv} & \text{if } uv \in E, \\ 0 & \text{otherwise} \end{cases}$$

Let us observe that a perfect 1-matching in G_{12} corresponds to a set of cycles in the graph G . This leads to the following observation.

Lemma 44. *The graph G contains a negative length cycle if and only if the weight of minimum cost perfect 1-matching in G_{12} is negative.*

Hence, if the weight of minimum cost perfect 1-matching in G_{12} is negative we conclude that there is a negative weight cycle in the graph G . Hence, we assume that the weight of minimum cost 1-matching is equal to 0. Moreover, the dual solution to 1-matchings, i.e., 1-vertex packing in G_{12} induces a potential function on G . Let $y : V_1 \cup V_2 \rightarrow \mathbb{R}$ be the maximum weight vertex 1-packing.

Lemma 45. *Let y be maximum vertex 1-packing in G_{12} . If $y(V_1 \cup V_2) = 0$ then $p_v := y(v_1)$ is a potential function on G , i.e., $c_{uv} + p_u - p_v \geq 0$.*

Proof. As $y(V_1 \cup V_2) = c(M) = 0$, i.e., the minimum weight perfect matching in G_{12} has weight zero. This means that all edges v_1v_2 for $v \in V$ are tight as they form a perfect 1-matching of weight 0. This in turn implies that $y(v_1) = -y(v_2)$. Thus for an edge $uv \in E$ we have

$$\begin{aligned} y(u_2) + y(v_1) &\leq c_{12}(u_2v_1) \\ -y(u_1) + y(v_1) &\leq c_{12}(u_2v_1) \\ -p_u + p_v &\leq c_{uv}. \end{aligned}$$

□

By combining Theorem 41 with the above lemma we obtain the following, i.e., we use the potential function to obtain a reweighed non-negative instance that can be solved using Dijkstra's algorithm.

Corollary 46. *Single source shortest paths in a graph with negative weights can be computed in $\tilde{O}(m^{10/7} \log W)$ time.*

A Proof of Lemma 23

Proof. We require a simple refinement of the proof of Lemma 18. The crucial aspect of this proof is that it does not require any control over how the energy changes between two iterations belonging to different phases. Just as we did in (5.7), let us bound the energy increase during a phase that starts at iteration $t_0 + 1$. Let $\hat{T} = m^{2\eta}$ represent the length of a phase.

$$\begin{aligned}
\sum_{t=t_0+1}^{t_0+\hat{T}} \Delta^t &\leq \sum_{t=t_0+1}^{t_0+\hat{T}} (\mathcal{E}^t - \mathcal{E}^{t-1} + 5c_\rho^2 \cdot m^{1-2\eta}) \\
&\leq \mathcal{E}^{t_0+\hat{T}} + \hat{T} \cdot 5c_\rho^2 \cdot m^{1-2\eta} \\
&\leq \left\| \nu^{t_0+\hat{T}} \right\|_1 + \hat{T} \cdot 5c_\rho^2 \cdot m^{1-2\eta} \\
&\leq \left(\left\| \nu^{t_0} \right\|_1 + \sum_{t=t_0+1}^{t_0+\hat{T}} 2C \cdot \Delta^t \right) + \hat{T} \cdot 5c_\rho^2 \cdot m^{1-2\eta} \\
&= \left\| \nu^{t_0} \right\|_1 + 2C \cdot \sum_{t=t_0+1}^{t_0+\hat{T}} \Delta^t + \hat{T} \cdot 5c_\rho^2 \cdot m^{1-2\eta}
\end{aligned}$$

Therefore

$$\sum_{t=t_0+1}^{t_0+\hat{T}} \Delta^t \leq \frac{1}{1-2C} \left(\left\| \nu^{t_0} \right\|_1 + \hat{T} \cdot 5m^{1-2\eta} \right) = 2 \cdot \left(\left\| \nu^{t_0} \right\|_1 + 5c_\rho^2 \cdot m \right) \quad (\text{A.1})$$

Also, the increase in measure during this phase satisfies

$$\left\| \nu^{t_0+\hat{T}} \right\|_1 - \left\| \nu^{t_0} \right\|_1 \leq 2C \cdot \left(\sum_{t=t_0+1}^{t_0+\hat{T}} \Delta^t \right) \leq 4C \left(\left\| \nu^{t_0} \right\|_1 + 5c_\rho^2 \cdot m \right) = 4C \left\| \nu^{t_0} \right\|_1 + m^{1/2+3\eta}/c_T \quad (\text{A.2})$$

Using the fact that initially $\left\| \nu^0 \right\| = m$, and applying (A.2) we obtain by induction that, for the first $K = c_T \cdot m^{1/2-3\eta}$ phases, the total measure increase during a phase satisfies

$$\left\| \nu^{t_0+\hat{T}} \right\|_1 - \left\| \nu^{t_0} \right\|_1 \leq \frac{2}{c_T} \cdot m^{1/2+3\eta}$$

This can be easily verified for the base case: $4Cm + m^{1/2+3\eta}/c_T = 4 \cdot \frac{m^{-1/2+3\eta}}{20c_T c_\rho^2} \cdot m + \frac{m^{1/2+3\eta}}{c_T} \leq \frac{2}{c_T} \cdot m^{1/2+3\eta}$. For the induction step, we have by the induction hypothesis that after K phases $\left\| \nu^{t_0} \right\|_1 \leq m + K \cdot \frac{2}{c_T} \cdot m^{1/2+3\eta}$. Therefore the measure increase in the new phase is at most $4C \left(m + K \cdot \frac{2}{c_T} \cdot m^{1/2+3\eta} \right) + m^{1/2+3\eta}/c_T \leq 4 \cdot \frac{m^{-1/2+3\eta}}{20c_T c_\rho^2} \cdot (m + 2m) + \frac{m^{1/2+3\eta}}{c_T} \leq \frac{2}{c_T} \cdot m^{1/2+3\eta}$.

Therefore, when the phases are over, we have

$$\left\| \nu \right\|_1 \leq m + c_T \cdot m^{1/2-3\eta} \cdot \frac{2}{c_T} \cdot m^{1/2+3\eta} \leq 3m \quad (\text{A.3})$$

which shows that Invariant 1 holds.

Also, plugging in (A.1) we obtain that during a phase the total energy increase is at most

$$2 \cdot (3m + 5c_\rho^2 \cdot m) \leq 16c_\rho^2 \cdot m$$

□

B Proof of Lemma 29

Proof. Let \hat{f} be an electrical flow in the graph with auxiliary arcs, and let $\hat{\phi}$ be the corresponding vertex potentials. Let $\hat{\phi}^{(P)}$ be a different set of vertex potentials satisfying

$$\hat{\phi}_v^{(P)} = \begin{cases} \hat{\phi}_v & v \in P \\ \frac{r_{\bar{u}v}\hat{\phi}_u + r_{uv}\hat{\phi}_{\bar{u}}}{r_{uv} + r_{\bar{u}v}} & v \in Q, (u, v), (\bar{u}, v) \in E \\ \hat{\phi}_{v_0} & v = v_0 \end{cases}$$

Let $\hat{\phi}^{(Q)} = \hat{\phi} - \hat{\phi}^{(P)}$. Also let $\hat{f}^{(P)}$ and $\hat{f}^{(Q)}$ be the electrical flows corresponding to potentials $\hat{\phi}^{(P)}$, respectively $\hat{\phi}^{(Q)}$ with the same set of resistances r . So these flows can be constructed directly from the vertex potentials $\hat{\phi}$.

Now let us verify that $\hat{f}^{(P)}$ satisfies flow conservation on all vertices in $Q \cup \{v_0\}$. Given any vertex $v \in Q$, let u and \bar{u} be the vertices corresponding to partner arcs (u, v) and (\bar{u}, v) . By Ohm's law, we have

$$\hat{f}_{uv}^{(P)} = \frac{\hat{\phi}_v^{(P)} - \hat{\phi}_u^{(P)}}{r_{uv}} = \frac{\frac{r_{\bar{u}v}\hat{\phi}_u + r_{uv}\hat{\phi}_{\bar{u}}}{r_{uv} + r_{\bar{u}v}} - \hat{\phi}_u}{r_{uv}} = \frac{r_{uv}(\hat{\phi}_{\bar{u}} - \hat{\phi}_u)}{r_{uv}(r_{uv} + r_{\bar{u}v})} = \frac{\hat{\phi}_{\bar{u}} - \hat{\phi}_u}{r_{uv} + r_{\bar{u}v}}$$

and similarly

$$\hat{f}_{\bar{u}v}^{(P)} = \frac{\hat{\phi}_v^{(P)} - \hat{\phi}_{\bar{u}}^{(P)}}{r_{\bar{u}v}} = \frac{\frac{r_{\bar{u}v}\hat{\phi}_u + r_{uv}\hat{\phi}_{\bar{u}}}{r_{uv} + r_{\bar{u}v}} - \hat{\phi}_{\bar{u}}}{r_{\bar{u}v}} = \frac{r_{\bar{u}v}(\hat{\phi}_u - \hat{\phi}_{\bar{u}})}{r_{\bar{u}v}(r_{uv} + r_{\bar{u}v})} = \frac{\hat{\phi}_u - \hat{\phi}_{\bar{u}}}{r_{uv} + r_{\bar{u}v}}$$

Hence $\hat{f}_{uv}^{(P)} + \hat{f}_{\bar{u}v}^{(P)} = 0$, so flow is conserved at v . The fact that flow is also conserved at v_0 is immediate, since potentials on $P \cup \{v_0\}$ are identical to those in $\hat{\phi}$, and the flow \hat{f} corresponding to these is always conserved at v_0 .

Furthermore, since f satisfies the same demands as \hat{f} , $f^{(P)}$ also obeys flow conservation on $Q \cup \{v_0\}$. Now, consider restricting the flows $f^{(P)}$ and $\hat{f}^{(Q)}$ to a pair of partner edges, $e = (u, v)$ and $\bar{e} = (\bar{u}, v)$ (i.e. zeroing out the flows on all other edges). This restriction of $\hat{f}^{(Q)}$ is an electrical flow induced by a voltage vector nonzero only at v , so the \mathbf{R} -inner product with it simply measures net flow to v . As $f^{(P)}$ obeys flow conservation at v , this implies that the restrictions of $f^{(P)}$ and $\hat{f}^{(Q)}$ are \mathbf{R} -orthogonal. That in turn implies that the energy of each of these restricted flows is at most the energy of the same restriction of f , which is equal to $\nu_e + \nu_{\bar{e}}$. \square

C Proof of Theorem 6

Proving this theorem is done in two parts. In the first one we analyze the predictor step, which produces a new iterate that is close to the central path ℓ_2 norm, and track the change in resistance. In the second part we analyze the centering steps, which - starting with a solution that is close to the central path - produce a centered solution. In both cases, we will also show that the iterates stay feasible at all times.

We start with analyzing the predictor step, described in Section 4.2:

$$\begin{aligned} f'_e &:= (1 - \delta)f_e + \delta\hat{f}_e, \\ s'_e &:= s_e - \frac{\delta}{1 - \delta}(\hat{\phi}_v - \hat{\phi}_u) \end{aligned}$$

for all arcs $e = (u, v)$.

We verify feasibility of f' and s' , and bound the change in resistance.

Proposition 47. *Both f' and s' are feasible, i.e. $f', s' > 0$. Furthermore, $\frac{r_e}{r'_e} = \frac{s_e}{f_e} \cdot \frac{f'_e}{s'_e} \leq 1 + 4\delta\rho_e$.*

Proof. We obtain

$$\begin{aligned}
f_{e'} &= (1 - \delta)f_e + \delta\hat{f}_e \geq (1 - \delta)f_e - \delta|\hat{f}_e| = (1 - \delta - \delta\rho_e)f_e \geq \left(1 - \delta - \frac{1}{8\|\rho\|_{\nu,4}}\rho_e\right)f_e \\
&\geq \left(1 - \frac{1}{8} - \frac{1}{8\|\rho\|_{\infty}}\rho_e\right)f_e \geq \left(1 - \frac{1}{8} - \frac{1}{8}\right)f_e > 0 \\
s_{e'} &= s_e - \frac{\delta}{1 - \delta}(\hat{\phi}_v - \hat{\phi}_u) = s_e - \frac{\delta}{1 - \delta} \cdot \frac{s_e}{f_e} \cdot \hat{f}_e = s_e \left(1 - \frac{\delta}{1 - \delta} \cdot \frac{\hat{f}_e}{f_e}\right) \\
&\geq s_e \left(1 - \frac{\delta}{1 - \delta} \cdot \rho_e\right) \geq s_e(1 - \delta \cdot \rho_e) \geq s_e \left(1 - \frac{1}{8\|\rho\|_{\nu,4}}\rho_e\right) \\
&\geq s_e \left(1 - \frac{1}{8\|\rho\|_{\infty}}\rho_e\right) \geq s_e \left(1 - \frac{1}{8}\right) > 0
\end{aligned}$$

Next, we analyze the change in resistance:

$$\begin{aligned}
\frac{r_e}{r'_e} &= \frac{s_e}{f_e} \cdot \frac{f'_e}{s'_e} = \frac{s_e}{f_e} \cdot \frac{(1 - \delta)f_e + \delta\hat{f}_e}{s_e \left(1 - \frac{\delta}{1 - \delta} \cdot \frac{\hat{f}_e}{f_e}\right)} = \frac{1 - \delta + \delta\frac{\hat{f}_e}{f_e}}{1 - \frac{\delta}{1 - \delta} \cdot \frac{\hat{f}_e}{f_e}} \leq (1 - \delta) \frac{1 - \delta + \delta\rho_e}{1 - \delta - \delta\rho_e} \\
&\leq \frac{1 - \delta + \delta\rho_e}{1 - \delta - \delta\rho_e} = 1 + \frac{2\delta\rho_e}{1 - \delta - \delta\rho_e} \leq 1 + \frac{2\delta\rho_e}{1 - \frac{1}{8} - \frac{1}{8\|\rho\|_{\nu,4}} \cdot \rho_e} \\
&\leq 1 + \frac{2\delta\rho_e}{1 - \frac{1}{8} - \frac{1}{8\|\rho\|_{\infty}} \cdot \rho_e} \leq 1 + \frac{2\delta\rho_e}{1 - \frac{1}{8} - \frac{1}{8}} \leq 1 + 4\delta\rho_e
\end{aligned}$$

□

As we saw in Section 4.2, taking the predictor step produces a solution that is close to the central path in ℓ_2 norm. We will now show that once this happens, only a few centering steps are required.

For simplicity, we will overload notation for the rest of the section, in order to measure centrality with respect to the following quantity:

Definition 48. $\left\| \frac{fs}{\mu} - 1 \right\|_{\nu,2} := \sqrt{\sum_e \nu_e \left(\frac{f_e s_e}{\mu_e} - 1 \right)^2}$

Now let us describe a centering step. Given an instance (f, s, ν) such that $\left\| \frac{fs}{\mu} - 1 \right\|_{\nu,2}^2 = \sum_e \nu_e \left(\frac{f_e s_e}{\mu_e} - 1 \right)^2 \leq \frac{1}{256}$, we show how to produce a new instance with better centrality. In order to do so, consider a flow $f^\#$ such that $f_e^\# = \frac{\mu_e}{s_e}$. This is clearly central, but it routes a different demand. In order to produce a flow that routes the correct demand, we take the demand of $f - f^\#$ and route it electrically with resistances $\tilde{r}_e = \frac{s_e}{f_e^\#} = \frac{s_e^2}{\mu_e}$ producing an electrical flow \tilde{f} . Then we set the new flow

$$\begin{aligned}
f'_e &= f_e^\# + \tilde{f}_e \\
s'_e &= s_e - \frac{s_e}{f_e^\#} \cdot \tilde{f}_e
\end{aligned}$$

We first verify that these flows stay feasible. In order to do so, we will require a proposition that we will use in upper bounding the ratio $\left| \frac{\tilde{f}_e}{f_e^\#} \right|$.

Proposition 49. $\left\| \frac{\tilde{f}_e}{f_e^\#} \right\|_{\nu,2} \leq \left\| \frac{f^s}{\mu} - 1 \right\|_{\nu,2}$

Proof. This is a straightforward energy minimization argument. We first express $\left\| \frac{\tilde{f}}{f^\#} \right\|_{\nu,2}^2$ in terms of the energy of the electrical flow \tilde{f} . Then, using the fact that $f - f^\#$ routes the same demand as the electrical flow \tilde{f} , the energy of $f - f^\#$ can only be larger than that of \tilde{f} . This yields:

$$\begin{aligned} \left\| \frac{\tilde{f}}{f^\#} \right\|_{\nu,2}^2 &= \sum_e \nu_e \left(\frac{\tilde{f}_e}{f_e^\#} \right)^2 = \sum_e \frac{\nu_e}{s_e f_e^\#} \cdot \tilde{r}_e \tilde{f}_e^2 = \sum_e \frac{\nu_e}{\mu_e} \cdot \tilde{r}_e \tilde{f}_e^2 = \frac{1}{\hat{\mu}} \cdot \sum_e \tilde{r}_e \tilde{f}_e^2 \\ &\leq \frac{1}{\hat{\mu}} \sum_e \tilde{r}_e (f_e - f_e^\#)^2 = \sum_e \frac{\nu_e}{(f_e^\#)^2} \cdot (f_e - f_e^\#)^2 \\ &= \sum_e \nu_e \left(\frac{f_e}{f_e^\#} - 1 \right)^2 = \sum_e \nu_e \left(\frac{f_e s_e}{\mu_e} - 1 \right)^2 \\ &= \left\| \frac{f^s}{\mu} - 1 \right\|_{\nu,2}^2 \end{aligned}$$

□

Now we can prove the following proposition:

Proposition 50. *After taking a centering step, f' and s' stay feasible, i.e. $f', s' > 0$. Furthermore $\left\| \frac{f' s'}{\mu} - 1 \right\|_{\nu,2} \leq \left\| \frac{f^s}{\mu} - 1 \right\|_{\nu,2}^2$. Also, $\frac{r_e}{r_{e'}} = \frac{s_e}{f_e} \cdot \frac{f'_e}{s'_e} \leq \left(1 + 2 \left| \frac{s_e f_e}{\mu} - 1 \right| \right) \left(1 + 4 \left| \frac{\tilde{f}_e}{f_e^\#} \right| \right)$.*

Proof. To lower bound f'_e , we can simply use Proposition 49 to upper bound $\left| \frac{\tilde{f}_e}{f_e^\#} \right| \leq \left\| \frac{f^s}{\mu} - 1 \right\|_{\nu,2} \leq \frac{1}{16}$. This yields $\left| \frac{\tilde{f}_e}{f_e^\#} \right| \leq f_e^\# \cdot \frac{1}{16}$. Plugging in, we obtain:

$$f'_e = f_e^\# + \tilde{f}_e \geq f_e^\# - f_e^\# \cdot \left\| \frac{f^s}{\mu} - 1 \right\|_{\nu,2} \geq f_e^\# - f_e^\# \cdot \frac{1}{16} > 0$$

In order to prove feasibility for the new slack iterate, we observe that

$$s'_e = s_e \left(1 - \frac{\tilde{f}_e}{f_e^\#} \right) \geq s_e \left(1 - \frac{1}{16} \right) > 0$$

Then we verify that this step improves centrality.

$$\begin{aligned}
\sum_e \nu_e \left(\frac{f'_e s'_e}{\mu_e} - 1 \right)^2 &= \sum_e \nu_e \left(\frac{(f_e^\# + \tilde{f}_e) \left(s_e - \frac{s_e}{f_e^\#} \cdot \tilde{f}_e \right)}{\mu_e} - 1 \right)^2 = \sum_e \nu_e \left(\frac{f_e^\# s_e - \frac{s_e}{f_e^\#} \cdot \tilde{f}_e^2}{\mu_e} - 1 \right)^2 \\
&= \sum_e \nu_e \left(\frac{\mu_e - \mu_e \left(\frac{\tilde{f}_e}{f_e^\#} \right)^2}{\mu_e} - 1 \right)^2 = \sum_e \nu_e \left(\frac{\tilde{f}_e}{f_e^\#} \right)^4
\end{aligned}$$

Applying Cauchy-Schwarz, we see that this is upper bounded by

$$\begin{aligned}
\left(\sum_e \nu_e \cdot \left(\frac{\tilde{f}_e}{f_e^\#} \right)^2 \right)^2 &= \left(\sum_e \nu_e \cdot \tilde{r}_e \frac{f_e^\#}{s_e} \cdot \left(\frac{\tilde{f}_e}{f_e^\#} \right)^2 \right)^2 = \left(\sum_e \nu_e \cdot \frac{\tilde{r}_e}{\mu_e} \cdot \tilde{f}_e^2 \right)^2 \\
&= \left(\sum_e \nu_e \cdot \frac{\tilde{r}_e}{\hat{\mu} \cdot \nu_e} \cdot \tilde{f}_e^2 \right)^2 = \left(\frac{1}{\hat{\mu}} \cdot \sum_e \tilde{r}_e \tilde{f}_e^2 \right)^2
\end{aligned}$$

Now note that the last term contains the energy of the electrical flow \tilde{f} with respect to resistances \tilde{r} . Since \tilde{f} and $f - f^\#$ route the same demand, but \tilde{f} is an electrical flow, so it minimizes energy, we know that this is upper bounded by

$$\begin{aligned}
\left(\frac{1}{\hat{\mu}} \cdot \sum_e \tilde{r}_e (f_e - f_e^\#)^2 \right)^2 &= \left(\frac{1}{\hat{\mu}} \cdot \sum_e \frac{s_e^2}{\mu_e} \left(f_e - \frac{\mu_e}{s_e} \right)^2 \right)^2 \\
&= \left(\frac{1}{\hat{\mu}} \cdot \sum_e \mu_e \left(\frac{f_e s_e}{\mu_e} - 1 \right)^2 \right)^2 \\
&= \left(\sum_e \nu_e \left(\frac{f_e s_e}{\mu_e} - 1 \right)^2 \right)^2
\end{aligned}$$

This shows that during every such step, centrality improves at a quadratic rate.

Next we bound how much resistances get changed by a single centering step:

$$\begin{aligned}
\frac{r_e}{r'_e} &= \frac{s_e}{f_e} \cdot \frac{f'_e}{s'_e} = \frac{s_e}{f_e} \cdot \frac{f_e^\# + \tilde{f}_e}{s_e \left(1 - \frac{\tilde{f}_e}{f_e^\#} \right)} = \frac{f_e^\#}{f_e} \cdot \frac{1 + \frac{\tilde{f}_e}{f_e^\#}}{1 - \frac{\tilde{f}_e}{f_e^\#}} = \frac{\mu_e}{s_e f_e} \cdot \frac{1 + \frac{\tilde{f}_e}{f_e^\#}}{1 - \frac{\tilde{f}_e}{f_e^\#}} \leq \frac{\mu_e}{s_e f_e} \cdot \left(1 + 2 \left| \frac{\tilde{f}_e}{f_e^\#} \right| \right) \\
&\leq \left(1 + 2 \left| \frac{s_e f_e}{\mu_e} - 1 \right| \right) \left(1 + 4 \left| \frac{\tilde{f}_e}{f_e^\#} \right| \right)
\end{aligned}$$

We used the inequalities $\frac{1}{1+x} \leq 1 + 2|x|$, and $\frac{1+x}{1-x} \leq 1 + 4|x|$, for $|x| \leq 1/2$. □

Now we can track how resistances change after restoring to near-perfect centrality. This is highlighted by the following proposition.

Proposition 51. *After a predictor step, followed by any number of centering steps, the new set of resistances $r^{(k)}$ satisfy $\frac{r_e}{r_e^{(k)}} \leq 1 + 4\delta\rho_e + \kappa_e$, for some κ such that $\|\kappa\|_{\nu,2} \leq 1$.*

Proof. Let $f^{(i)}$ and $s^{(i)}$ be the flow and slack iterates obtained after applying a predictor step to (f, s, ν) , followed by i centering steps. Let $\tilde{f}^{(i)}$ and $f^\sharp(i)$ be defined similarly. Combining Proposition 47 with Proposition 50 we can upper bound the resistance $r_e^{(k)} = \frac{s^{(k)}}{f^{(k)}}$: \square

$$\begin{aligned}
\frac{r_e}{r_e^{(k)}} &\leq (1 + 4\delta\rho_e) \cdot \prod_{i=0}^{k-1} \left(1 + 2 \left| \frac{f_e^{(i)} s_e^{(i)}}{\mu} - 1 \right| \right) \left(1 + 4 \left| \frac{\tilde{f}_e^{(i)}}{f_e^\sharp(i)} \right| \right) \\
&\leq (1 + 4\delta\rho_e) \cdot \exp \left(2 \left(\sum_{i=0}^{k-1} 2 \left| \frac{f_e^{(i)} s_e^{(i)}}{\mu} - 1 \right| + 4 \left| \frac{\tilde{f}_e^{(i)}}{f_e^\sharp(i)} \right| \right) \right) \\
&\leq (1 + 4\delta\rho_e) \cdot \left(1 + \left(\sum_{i=0}^{k-1} 8 \left| \frac{f_e^{(i)} s_e^{(i)}}{\mu} - 1 \right| + 16 \left| \frac{\tilde{f}_e^{(i)}}{f_e^\sharp(i)} \right| \right) \right) \\
&\leq 1 + 4\delta\rho_e + \frac{1}{2} \cdot \sum_{i=0}^{k-1} \left(8 \left| \frac{f_e^{(i)} s_e^{(i)}}{\mu} - 1 \right| + 16 \left| \frac{\tilde{f}_e^{(i)}}{f_e^\sharp(i)} \right| \right) \\
&= 1 + 4\delta\rho_e + \sum_{i=0}^{k-1} \left(4 \left| \frac{f_e^{(i)} s_e^{(i)}}{\mu} - 1 \right| + 8 \left| \frac{\tilde{f}_e^{(i)}}{f_e^\sharp(i)} \right| \right) \\
&= 1 + 4\delta\rho_e + \kappa_e
\end{aligned}$$

We used the fact that $\sum_{i=0}^{k-1} 2 \left| \frac{f_e^{(i)} s_e^{(i)}}{\mu} - 1 \right| + 4 \left| \frac{\tilde{f}_e^{(i)}}{f_e^\sharp(i)} \right| \leq \sum_{i=0}^{k-1} 6 \left\| \frac{f^{(i)} s^{(i)}}{\mu} - 1 \right\|_{\nu,2} \leq 6 \cdot \frac{1}{16} \leq 1$, and $\exp(x) \leq 1 + 2|x|$ for $|x| \leq 1$.

Now to bound the norm $\|\kappa\|_{\nu,2}$, we use triangle inequality and bound the contribution of each centering step separately:

$$\begin{aligned}
\|\kappa\|_{\nu,2} &\leq \sum_{i=0}^{k-1} \left(4 \left\| \frac{f^{(i)} s^{(i)}}{\mu} - 1 \right\|_{\nu,2} + 8 \left\| \frac{\tilde{f}^{(i)}}{f^\sharp(i)} \right\|_{\nu,2} \right) \\
&\leq 12 \sum_{i=0}^{k-1} \left\| \frac{f^{(i)} s^{(i)}}{\mu} - 1 \right\|_{\nu,2} \\
&\leq 12 \sum_{i=0}^{k-1} \left(\frac{1}{16} \right)^{2^i} \\
&\leq 12 \cdot \frac{1}{15} \\
&\leq 1
\end{aligned}$$

Remark 52. In Proposition 50 we saw that centrality decreases at a quadratic rate each iteration. Therefore centrality decreases to machine precision in $\tilde{O}(\log \log U)$ iterations. For this reason, we perform centering steps until we restore centrality to machine precision $O(\epsilon_{\text{mach}})$ (the number of such iterations will be absorbed by the \tilde{O} in the running time). Although throughout the paper we work with exact centrality, rather than having it set to $O(\epsilon_{\text{mach}})$, fixing this can easily be done by perturbing the measures at the end of the progress step in order to guarantee exact centrality. This changes all measures by $O(\epsilon_{\text{mach}})$; all the analysis we do is still valid, but this extra change needs to be carried over throughout the computations. We chose to ignore it in order to simplify the presentation.

D Preconditioning proofs.

The purpose of this section is to prove the core result, Lemma 32, on the guarantees of preconditioning. However, we begin with a simpler, more general statement:

Lemma 53. *Let N be an electrical network, let $f^{(N)}$ be a flow on N , and let $\hat{f}^{(N)}$ be the electrical flow on N satisfying the same demands as $f^{(N)}$, induced by voltages $\hat{\phi}^{(N)}$. Let ϕ_0 be any fixed real number (representing an absolute voltage). Then*

$$\sum_{\substack{e=(u,v), \\ \max(|\hat{\phi}_u^{(N)} - \phi_0|, |\hat{\phi}_v^{(N)} - \phi_0|) \geq V/4}} r_e (f_e^{(N)})^2 \geq \frac{1}{2} \sum_{\substack{e=(u,v), \\ |\hat{\phi}_u^{(N)} - \hat{\phi}_v^{(N)}| \geq V}} r_e (\hat{f}_e^{(N)})^2.$$

Proof. Given a set of vertex potentials $\hat{\phi}^{(N)}$, a flow f and an absolute voltage ϕ , we define

$$F(\hat{\phi}^{(N)}, f, \phi) = \sum_{\substack{e=(u,v) \\ \min(\hat{\phi}_u^{(N)}, \hat{\phi}_v^{(N)}) \leq \phi \leq \max(\hat{\phi}_u^{(N)}, \hat{\phi}_v^{(N)})}} |f_e|.$$

That is, $F(\hat{\phi}^{(N)}, f, \phi)$ is the sum of the absolute values of the flows of edges crossing the cut induced by the voltage ϕ . Notably, $\int_{-\infty}^{\infty} F(\hat{\phi}^{(N)}, \hat{f}^{(N)}, \phi) d\phi$ is equal to the energy of the electrical flow $\hat{f}^{(N)}$, since the contribution of each edge $e = (u, v)$ is $|\hat{f}_e^{(N)}| |\hat{\phi}_v^{(N)} - \hat{\phi}_u^{(N)}| = r_e |\hat{f}_e^{(N)}|^2$. Furthermore, for all ϕ we have $F(\hat{\phi}^{(N)}, \hat{f}^{(N)}, \phi) \leq F(\hat{\phi}^{(N)}, f^{(N)}, \phi)$. This is because $\hat{f}^{(N)}$ satisfies the same demands as $f^{(N)}$, and thus they have the same *net* flow across each cut, including the cuts induced by the voltage ϕ . In the electrical flow, all of the flow is in the direction of increasing $\hat{\phi}^{(N)}$ by definition; thus the sum of the absolute values of the edge flows, $F(\hat{\phi}^{(N)}, \hat{f}^{(N)}, \phi)$, is equal to that net flow value. In $f^{(N)}$, on the other hand, that net flow value is a lower bound for $F(\hat{\phi}^{(N)}, f^{(N)}, \phi)$.

We now define

$$X = \left(\int_{-\infty}^{\phi_0 - \frac{V}{4}} F(\hat{\phi}^{(N)}, \hat{f}^{(N)}, \phi) d\phi \right) + \left(\int_{\phi_0 + \frac{V}{4}}^{\infty} F(\hat{\phi}^{(N)}, \hat{f}^{(N)}, \phi) d\phi \right)$$

$$Y = \left(\int_{-\infty}^{\phi_0 - \frac{V}{4}} F(\hat{\phi}^{(N)}, f^{(N)}, \phi) d\phi \right) + \left(\int_{\phi_0 + \frac{V}{4}}^{\infty} F(\hat{\phi}^{(N)}, f^{(N)}, \phi) d\phi \right).$$

Because the inequality holds for the integrands everywhere, we have that $X \leq Y$. Alternatively, we may define t_e for an edge (u, v) to be the fraction of the interval $[\min(\hat{\phi}_u^{(N)}, \hat{\phi}_v^{(N)}), \max(\hat{\phi}_u^{(N)}, \hat{\phi}_v^{(N)})]$ that is not contained in the interval $[\phi_0 - \frac{V}{4}, \phi_0 + \frac{V}{4}]$. We can then expand out X and Y as sums over edges, and define a new Z :

$$X = \sum_e t_e \left| \hat{f}_e^{(N)} \right| \left| \hat{\phi}_v^{(N)} - \hat{\phi}_u^{(N)} \right|$$

$$= \sum_e t_e r_e \left| \hat{f}_e^{(N)} \right|^2$$

$$Y = \sum_e t_e \left| f_e^{(N)} \right| \left| \hat{\phi}_v^{(N)} - \hat{\phi}_u^{(N)} \right|$$

$$= \sum_e t_e r_e \left| \hat{f}_e^{(N)} \right| \left| f_e^{(N)} \right|$$

$$Z = \sum_e t_e r_e \left| f_e^{(N)} \right|^2.$$

Then by Cauchy-Schwarz on the sum expressions, we have $Y \leq \sqrt{XZ}$; on the other hand, since $X \leq Y$ this means that we have $X \leq \sqrt{XZ}$ and so $Z \geq X$. Finally,

$$\sum_{\substack{e=(u,v), \\ \max(|\hat{\phi}_u^{(N)} - \phi_0|, |\hat{\phi}_v^{(N)} - \phi_0|) \geq V/4}} r_e (f_e^{(N)})^2 \geq Z$$

since only those edges with an endpoint outside of $[\phi_0 - \frac{V}{4}, \phi_0 + \frac{V}{4}]$ have $t_e > 0$, while we have

$$X \geq \frac{1}{2} \sum_{\substack{e=(u,v), \\ |\hat{\phi}_u^{(N)} - \hat{\phi}_v^{(N)}| \geq V}} r_e (\hat{f}_e^{(N)})^2$$

since each of the edges in the latter sum, with their interval of length $|\hat{\phi}_u^{(N)} - \hat{\phi}_v^{(N)}| \geq V$, has $t_e \geq \frac{1}{2}$. \square

We can now prove Lemma 32:

Proof of Lemma 32. Apply Lemma 53 to N and $f^{(N)}$, with $\phi_0 = \hat{\phi}_{v_0}^{(N)}$. Now, note that since $r_e (f_e^{(N)})^2 \leq \nu_e$, the total energy of $f^{(N)}$, $\sum_e r_e (f_e^{(N)})^2$, is at most $\sum_e \nu_e$. As an electrical flow satisfying the same demands, the total energy of $\hat{f}^{(N)}$ is also at most $\sum_e \nu_e$.

On the other hand, every node v is connected to v_0 by an edge with resistance $\frac{R}{a(v)}$, which contributes $\frac{(\hat{\phi}_v^{(N)} - \hat{\phi}_{v_0}^{(N)})^2 a(v)}{R}$. Thus the total energy is at least

$$\begin{aligned} \sum_v \frac{(\hat{\phi}_v^{(N)} - \hat{\phi}_{v_0}^{(N)})^2 a(v)}{R} &= \sum_v \frac{(\hat{\phi}_v^{(N)} - \phi_0)^2 a(v)}{R} \\ &\geq \frac{V^2}{16R} \sum_{v, |\hat{\phi}_v^{(N)} - \phi_0| \geq \frac{V}{4}} a(v) \\ &\geq \frac{V^2}{16R} \sum_{\substack{e=(u,v), \\ \max(|\hat{\phi}_u^{(N)} - \phi_0|, |\hat{\phi}_v^{(N)} - \phi_0|) \geq \frac{V}{4}}} \nu'_e \\ &\geq \frac{V^2}{16R} \left(\left(\sum_{\substack{e=(u,v), \\ \max(|\hat{\phi}_u^{(N)} - \phi_0|, |\hat{\phi}_v^{(N)} - \phi_0|) \geq \frac{V}{4}}} \nu_e \right) - \sum_e (\nu_e - \nu'_e) \right) \\ &\geq \frac{V^2}{16R} \left(\left(\sum_{\substack{e=(u,v), \\ \max(|\hat{\phi}_u^{(N)} - \phi_0|, |\hat{\phi}_v^{(N)} - \phi_0|) \geq \frac{V}{4}}} r_e (f_e^{(N)})^2 \right) - M \right) \\ &\geq \frac{V^2}{16R} \left(\frac{1}{2} \left(\sum_{\substack{e=(u,v), \\ |\hat{\phi}_u^{(N)} - \hat{\phi}_v^{(N)}| \geq V}} r_e (\hat{f}_e^{(N)})^2 \right) - M \right) \\ &= \frac{V^2}{16R} \left(\frac{1}{2} \left(\sum_{r_e |\hat{f}_e^{(N)}| \geq V} r_e (\hat{f}_e^{(N)})^2 \right) - M \right). \end{aligned}$$

We therefore have

$$\frac{V^2}{16R} \left(\frac{1}{2} \left(\sum_{r_e |\hat{f}_e^{(N)}| \geq V} r_e (\hat{f}_e^{(N)})^2 \right) - M \right) \leq \sum_e \nu_e$$

$$\sum_{r_e |\hat{f}_e^{(N)}| \geq V} r_e (\hat{f}_e^{(N)})^2 \leq \frac{32R \sum_e \nu_e}{V^2} + 2M.$$

This completes the proof. □

References

- [1] R. Bellman. On a Routing Problem. *Quarterly of Applied Mathematics*, 16(1):87–90, 1958.
- [2] Bela Bollobas. *Modern Graph Theory*. Springer, 1998.
- [3] Michael B. Cohen, Rasmus Kyng, Gary L. Miller, Jakub W. Pachocki, Richard Peng, Anup B. Rao, and Shen Chen Xu. Solving sdd linear systems in nearly $m \log^{1/2} n$ time. In *STOC'14: Proceedings of the 46th Annual ACM Symposium on Theory of Computing*, pages 343–352, 2014.
- [4] Samuel I. Daitch and Daniel A. Spielman. Faster approximate lossy generalized flow via interior point algorithms. In *Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing*, STOC '08, pages 451–460, New York, NY, USA, 2008. ACM.
- [5] E.A. Dinic and M. A. Kronrod. An Algorithm for the Solution of the Assignment Problem. *Soviet Math. Dokl.*, 10:1324–1326, 1969.
- [6] J. Edmonds and R.M. Karp. Theoretical Improvements in Algorithmic Efficiency for Network Flow Problems. *J. ACM*, 19(2):248–264, 1972.
- [7] E. Egerváry. Matrixok kombinatorius tulajdonságairól (hungarian) on combinatorial properties of matrices. *Matematikai és Fizikai Lapok*, 38:16–28, 1931.
- [8] H. N. Gabow and P. Sankowski. Algebraic algorithms for b-matching, shortest undirected paths, and f-factors. In *Foundations of Computer Science (FOCS), 2013 IEEE 54th Annual Symposium on*, pages 137–146, Oct 2013.
- [9] H.N. Gabow. Scaling Algorithms for Network Problems. *J. Comput. Syst. Sci.*, 31(2):148–168, 1985.
- [10] H.N. Gabow and R.E. Tarjan. Faster Scaling Algorithms for Network Problems. *SIAM J. Comput.*, 18(5):1013–1036, 1989.
- [11] A. Goldberg and R. Tarjan. Solving minimum-cost flow problems by successive approximation. In *Proceedings of the Nineteenth Annual ACM Symposium on Theory of Computing*, STOC '87, pages 7–18, New York, NY, USA, 1987. ACM.
- [12] Andrew V. Goldberg. Scaling algorithms for the shortest paths problem. In *SODA '93: Proceedings of the fourth annual ACM-SIAM Symposium on Discrete algorithms*, pages 222–231. Society for Industrial and Applied Mathematics, 1993.
- [13] M. Iri. A new method for solving transportation-network problems. *Journal of the Operations Research Society of Japan*, 3:27–87, 1960.

- [14] L.R. Ford Jr. Network Flow Theory. Paper P-923, The RAND Corporation, Santa Monica, California, August 1956.
- [15] Donggu Kang and James Payor. Flow rounding. *CoRR*, abs/1507.08139, 2015.
- [16] M.-Y. Kao, T. W. Lam, W.-K. Sung, and H.-F. Ting. A decomposition theorem for maximum weight bipartite matchings with applications to evolutionary trees. In *Proceedings of the 7th Annual European Symposium on Algorithms*, pages 438–449, 1999.
- [17] Jonathan A. Kelner, Lorenzo Orecchia, Aaron Sidford, and Zeyuan Allen Zhu. A simple, combinatorial algorithm for solving SDD systems in nearly-linear time. In *STOC'13: Proceedings of the 45th Annual ACM Symposium on the Theory of Computing*, pages 911–920, 2013.
- [18] Ioannis Koutis, Gary L. Miller, and Richard Peng. Approaching optimality for solving SDD systems. In *FOCS'10: Proceedings of the 51st Annual IEEE Symposium on Foundations of Computer Science*, pages 235–244, 2010.
- [19] Ioannis Koutis, Gary L. Miller, and Richard Peng. A nearly $m \log n$ -time solver for SDD linear systems. In *FOCS'11: Proceedings of the 52nd Annual IEEE Symposium on Foundations of Computer Science*, pages 590–598, 2011.
- [20] H.W Kuhn. The Hungarian Method for the Assignment Problem. *Naval Research Logistics Quarterly*, 2:83–97, 1955.
- [21] Rasmus Kyng, Yin Tat Lee, Richard Peng, Sushant Sachdeva, and Daniel A. Spielman. Sparsified cholesky and multigrid solvers for connection laplacians. In *STOC'16: Proceedings of the 48th Annual ACM Symposium on Theory of Computing*, 2016.
- [22] Rasmus Kyng and Sushant Sachdeva. Approximate gaussian elimination for laplacians: Fast, sparse, and simple. In *FOCS'16: Proceedings of the 57th Annual IEEE Symposium on Foundations of Computer Science*, 2016.
- [23] Y. T. Lee and A. Sidford. Path finding methods for linear programming: Solving linear programs in $\sqrt{\text{rank}}$ iterations and faster algorithms for maximum flow. In *Foundations of Computer Science (FOCS), 2014 IEEE 55th Annual Symposium on*, pages 424–433, Oct 2014.
- [24] Aleksander Mądry. Navigating central path with electrical flows: from flows to matchings, and back. In *FOCS'13: Proceedings of the 54th Annual IEEE Symposium on Foundations of Computer Science*, pages 253–262, 2013.
- [25] E. F. Moore. The Shortest Path Through a Maze. In *Proceedings of the International Symposium on the Theory of Switching*, pages 285–292. Harvard University Press, 1959.
- [26] J. Munkres. Algorithms for the Assignment and Transportation Problems. *Journal of SIAM*, 5(1):32–38, 1957.
- [27] Piotr Sankowski. *Algorithms – ESA 2005: 13th Annual European Symposium, Palma de Mallorca, Spain, October 3-6, 2005. Proceedings*, chapter Shortest Paths in Matrix Multiplication Time, pages 770–778. Springer Berlin Heidelberg, Berlin, Heidelberg, 2005.
- [28] Piotr Sankowski. *Automata, Languages and Programming: 33rd International Colloquium, ICALP 2006, Venice, Italy, July 10-14, 2006, Proceedings, Part I*, chapter Weighted Bipartite Matching in Matrix Multiplication Time, pages 274–285. Springer Berlin Heidelberg, Berlin, Heidelberg, 2006.
- [29] A. Schrijver. *Combinatorial Optimization - Polyhedra and Efficiency*. Springer, 2003.

- [30] A. Shimbel. Structure in Communication Nets. In *In Proceedings of the Symposium on Information Networks*, pages 199–203. Polytechnic Press of the Polytechnic Institute of Brooklyn, Brooklyn, 1955.
- [31] Daniel D. Sleator and Robert Endre Tarjan. A data structure for dynamic trees. *J. Comput. Syst. Sci.*, 26(3):362–391, June 1983.
- [32] Daniel A. Spielman and Shang-Hua Teng. Solving sparse, symmetric, diagonally-dominant linear systems in time $O(m^{1.31})$. In *FOCS'03: Proceedings of the 44th Annual IEEE Symposium on Foundations of Computer Science*, pages 416–427, 2003.
- [33] Daniel A. Spielman and Shang-Hua Teng. Nearly-linear time algorithms for graph partitioning, graph sparsification, and solving linear systems. In *STOC'04: Proceedings of the 36th Annual ACM Symposium on the Theory of Computing*, pages 81–90, 2004.
- [34] Pravin M. Vaidya. Solving linear equations with symmetric diagonally dominant matrices by constructing good preconditioners. Unpublished manuscript, UIUC 1990. A talk based on the manuscript was presented at the IMA Workshop on Graph Theory and Sparse Matrix Computation, October 1991, Mineapolis.
- [35] R. Yuster and U. Zwick. Answering distance queries in directed graphs using fast matrix multiplication. In *Foundations of Computer Science, 2005. FOCS 2005. 46th Annual IEEE Symposium on*, pages 389–396, Oct 2005.