
Improved Convergence for ℓ_∞ and ℓ_1 Regression via Iteratively Reweighted Least Squares

Alina Ene^{*1} Adrian Vladu^{*1}

Abstract

The iteratively reweighted least squares method (IRLS) is a popular technique used in practice for solving regression problems. Various versions of this method have been proposed, but their theoretical analyses failed to capture the good practical performance.

In this paper we propose a simple and natural version of IRLS for solving ℓ_∞ and ℓ_1 regression, which provably converges to a $(1 + \varepsilon)$ -approximate solution in $O(m^{1/3} \log(1/\varepsilon)/\varepsilon^{2/3} + \log m/\varepsilon^2)$ iterations, where m is the number of rows of the input matrix. Interestingly, this running time is independent of the conditioning of the input, and the dominant term of the running time depends sublinearly in ε^{-1} , which is atypical for the optimization of non-smooth functions.

This improves upon the more complex algorithms of Chin et al. (ITCS '12), and Christiano et al. (STOC '11) by a factor of at least $1/\varepsilon^2$, and yields a truly efficient natural algorithm for the slime mold dynamics (Straszak-Vishnoi, SODA '16, ITCS '16, ITCS '17).

1. Introduction

Regression problems are fundamental primitives in scientific computing. Among these, ℓ_∞ - and ℓ_1 -regression are their hardest instantiations, since through standard reductions they can be shown to be equivalent to linear programming.

While the series of works on these topics is truly extensive and diverse, the simpler methods have pervaded into the realm of practical applications. Among these, an extremely popular scheme known for its simplicity is the iteratively re-weighted least squares (IRLS) method. The idea behind

it is to reduce optimization problems to iteratively solving a series of weighted ℓ_2 -minimization problems, where the weights are adaptively chosen in such a way that the resulting solutions from the sequence of least-squares problems converge to the sought optimal point. In particular, due to its relevance in signal processing, ℓ_1 regression is a very important application of IRLS (Candès et al., 2006; Chartrand & Yin, 2008).

Despite the fact that various versions of this method have been studied ever since the 60's (Lawson, 1961; Osborne, 1985) theoretical understanding of their convergence has lacked. Recent works have attempted to fill this gap, and offer provable guarantees (Daubechies et al., 2010; Straszak & Vishnoi, 2016a;b;c), some of them inspired from the interpretation of this method as a dynamical system. In particular, we note the *Physarum* dynamics, which have been studied in a completely different context (Ito et al., 2011; Johansson & Zou, 2012; Tero et al., 2007; Bonifaci et al., 2012; Becchetti et al., 2013) in order to justify an experiment which revealed that a unicellular organism, the slime mold, could solve the shortest path problem in a maze (Nakagaki et al., 2000). The fact that these dynamics are essentially just a version of the IRLS method was observed in (Straszak & Vishnoi, 2016a).

Returning to the more classical world of algorithm design and analysis, it is worth observing that existing analyses of IRLS methods fall into one of the following two categories: (i) they show convergence only when the problem is properly initialized, or (ii) the guaranteed running time is prohibitive in the sense that it is highly dependent on how the input is conditioned, or it has a high polynomial dependency on the desired solution accuracy.

In this paper, we focus on analyzing simple versions of IRLS which overcome both aforementioned obstacles. In particular, our methods always converge to $1 + \varepsilon$ multiplicative approximation for the objectives $\min_{\mathbf{x}: A\mathbf{x}=\mathbf{b}} \|\mathbf{x}\|_p$, $p \in \{1, \infty\}$, in $\tilde{O}(m^{1/3}/\varepsilon^{2/3} + 1/\varepsilon^2)$ iterations¹ of solving a weighted least squares problem, where m is the dimension of the sought vector \mathbf{x} .

^{*}Equal contribution ¹Boston University, Boston, USA. Correspondence to: Alina Ene <aene@bu.edu>, Adrian Vladu <avladu@mit.edu>.

¹We use \tilde{O} notation to suppress polylogarithmic factors in m/ε .

Inspiration for our methods is heavily drawn from the work of (Christiano et al., 2011), which offered a ground-breaking result by showing that in undirected graphs, a $(1 + \varepsilon)$ -approximate maximum flow can be found in $\tilde{O}(m^{1/3}/\varepsilon^{11/3})$ iterations (subsequently the ε dependence was improved to $1/\varepsilon^{8/3}$, see (Chin et al., 2013)) of solving a weighted least squares problem – which in conjunction with efficient Laplacian system solvers, broke a longstanding barrier for fast graph algorithms. While these algorithms generalize to arbitrary ℓ_1 and ℓ_∞ regression problems, they are somewhat involved, in particular due to the fact that they are the product of combining the multiplicative weights update method with a regularization technique, and a second potential function.²

Instead, our method attempts to directly solve the non-smooth objective while tracking a single potential function. The number of iterations looks surprising, since the dominant term is $\tilde{O}(m^{1/3}/\varepsilon^{2/3})$, whenever $\varepsilon \geq m^{-1/4}$, while classical techniques for optimizing non-smooth functions require a number of iterations that depends on the product between the function’s parameters (such as Lipschitz constant of the gradient or radius of the domain), and $1/\varepsilon$ in the best case, when accelerated methods are used; see (Nesterov, 2005) for more details.³

Interestingly, a line of work that yielded results very similar in spirit to ours is that of approximately solving positive linear and semidefinite programs (Young, 2001; Allen-Zhu & Orecchia, 2015; Allen-Zhu et al., 2016), where the goal was to produce a first order optimization method that can be implemented in a number of iterations independent of the conditioning of the input. Improving the ε dependence to $o(1/\varepsilon^2)$ is an important open problem in this subfield.

We believe that understanding the connection between these results can pave the way for designing new efficient optimization primitives.

²To be more specific, Christiano et al. solve the approximate maximum flow problem, which is a specific instance of ℓ_∞ regression. Chin et al. build on this work to solve ℓ_1 regression with block structure; the block structure is relevant for their specific applications, but is a direct extension of the method, so solving vanilla ℓ_1 regression is still the main problem tackled there.

³We emphasize that using off-the-shelf methods, without further assumptions on the input, the number of iterations of any standard optimization method would be $\Omega(\sqrt{m})$ even for the very special instances where the affine constraint corresponds to a flow satisfying a given demand in unweighted graphs, and in general will depend on how the input matrix is conditioned, since this conditioning determines the magnitude of the subgradients or the diameter of the domain we are optimizing over. The breakthrough of Christiano et al. was the first work that managed to reduce this dependence for maximum flow, which is a specific instance of the ℓ_∞ regression problem.

1.1. Main Theorem

We state the main theorem of this paper. It follows from the correctness proofs described in Sections 3.1 and 3.2, and the convergence proofs from Lemmas A.6 and A.8.

Theorem 1.1. *There exist algorithms ℓ_∞ -MINIMIZATION and ℓ_1 -MINIMIZATION such that, on input $(\mathbf{A}, \mathbf{b}, \varepsilon, M)$, where $\mathbf{A} \in \mathbb{R}^{n \times m}$ is a matrix, $\mathbf{b} \in \mathbb{R}^n$ is a vector which lies in the span of \mathbf{A} ’s columns, ε is an accuracy parameter, and M is a target value:*

1. ℓ_∞ -MINIMIZATION returns a solution \mathbf{x} such that $\mathbf{Ax} = \mathbf{b}$, and $\|\mathbf{x}\|_\infty \leq (1 + \varepsilon)M$, or certifies that $\min_{\mathbf{x}: \mathbf{Ax}=\mathbf{b}} \|\mathbf{x}\|_\infty \geq (1 - \varepsilon)M$.
2. ℓ_1 -MINIMIZATION returns a solution \mathbf{x} such that $\mathbf{Ax} = \mathbf{b}$, and $\|\mathbf{x}\|_1 \leq (1 + \varepsilon)M$, or certifies that $\min_{\mathbf{x}: \mathbf{Ax}=\mathbf{b}} \|\mathbf{x}\|_1 \geq (1 - \varepsilon)M$.

Furthermore both algorithms finish in

$$O\left(\frac{m^{1/3} \log(1/\varepsilon)}{\varepsilon^{2/3}} + \frac{\log m}{\varepsilon^2}\right)$$

iterations, each of which can be implemented in the time required to solve a linear system of the form $\mathbf{ADA}^\top \phi = \mathbf{b}$, where $\mathbf{D} \in \mathbb{R}^{m \times m}$ is an arbitrary nonnegative diagonal matrix.

While our theorem statements are concerned with approximately solving a decision problem which requires a guess M on the value of the objective, it follows from standard techniques that this can be used to find a good approximation to the optimal solution without paying more than a $\tilde{O}(1)$ overhead in the number of iterations. For completeness, we provide the details in Section D.

1.2. Relation to Previous IRLS Methods and Slime-Mold Dynamics

A popular method for solving ℓ_1 minimization is the iteratively re-weighted least squares method (IRLS). This is essentially based on the observation that whenever $\mathbf{x}^* = \arg \min_{\mathbf{Ax}=\mathbf{b}} \|\mathbf{x}\|_1$, one also has that this is the minimizer of the least squares problem $\arg \min_{\mathbf{x}: \mathbf{Ax}=\mathbf{b}} \langle 1/\mathbf{x}^*, \mathbf{x}^2 \rangle$.⁴ Hence one approach that has been employed ever since the 60’s (Lawson, 1961; Osborne, 1985; Daubechies et al., 2010) is to iteratively adjust the weighting of the coordinates and re-solve the least squares problem, until \mathbf{x} converges to a stationary point. This is rigorously described by the iteration

$$\mathbf{x}^{(t+1)} = \arg \min_{\mathbf{Ax}=\mathbf{b}} \left\langle \frac{1}{|\mathbf{x}^{(t)}|}, \mathbf{x}^2 \right\rangle,$$

⁴Throughout the paper we use the convention that $0/0 = 0$.

We abuse notation by applying scalar operations to vectors, with the meaning that they are applied element-wise.

Subsequent works attempted to rigorously analyze this iteration and prove convergence bounds. Oftentimes this relied on specific structure, such as \mathbf{x} being sparse (Daubechies et al., 2010). A recent series of works drew inspiration from convergence proofs for the slime-mold dynamics – a method which essentially solves ℓ_1 minimization, based on a model used to describe the evolution of a slime mold (*Physarum polycephalum*) as it spreads through its environment in order to optimize its access to food sources (Nakagaki et al., 2000; Tero et al., 2007). Based on the intuition that these dynamics yield a method for solving the transportation problem, Straszak and Vishnoi proved in a series of works (Straszak & Vishnoi, 2016a;b;c) that this is as a matter of fact equivalent to the IRLS method, and provided a rigorous convergence analysis for a damped version of it:

$$\mathbf{x}^{(t+1)} = \arg \min_{\mathbf{Ax}=\mathbf{b}} \left\langle \frac{1}{\sqrt{(\mathbf{x}^{(t)})^2 + \eta^2}}, \mathbf{x}^2 \right\rangle.$$

Unfortunately their convergence proof shows that this method is highly inefficient, and the time to convergence has a high polynomial dependence in the desired accuracy, and the structure of the linear constraint.

By comparison, what we describe in this work is an IRLS method where the weights are updated according to a thresholding rule. Given a guess M for the optimal value, we perform an iteration equivalent to:

$$c_i^{(t+1)} = c_i^{(t)} \cdot \psi_{1/(1-\varepsilon)} \left(\frac{x_i^{(t)}/c_i^{(t)}}{\left\langle \frac{1}{c^{(t)}}, (\mathbf{x}^{(t)})^2 \right\rangle} \cdot M \right)^2,$$

$$\mathbf{x}^{(t+1)} = \arg \min_{\mathbf{Ax}=\mathbf{b}} \left\langle \frac{1}{\mathbf{c}^{(t+1)}}, \mathbf{x}^2 \right\rangle,$$

where ψ is a thresholding operator i.e. $\psi_b(u) = u$, if $u \geq b$, and $\psi_b(u) = 1$ otherwise. Intuitively, this increases the weights c_i only for the elements where the corresponding component x_i^2/c_i of the quadratic objective contributes significantly, therefore we want to favor increasing it even more in the future by decreasing the weight $1/c_i$ we place on this coordinate.⁵

⁵Another way to think of this is that, ignoring the thresholding operator, the update would simply be $c_i^{(t+1)} = (x_i^{(t)})^2/c_i^{(t)} \cdot \gamma$, where γ is some normalization factor. What thresholding achieves here is to decide whether the contribution of a particular coordinate to the energy of the system is sufficiently large compared to the contributions of the entire vector \mathbf{x} .

2. Preliminaries

2.1. Basic Notation

Sets. We let \mathbb{R} be the set of real numbers. For any natural number n , we write $[n] := \{1, \dots, n\}$. We denote by Δ_m the m -dimensional simplex i.e. $\Delta_m = \{\mathbf{p} \in \mathbb{R}^m : \sum_{i=1}^m p_i = 1, p_i \geq 0 \text{ for all } i\}$.

Vectors. We let $\mathbf{0}, \mathbf{1} \in \mathbb{R}^n$ denote the all zeros and all ones vectors, respectively. When it is clear from the context, we apply scalar operations to vectors with the interpretation that they are applied coordinate-wise.

Matrices. We write matrices in bold. We use \mathbf{I} to denote the identity matrix. Given a vector \mathbf{x} we let $\mathbf{D}(\mathbf{x})$ be the diagonal matrix whose entries are given by \mathbf{x} . For a symmetric matrix \mathbf{A} , we let \mathbf{A}^+ be its Moore-Penrose pseudoinverse, i.e. $\mathbf{AA}^+ = \mathbf{A}^+\mathbf{A} = \mathbf{I}_{\text{Im}(\mathbf{A})}$. The pseudoinverse can be thought of as replacing all the nonzero eigenvalues of \mathbf{A} with their reciprocals.

Inner products. When it is convenient, we use $\langle \cdot, \cdot \rangle$ notation to denote inner products. Given two vectors \mathbf{x}, \mathbf{y} of equal dimensions, we let $\langle \mathbf{x}, \mathbf{y} \rangle = \mathbf{x}^\top \mathbf{y}$.

Norms. Given a vector \mathbf{x} , we denote the ℓ_p norm of \mathbf{x} by $\|\mathbf{x}\|_p = (\sum x_i^p)^{1/p}$. When the subscript is dropped, we refer to the ℓ_2 norm. From this definition, we can also see that $\|\mathbf{x}\|_\infty = \max_i |x_i|$.

2.2. Proof Technique

Let us first understand the idea behind our ℓ_∞ minimization algorithm. The problem we aim to solve is $\min_{\mathbf{x}:\mathbf{Ax}=\mathbf{b}} \|\mathbf{x}\|_\infty$. Letting Δ_m be the m -dimensional unit simplex, we can write our objective equivalently as

$$\begin{aligned} \min_{\mathbf{x}:\mathbf{Ax}=\mathbf{b}} \|\mathbf{x}^2\|_\infty &= \min_{\mathbf{x}:\mathbf{Ax}=\mathbf{b}} \max_{\mathbf{r} \in \Delta_m} \langle \mathbf{r}, \mathbf{x}^2 \rangle \\ &= \max_{\mathbf{r} \in \Delta_m} \left(\min_{\mathbf{x}:\mathbf{Ax}=\mathbf{b}} \langle \mathbf{r}, \mathbf{x}^2 \rangle \right) := \max_{\mathbf{r} \in \Delta_m} \mathcal{E}_{\mathbf{r}}(\mathbf{b}), \end{aligned}$$

where the second identity follows from Sion's theorem (Sion, 1958), which allows us to interchange min and max. The quantity between the parentheses has a very natural interpretation, in the case of electrical networks: it is precisely the electrical energy required to route a demand \mathbf{b} through an electrical network encoded in \mathbf{A} . Furthermore, we have an easy way to lower bound how this energy increases whenever resistances are increased, which is a finer quantitative version of Rayleigh's monotonicity principle. More precisely, we can easily certify a lower bound on the increase in energy determined by increasing a single coordinate of \mathbf{r} . Using this observation, which we make more precise in Section A.2, we can identify a set of coordinates

of \mathbf{r} to increase, guaranteeing that if \mathbf{r}' is the new vector with perturbed resistances, we have

$$\frac{\mathcal{E}_{\mathbf{r}'}(\mathbf{b}) - \mathcal{E}_{\mathbf{r}}(\mathbf{b})}{\|\mathbf{r}' - \mathbf{r}\|_1} \geq M^2, \quad (2.1)$$

for a fixed parameter M . In the case when no coordinates of \mathbf{r} can be increased, while preserving this property, this yields a certificate that \mathbf{r} is as a matter of fact (close to) optimal, and thus we are done (Lemma A.5). Hence our goal becomes that of guaranteeing that $\|\mathbf{r}\|_1$ increases very fast. Indeed, since the "electrical energy" increases at the right rate relative to $\|\mathbf{r}\|_1$, after the latter has increased sufficiently, we can safely guarantee that $\mathcal{E}_{\mathbf{r}}(\mathbf{b})/\|\mathbf{r}\|_1 \geq (1 - \varepsilon)M$, since the increase in $\|\mathbf{r}\|_1$ cancels out most of the initial error introduced by starting with a potentially poor solution.

The ℓ_1 minimization algorithm relies on squaring the objective, and then writing it equivalently as

$$\begin{aligned} \min_{\mathbf{x}: \mathbf{A}\mathbf{x}=\mathbf{b}} \|\mathbf{x}\|_1^2 &= \min_{\mathbf{x}: \mathbf{A}\mathbf{x}=\mathbf{b}} \left(\min_{\mathbf{c} \in \Delta_m} \left\langle \frac{1}{\mathbf{c}}, \mathbf{x}^2 \right\rangle \right) \\ &= \min_{\mathbf{c} \in \Delta_m} \left(\min_{\mathbf{x}: \mathbf{A}\mathbf{x}=\mathbf{b}} \left\langle \frac{1}{\mathbf{c}}, \mathbf{x}^2 \right\rangle \right) = \min_{\mathbf{c} \in \Delta_m} \mathcal{E}_{1/\mathbf{c}}(\mathbf{b}). \end{aligned}$$

For the first identity we used the fact that $\|\mathbf{x}\|_1^2 = \min_{\mathbf{c} \in \Delta_m} \langle 1/\mathbf{c}, \mathbf{x}^2 \rangle$, achieved at $\mathbf{c} = \mathbf{x}/\|\mathbf{x}\|_1$; see (Owen, 2007; Sun & Zhang, 2012) for further use of this trick.⁶ The second identity follows from joint convexity w.r.t. \mathbf{c} and \mathbf{x} , which can be verified by computing the Hessian of the function in (\mathbf{x}, \mathbf{c}) . So completely oppositely from the previous case, the objective of our problem becomes minimizing electrical energy with respect to a set of inverse resistances, which we will call conductances. Note that in this case the quantity that is invariant under scaling \mathbf{c} by a constant is $\mathcal{E}_{1/\mathbf{c}} \cdot \|\mathbf{c}\|_1$. Therefore, equivalently, our goal will be to find the set of conductances $\mathbf{c} \geq 0$ for which $(\mathcal{E}_{1/\mathbf{c}})^{-1}/\|\mathbf{c}\|_1 \geq \frac{1}{(1+\varepsilon)M}$. Similarly to the ℓ_∞ case, in this case we make progress by iteratively increasing conductances from \mathbf{c} to \mathbf{c}' in such a way that

$$\frac{\frac{1}{\mathcal{E}_{\mathbf{c}'}(\mathbf{b})} - \frac{1}{\mathcal{E}_{\mathbf{c}}(\mathbf{b})}}{\|\mathbf{c}' - \mathbf{c}\|_1} \geq \frac{1}{M^2}. \quad (2.2)$$

Just as before, we can prove that unless the value of the objective can not be made smaller than M , then \mathbf{c} can be increased while enforcing this invariant (Lemma A.7). Hence we can prove fast convergence by arguing that $\|\mathbf{c}\|_1$ increases very fast.

⁶Interestingly, this can also be thought of as achieving tightness for reverse Hölder's inequality whenever we are considering the dual 'norms' ℓ_{-1} and $\ell_{1/2}$.

2.3. Approximate Solutions and Infeasibility Certificates

ℓ_∞ minimization We consider the formulation

$$\min_{\mathbf{A}\mathbf{x}=\mathbf{b}} \|\mathbf{x}\|_\infty, \quad (2.3)$$

for which we seek an approximate solution in the following sense. Given a target value M , we aim to find one of the following:

1. an approximate solution \mathbf{x} in the sense that $\mathbf{A}\mathbf{x} = \mathbf{b}$ and $\|\mathbf{x}\|_\infty \leq (1 + \varepsilon)M$,
2. an approximate infeasibility certificate \mathbf{r} in the sense that $\mathbf{r} \in \Delta_m$ and $\mathbf{b}^\top (\mathbf{A}\mathbf{D}(\mathbf{r})^{-1} \mathbf{A}^\top) + \mathbf{b} \geq (1 - \varepsilon)^2 M^2$.

We prove in Lemma 2.1 that the latter is indeed an infeasibility certificate.

Lemma 2.1. *Let \mathbf{x}^* be the solution to the problem defined in Equation 2.3, and let $\mathbf{r} \in \Delta_m$. Then $\|\mathbf{x}^*\|_\infty^2 \geq \mathbf{b}^\top (\mathbf{A}\mathbf{D}(\mathbf{r})^{-1} \mathbf{A}^\top) + \mathbf{b}$.*

Proof. Using Lemma A.2 we can write

$$\begin{aligned} \mathbf{b}^\top (\mathbf{A}\mathbf{D}(\mathbf{r})^{-1} \mathbf{A}^\top) + \mathbf{b} &= \min_{\mathbf{x}: \mathbf{A}\mathbf{x}=\mathbf{b}} \langle \mathbf{r}, \mathbf{x}^2 \rangle \leq \langle \mathbf{r}, (\mathbf{x}^*)^2 \rangle \\ &\leq \|\mathbf{r}\|_1 \|\mathbf{x}^*\|_\infty^2 = \|\mathbf{x}^*\|_\infty^2, \end{aligned}$$

which gives us what we needed. \square

ℓ_1 minimization We consider the formulation

$$\min_{\mathbf{A}\mathbf{x}=\mathbf{b}} \|\mathbf{x}\|_1, \quad (2.4)$$

for which seek an approximate solution in the following sense. Given a target value M , we seek one of the following:

1. an approximate infeasibility certificate $\phi \in \mathbb{R}^n$ in the sense that $\frac{\mathbf{b}^\top \phi}{\|\mathbf{A}^\top \phi\|_\infty} \geq (1 - \varepsilon)M$,
2. an approximate feasibility certificate \mathbf{c} in the sense that $\mathbf{c} \in \Delta_m$ and $\mathbf{b}^\top (\mathbf{A}\mathbf{D}(\mathbf{c}) \mathbf{A}^\top) + \mathbf{b} \leq (1 + \varepsilon)^2 M^2$, which yields an approximately feasible solution $\mathbf{x} = \mathbf{D}(\mathbf{c}) \mathbf{A}^\top (\mathbf{A}\mathbf{D}(\mathbf{c}) \mathbf{A}^\top) + \mathbf{b}$ in the sense that $\mathbf{A}\mathbf{x} = \mathbf{b}$ and $\|\mathbf{x}\|_1 \leq (1 + \varepsilon)M$.

The fact that the former is an approximate infeasibility certificate follows from convex duality. Indeed, one can see that the dual of the minimization problem is $\max_{\phi: \|\mathbf{A}^\top \phi\|_\infty \leq 1} \mathbf{b}^\top \phi$, so exhibiting a solution as above implies that the value of this objective is at least $(1 - \varepsilon)M$. A proof for the fact that the latter is indeed an approximate feasibility certificate, and that it yields an approximately feasible solution can be found in Lemma 2.2.

Lemma 2.2. Given $\mathbf{c} \in \Delta_m$, the vector $\mathbf{x} = \mathbf{D}(\mathbf{c})\mathbf{A}^\top(\mathbf{AD}(\mathbf{c})\mathbf{A}^\top)^+\mathbf{b}$ satisfies $\mathbf{Ax} = \mathbf{b}$, and $\|\mathbf{x}\|_1^2 \leq \mathbf{b}^\top(\mathbf{AD}(\mathbf{c})\mathbf{A}^\top)^+\mathbf{b}$.

Proof. The fact that $\mathbf{Ax} = \mathbf{b}$ follows directly by substitution, and using the fact that $\mathbf{b} \in \text{Im}(\mathbf{A})$. Using Lemma A.2 and the definition in (A.1) we write

$$\begin{aligned} & \mathbf{b}^\top(\mathbf{AD}(\mathbf{c})\mathbf{A}^\top)^+\mathbf{b} \\ &= \mathbf{b}^\top(\mathbf{AD}(\mathbf{c})\mathbf{A}^\top)^+(\mathbf{AD}(\mathbf{c})\mathbf{A}^\top)(\mathbf{AD}(\mathbf{c})\mathbf{A}^\top)^+\mathbf{b} \\ &= \sum_{i=1}^m \frac{1}{c_i} \cdot \left(\mathbf{D}(\mathbf{c})\mathbf{A}^\top(\mathbf{AD}(\mathbf{c})\mathbf{A}^\top)^+\mathbf{b} \right)^2 = \sum_{i=1}^m \frac{1}{c_i} \cdot x_i^2. \end{aligned}$$

We can use this identity inside the following upper bound, which we obtain by applying Cauchy-Schwarz:

$$\begin{aligned} \|\mathbf{x}\|_1 &= \sum_{i=1}^m \frac{|x_i|}{\sqrt{c_i}} \cdot \sqrt{c_i} \leq \sqrt{\left(\sum_{i=1}^m \frac{x_i^2}{c_i} \right) \left(\sum_{i=1}^m c_i \right)} \\ &= \sqrt{\mathbf{b}^\top(\mathbf{AD}(\mathbf{c})\mathbf{A}^\top)^+\mathbf{b}}. \end{aligned}$$

This yields our claim. \square

3. The Algorithms

Having introduced the necessary notation, we can describe our simple IRLS routine. We prove convergence in Section A.

3.1. The ℓ_∞ Minimization Algorithm

We first present the algorithm for the ℓ_∞ version of the problem, since it is the most intuitive. The method attempts to find a weighting of the columns of \mathbf{A} i.e. a vector $\mathbf{r} \in \mathbb{R}^m$ for which the corresponding least squares solution has a small ℓ_∞ norm; more precisely $\|\mathbf{x}\|_\infty / \|\mathbf{r}\|_1 \leq (1 + \varepsilon)M$ for some chosen target value M .

Then the weighting is updated via the following simple thresholding rule. Elements for which the corresponding coordinate of the least squares solution x_i is below the desired target value are left unchanged. The others are scaled exactly by the amount by which the square of the corresponding coordinate x_i violates the desired threshold i.e. x_i^2/M^2 .

Note that the iteration defined here simply attempts to construct an infeasibility certificate for the problem defined in Equation 2.3. Building the feasible solution involves maintaining a solution obtained by uniformly averaging a subset of the iterates \mathbf{x} witnessed so far. These are used to return the approximately feasible solution in case the algorithm fails to quickly produce an (approximate) infeasibility certificate. The details referring to how and why we perform

Algorithm 1 ℓ_∞ -MINIMIZATION($\mathbf{A}, \mathbf{b}, \varepsilon, M$)

```

1: Input: Matrix  $\mathbf{A} \in \mathbb{R}^{n \times m}$ , vector  $\mathbf{b} \in \mathbb{R}^n$ , accuracy  $\varepsilon$ , target value  $M$ .
2: Output: Vector  $\mathbf{x}$  such that  $\mathbf{Ax} = \mathbf{b}$  and  $\|\mathbf{x}\|_\infty \leq (1 + \varepsilon)M$ , or approximate infeasibility certificate  $\mathbf{r} \in \Delta_m$ .
3:  $t = 0$ ,  $\mathbf{r}^{(0)} = \mathbf{1}/m$ .
4:  $t' = 0$ ,  $\mathbf{s}^{(t')} = \vec{0}$ .
5: while  $\|\mathbf{r}^{(t)}\|_1 \leq 1/\varepsilon$  do
6:    $\mathbf{x}^{(t)} = \arg \min_{\mathbf{x}: \mathbf{Ax}=\mathbf{b}} \langle \mathbf{r}, \mathbf{x}^2 \rangle$ .
   // Equivalently,
    $\mathbf{x}^{(t)} = \mathbf{D}(\mathbf{r})^{-1} \mathbf{A}^\top \left( \mathbf{AD}(\mathbf{r})^{-1} \mathbf{A}^\top \right)^+ \mathbf{b}$ .
7:   if  $\|\mathbf{x}^{(t)}\|_\infty \leq m^{1/3} \cdot M$  then
8:      $t' = t' + 1$ ,  $\mathbf{s}^{(t')} = \mathbf{s}^{(t'-1)} + \mathbf{x}^{(t)}$ .
9:   end if
10:  if  $\|\mathbf{s}^{(t')}\|_\infty / t' \leq (1 + \varepsilon)M$  then
11:    return  $\mathbf{s}^{(t')}/t'$ .
12:  end if
13:   $\alpha_i^{(t)} = \begin{cases} 1 & \text{if } |x_i^{(t)}| < (1 + \varepsilon)M, \\ \frac{(x_i^{(t)})^2}{M^2} & \text{otherwise.} \end{cases}$ 
14:  if  $\alpha^{(t)} = \mathbf{1}$  then
15:    return  $\mathbf{x}^{(t)}$ .
16:  end if
17:   $\mathbf{r}^{(t+1)} = \mathbf{r}^{(t)} \cdot \alpha^{(t)}$ .
18:   $t = t + 1$ .
19: end while
20: return  $\mathbf{r} / \|\mathbf{r}\|_1$ .

```

this specific set of updates are explained in the convergence proof. The steps involved in building this feasible solution are written in blue. They can be ignored if the goal is simply that of returning a yes/no answer.

Correctness. We notice that Algorithm 1 has two possible outcomes. Either it returns a primal approximately feasible vector (lines 11 and 15), or returns a dual certificate (line 20). In the former case, it is clear from the description of the algorithm that the returned vector is indeed approximately feasible: line 11 returns a uniform average of vectors satisfying the linear constraint with small ℓ_∞ norm; line 20 returns the $\mathbf{x}^{(t)}$ computed within the corresponding iteration, whenever $\alpha^{(t)} = \vec{1}$, i.e. $\|\mathbf{x}^{(t)}\|_\infty < (1 + \varepsilon)M$.

Also, note that in case none of these stopping conditions is triggered, the algorithm returns a dual certificate on line 20 after a finite number of iterations. Indeed, note that every iteration where $\alpha_i^{(t)} \neq \vec{1}$, at least one element of $\mathbf{r}^{(t)}$ gets increased by a factor of at least $(1 + \varepsilon)^2$, due to way $\alpha^{(t)}$ is defined. Since the algorithm stops when $\|\mathbf{r}^{(t)}\|_1 = 1/\varepsilon$, no element of \mathbf{r} can be scaled more than $O(\log_{(1+\varepsilon)}(m/\varepsilon))$ times, hence the total number of iterations is very roughly upper bounded by $O(m \log(m/\varepsilon)/\varepsilon)$. We will see in Sec-

Algorithm 2 ℓ_1 -MINIMIZATION($\mathbf{A}, \mathbf{b}, \varepsilon, M$)

```

1: Input: Matrix  $\mathbf{A} \in \mathbb{R}^{n \times m}$ , vector  $\mathbf{b} \in \mathbb{R}^n$ , accuracy  $\varepsilon$ ,
   target value  $M$ .
2: Output: Vector  $\mathbf{x}$  such that  $\mathbf{Ax} = \mathbf{b}$  and  $\|\mathbf{x}\|_1 \leq (1 + \varepsilon)M$ ,
   or approximate infeasibility certificate  $\phi \in \Delta_n$ .
3:  $t = 0$ ,  $\mathbf{c}^{(0)} = \mathbf{1}/m$ .
4:  $t' = 0$ ,  $\mathbf{s}^{(t')} = \mathbf{0}$ ,  $\Phi^{(0)} = \mathbf{0}$ .
5: while  $\|\mathbf{c}^{(t)}\|_1 \leq 1 + \frac{1}{(1+\varepsilon)^2-1}$  do
6:    $\phi^{(t)} = \left(\mathbf{AD}(\mathbf{c})\mathbf{A}^\top\right)^+ \mathbf{b}$ . // Equivalently,
      $\phi^{(t)}$  is the vector of potentials
     which induce the electrical
     flow  $\mathbf{x} = \arg \min_{\mathbf{Ax}=\mathbf{b}} \langle \mathbf{1}/\mathbf{c}, \mathbf{x}^2 \rangle$  via
      $\mathbf{x} = \mathbf{D}(\mathbf{c})\mathbf{A}^\top \phi$ .
7:   if  $\left\| \frac{\mathbf{A}^\top \phi^{(t)}}{\mathbf{b}^\top \phi^{(t)}} \right\|_\infty \leq m^{1/3} \cdot \frac{1}{M}$  then
8:      $t' = t' + 1$ ,  $\mathbf{s}^{(t')} = \mathbf{s}^{(t'-1)} + \left| \frac{\mathbf{A}^\top \phi^{(t)}}{\mathbf{b}^\top \phi^{(t)}} \right|$ ,  $\Phi^{(t')} =$ 
        $\Phi^{(t'-1)} + \frac{\phi^{(t)}}{\mathbf{b}^\top \phi^{(t)}}$ .
9:   end if
10:  if  $\|\mathbf{s}^{(t')}\|_\infty / t' \leq \frac{1}{(1-\varepsilon)M}$  then
11:    return  $\Phi^{(t')}/t'$ .
12:  end if
13:   $\alpha_i^{(t)} = \begin{cases} 1 & \text{if } \frac{|\mathbf{A}^\top \phi^{(t)}|_i}{\mathbf{b}^\top \phi^{(t)}} \leq \frac{1}{(1-\varepsilon)M}, \\ \left( \frac{(\mathbf{A}^\top \phi^{(t)})_i}{\mathbf{b}^\top \phi^{(t)}} \right)^2 \cdot M^2 & \text{otherwise.} \end{cases}$ 
14:  if  $\alpha^{(t)} = \mathbf{1}$  then
15:    return  $\phi^{(t)}$ .
16:  end if
17:   $\mathbf{c}^{(t+1)} = \mathbf{c}^{(t)} \cdot \alpha^{(t)}$ .
18:   $t = t + 1$ .
19: end while
20: return  $\mathbf{x} = \mathbf{D}(\mathbf{c})\mathbf{A}^\top \phi^{(t)}$ .

```

tion A that we can prove a much finer upper bound.

Finally, we need to argue that whenever the algorithm returns on line 20, it returns an infeasibility certificate as per Lemma 2.1. We defer the proof to Lemma A.5 in Section A.

3.2. The ℓ_1 Minimization Algorithm

The ℓ_1 version is very similar. As a matter of fact, it can be re-derived simply by attempting to solve the convex dual of the problem from (2.3), which is an ℓ_∞ minimization problem, by using the routine from Figure 1. However, since the reduction requires several, and previous works attempted to solve this directly using various versions of IRLS, we provide a natural iteration which does not involve any reductions.

Correctness. We notice that Algorithm 2 has two possible outcomes. Either it returns an approximate infeasibility

certificate (lines 11 and 15), or returns an approximately feasible solution (line 20).

Let us verify that in the former case the returned vector is indeed an approximate infeasibility certificate. Line 11 returns $\Phi^{(t')} = \sum_{t \in S} \frac{\phi^{(t)}}{\mathbf{b}^\top \phi^{(t)}}$, where we know that S is a set for which

$$\begin{aligned} \left\| \mathbf{A}^\top \Phi^{(t')} \right\|_\infty &= \left\| \mathbf{A}^\top \cdot \sum_{t \in S} \frac{\phi^{(t)}}{\mathbf{b}^\top \phi^{(t)}} \right\|_\infty \\ &= \left\| \sum_{t \in S} \frac{\mathbf{A}^\top \phi^{(t)}}{\mathbf{b}^\top \phi^{(t)}} \right\|_\infty \leq \left\| \sum_{t \in S} \left| \frac{\mathbf{A}^\top \phi^{(t)}}{\mathbf{b}^\top \phi^{(t)}} \right| \right\|_\infty \leq \frac{t'}{(1-\varepsilon)M}. \end{aligned}$$

Since $\mathbf{b}^\top \Phi^{(t')} = t'$, we see that returned vector $\Phi^{(t')}/t'$ is an approximate infeasibility certificate, as defined in Section 2.3. If the algorithm returns on line 15, we get that $\left\| \frac{\mathbf{A}^\top \phi^{(t)}}{\mathbf{b}^\top \phi^{(t)}} \right\|_\infty \leq \frac{1}{(1-\varepsilon)M}$, hence $\phi^{(t)}$ is an approximate infeasibility certificate.

Also, note that in case none of these stopping conditions is triggered, the algorithm returns a solution on line 20 after a finite number of iterations. Indeed, just as in the ℓ_∞ case, in every iteration some conductance gets increased by a factor of at least $\Omega(1 + \varepsilon)$, hence the algorithm must stop in finite time. We provide a rigorous analysis of the time required for convergence in Section A.

Finally, we need to argue that whenever the algorithm returns a solution on line 20, it is indeed an approximately feasible solution. We defer the proof to Lemma A.7 in Section A.

4. Experiments

We test both our resistance/conductance update schemes in order to verify that the resulting algorithms converge fast in practice. We slightly modify the schemes such that they always update their target value M depending on the value of the objective they have achieved so far. We stop when given the history of witnessed iterates, we can certify a sufficiently small duality gap. For solving linear systems, we used the conjugate gradient implementation from the ℓ_1 -MAGIC optimization suite (Candès & Romberg).

We test both algorithms while varying ε , and varying m . We consider both the update scheme given by our algorithms from Section 3, and one where we attempt to double the length of the step for as long as the invariants from (2.1) and (2.2), respectively, are maintained. We notice that in general, using this long step strategy, we improve both the number of iterations and the running time.

The plots corresponding to the standard update scheme (short-steps) are drawn in red, those corresponding to the long-step version are drawn in blue.

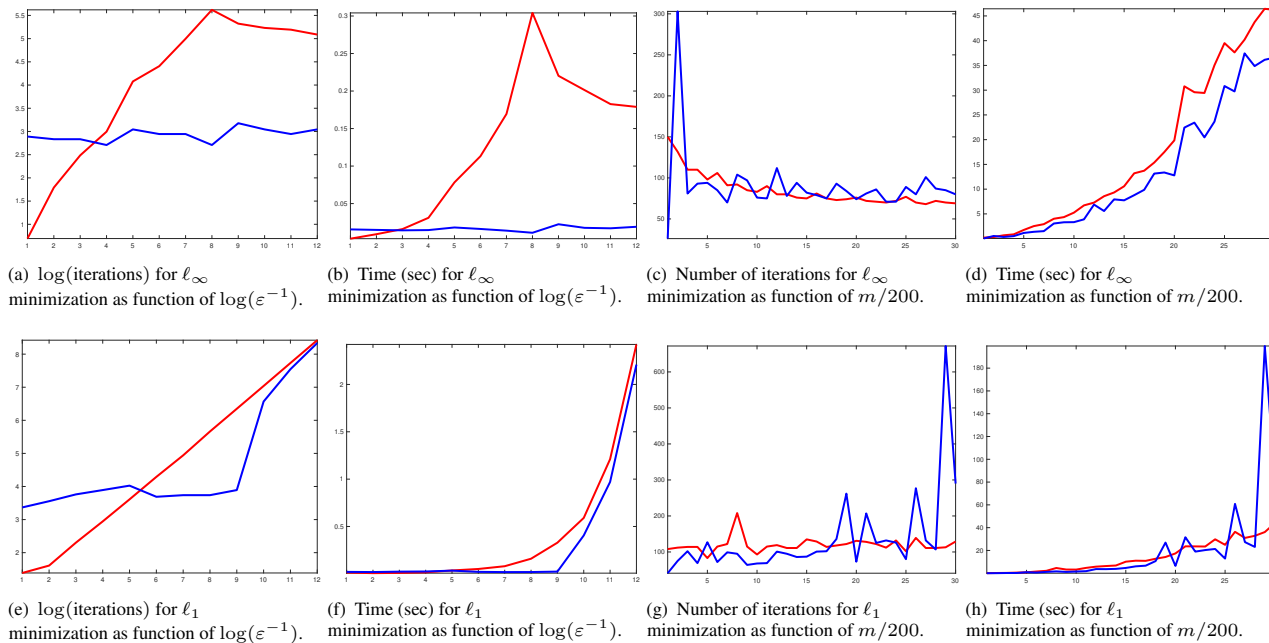


Figure 3.1: Experimental results.

The experiments where we varied ε are reported in figures 1(a), 1(b), 1(e), and 1(f). For all these experiments, the input consists a random 150×200 matrix \mathbf{A} with orthogonal rows, and a vector \mathbf{b} obtained from applying \mathbf{A} to a ± 1 -vector of sparsity 15. We plot the number of iterations/running time of the algorithm for $\varepsilon = 1/2^k$, where $k \in \{1, \dots, 12\}$.

We notice that for these experiments, the number of iterations for the short-step version does indeed scale linearly with ε^{-1} ; the long-step version makes significant gains in the ℓ_∞ case.

The experiments where we varied m are reported in figures 1(c), 1(d), 1(g), and 1(h). For all these experiments, the input consists of a random $150 \times (200 \cdot k)$ matrix \mathbf{A} with orthogonal vectors, and a vector \mathbf{b} obtained from applying \mathbf{A} to a ± 1 -vector of sparsity 15, and a fixed accuracy $\varepsilon = .01$. We plot the number of iterations required by the algorithm for $k \in \{1, \dots, 30\}$.

We notice that for these experiments, both the number of iterations and the running time scale significantly better than by $m^{1/3}$, which suggests that this polynomial dependence in m depends on the input structure, and can be avoided in practice.

Acknowledgements

AE was partially supported by NSF CAREER grant CCF-1750333 and NSF grant CCF-1718342. AV was partially

supported by NSF grant CCF-1718342.

References

Allen-Zhu, Z. and Orecchia, L. Using optimization to break the epsilon barrier: A faster and simpler width-independent algorithm for solving positive linear programs in parallel. In *Proceedings of the twenty-sixth annual ACM-SIAM symposium on Discrete algorithms*, pp. 1439–1456. Society for Industrial and Applied Mathematics, 2015.

Allen-Zhu, Z., Lee, Y. T., and Orecchia, L. Using optimization to obtain a width-independent, parallel, simpler, and faster positive sdp solver. In *Proceedings of the twenty-seventh annual ACM-SIAM symposium on Discrete algorithms*, pp. 1824–1831. Society for Industrial and Applied Mathematics, 2016.

Becchetti, L., Bonifaci, V., Dirnberger, M., Karrenbauer, A., and Mehlhorn, K. Physarum can compute shortest paths: Convergence proofs and complexity bounds. In *International Colloquium on Automata, Languages, and Programming*, pp. 472–483. Springer, 2013.

Bonifaci, V., Mehlhorn, K., and Varma, G. Physarum can compute shortest paths. *Journal of theoretical biology*, 309:121–133, 2012.

Candès, E. and Romberg, J. ℓ_1 -MAGIC: Recovery of sparse signals via convex programming. URL <https://statweb.stanford.edu/~candes/llmagic/>.

- Candès, E. J. et al. Compressive sampling. In *Proceedings of the international congress of mathematicians*, volume 3, pp. 1433–1452. Madrid, Spain, 2006.
- Chartrand, R. and Yin, W. Iterative reweighted algorithms for compressive sensing. Technical report, 2008.
- Chin, H. H., Madry, A., Miller, G. L., and Peng, R. Runtime guarantees for regression problems. In Kleinberg, R. D. (ed.), *Innovations in Theoretical Computer Science, ITCS '13, Berkeley, CA, USA, January 9-12, 2013*, pp. 269–282. ACM, 2013. ISBN 978-1-4503-1859-4. doi: 10.1145/2422436.2422469. URL <https://doi.org/10.1145/2422436.2422469>.
- Christiano, P., Kelner, J. A., Madry, A., Spielman, D. A., and Teng, S. Electrical flows, laplacian systems, and faster approximation of maximum flow in undirected graphs. In Fortnow, L. and Vadhan, S. P. (eds.), *Proceedings of the 43rd ACM Symposium on Theory of Computing, STOC 2011, San Jose, CA, USA, 6-8 June 2011*, pp. 273–282. ACM, 2011. ISBN 978-1-4503-0691-1. doi: 10.1145/1993636.1993674. URL <https://doi.org/10.1145/1993636.1993674>.
- Daubechies, I., DeVore, R., Fornasier, M., and Güntürk, C. S. Iteratively reweighted least squares minimization for sparse recovery. *Communications on Pure and Applied Mathematics: A Journal Issued by the Courant Institute of Mathematical Sciences*, 63(1):1–38, 2010.
- Ito, K., Johansson, A., Nakagaki, T., and Tero, A. Convergence properties for the physarum solver. *arXiv preprint arXiv:1101.5249*, 2011.
- Johansson, A. and Zou, J. A slime mold solver for linear programming problems. In *Conference on Computability in Europe*, pp. 344–354. Springer, 2012.
- Lawson, C. *Contributions to the Theory of Linear Least Maximum Approximation*. University of California, Los Angeles, 1961. URL https://books.google.at/books?id=b_CtbwAACAAJ.
- Nakagaki, T., Yamada, H., and Tóth, Á. Intelligence: Maze-solving by an amoeboid organism. *Nature*, 407(6803):470, 2000.
- Nesterov, Y. Smooth minimization of non-smooth functions. *Mathematical programming*, 103(1):127–152, 2005.
- Osborne, M. R. *Finite algorithms in optimization and data analysis*. 1985.
- Owen, A. B. A robust hybrid of lasso and ridge regression. *Contemporary Mathematics*, 443(7):59–72, 2007.
- Sion, M. On general minimax theorems. *Pacific Journal of mathematics*, 8(1):171–176, 1958.
- Straszak, D. and Vishnoi, N. K. On a natural dynamics for linear programming. In Sudan, M. (ed.), *Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science, Cambridge, MA, USA, January 14-16, 2016*, pp. 291. ACM, 2016a. ISBN 978-1-4503-4057-1. doi: 10.1145/2840728.2840762. URL <https://doi.org/10.1145/2840728.2840762>.
- Straszak, D. and Vishnoi, N. K. Natural algorithms for flow problems. In Krauthgamer, R. (ed.), *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016, Arlington, VA, USA, January 10-12, 2016*, pp. 1868–1883. SIAM, 2016b. ISBN 978-1-61197-433-1. doi: 10.1137/1.9781611974331.ch131. URL <https://doi.org/10.1137/1.9781611974331.ch131>.
- Straszak, D. and Vishnoi, N. K. IRLS and slime mold: Equivalence and convergence. *CoRR*, abs/1601.02712, 2016c. URL <http://arxiv.org/abs/1601.02712>.
- Sun, T. and Zhang, C.-H. Scaled sparse linear regression. *Biometrika*, 99(4):879–898, 2012.
- Tero, A., Kobayashi, R., and Nakagaki, T. A mathematical model for adaptive transport network in path finding by true slime mold. *Journal of theoretical biology*, 244(4):553–564, 2007.
- Young, N. E. Sequential and parallel algorithms for mixed packing and covering. In *Foundations of Computer Science, 2001. Proceedings. 42nd IEEE Symposium on*, pp. 538–546. IEEE, 2001.

A. The Algorithm Analyses

A.1. The Flow/Potential Interpretation

While we study a very general problem, it is very useful to develop intuition based on the case where \mathbf{A} is the vertex-edge incidence matrix of a graph. In this case we will always think of the sought solution \mathbf{x} as a flow on the graph's edges. The corresponding dual object is a set of potentials ϕ defined on the graph's vertices.

To be more precise, we consider the following setting. Let $G = (V, E)$ be an undirected graph. For each edge we choose an arbitrary orientation, and define $E^+(v)$ be the set of arcs leaving vertex v , and $E^-(v)$ the set of arcs entering vertex v , for all v .

Letting $m = |E|$, $n = |V|$, we consider the matrix $\mathbf{A} \in \mathbb{R}^{n \times m}$ where

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

One can easily verify that given a vector $\mathbf{x} \in \mathbb{R}^m$ defined on the arcs of the graph (which we will think of as a **flow**), after applying the operator \mathbf{A} we obtain the demand routed by this flow $\mathbf{Ax} \in \mathbb{R}^n$, which lives in the space of **potentials** defined on the graph's vertices.

Therefore the ℓ_∞ minimization problem from 2.3 can be interpreted as finding the flow \mathbf{x} with minimum congestion which routes the demand \mathbf{b} , while the ℓ_1 minimization problem from 2.4 corresponds to finding the minimum cost flow routing the demand \mathbf{b} .

With this interpretation in mind, we proceed to define some objects that in the case of electrical networks correspond to energy and electrical flows.

We use weightings of \mathbf{A} 's columns $\mathbf{c} \in \mathbb{R}^m$ which we refer to as **conductances**. We equivalently refer to the reciprocals $\mathbf{r} \in \mathbb{R}^m$, with $r_i = 1/c_i$, which we call **resistances**. Our analysis is exclusively based on tracking a potential function which corresponds to the electrical energy of a flow.

Definition A.1 (Energy of a flow). Given a flow $\mathbf{x} \in \mathbb{R}^m$, along with a vector of resistances $\mathbf{r} \in \mathbb{R}^m$, we let the energy of \mathbf{x} be

$$\mathcal{E}_r(\mathbf{x}) = \langle \mathbf{r}, \mathbf{x}^2 \rangle.$$

Overloading this notation, given a vector $\mathbf{b} \in \mathbb{R}^n$, let the **electrical energy** be

$$\mathcal{E}_r(\mathbf{b}) = \min_{\mathbf{x}: \mathbf{Ax}=\mathbf{b}} \mathcal{E}_r(\mathbf{x}), \quad (\text{A.1})$$

in other words this is the minimum energy over all flows satisfying $\mathbf{Ax} = \mathbf{b}$. We drop the argument whenever \mathbf{b} is clear from the context.

A.2. Preliminaries on Electrical Energy

Throughout the paper, our analyses will rely on a potential function, which in the case of resistor networks corresponds to the electrical energy. In this section we provide a few useful facts.

Lemma A.2 (Characterization of Electrical Energy). *Given a vector of resistances $\mathbf{r} \in \mathbb{R}^m$, we have the following equivalent characterizations for the electrical energy.*

$$\mathcal{E}_r(\mathbf{b}) = \mathbf{b}^\top \left(\mathbf{AD}(\mathbf{r})^{-1} \mathbf{A}^\top \right) \mathbf{b} \quad (\text{A.2})$$

$$= \max_{\phi} 2 \cdot \mathbf{b}^\top \phi - \sum_{i=1}^m \frac{(\mathbf{A}^\top \phi)_i^2}{r_i} \quad (\text{A.3})$$

$$= \left(\min_{\phi: \mathbf{b}^\top \phi = 1} \sum_{i=1}^m \frac{(\mathbf{A}^\top \phi)_i^2}{r_i} \right)^{-1}. \quad (\text{A.4})$$

Furthermore, if \mathbf{x} is the minimizing flow for the expression in (A.1), and ϕ is the maximizing set of potentials for the expression in (A.3), then for all i :

$$x_i = (\mathbf{A}^\top \phi)_i / r_i. \quad (\text{A.5})$$

Since the proof is standard, we defer it to Section B.1.

As a corollary, we can derive a lower bound on the increase in energy after increasing resistances.

Lemma A.3. *Let \mathbf{r}, \mathbf{r}' , and let $\mathbf{x} = \arg \min_{\mathbf{x}: \mathbf{Ax}=\mathbf{b}} \langle \mathbf{r}, \mathbf{x}^2 \rangle$. Then, one has that*

$$\mathcal{E}_{\mathbf{r}'}(\mathbf{b}) \geq \mathcal{E}_r(\mathbf{b}) + \sum_{i=1}^m r_i x_i^2 \left(1 - \frac{r_i}{r'_i} \right).$$

The proof can be found in Section B.3.

We can derive a similar lower bound on the inverse energy, after increasing conductances.

Lemma A.4. *Let $\phi = \arg \min_{\phi: \langle \mathbf{b}, \phi \rangle = 1} \langle \mathbf{c}, (\mathbf{A}^\top \phi)^2 \rangle$. Then one has that*

$$\frac{1}{\mathcal{E}_{1/c'}(\mathbf{b})} \geq \frac{1}{\mathcal{E}_{1/c}(\mathbf{b})} + \frac{1}{\mathcal{E}_{1/c}(\mathbf{b})^2} \cdot \sum_{i=1}^m c_i (\mathbf{A}^\top \phi)_i^2 \left(1 - \frac{c_i}{c'_i} \right).$$

We defer the proof to Section B.2.

A.3. Convergence Proof for ℓ_∞ Minimization

Having put together all these tools, we are ready to analyze the algorithms presented in Section 3. We first prove that ℓ_∞ -MINIMIZATION returns a correct infeasibility certificate, whenever it returns on line 20. This lemma is key to understanding the intuition behind the algorithm.

Lemma A.5. *Whenever ℓ_∞ -MINIMIZATION returns on line 20, $\mathbf{r}/\|\mathbf{r}\|_1$ is a correct approximate infeasibility certificate in the sense that*

$$\mathcal{E}_{\mathbf{r}/\|\mathbf{r}\|_1}(\mathbf{d}) \geq (1 - \varepsilon)^2 M^2.$$

Proof. First notice that by Lemma 2.1, the lower bound on energy is indeed an approximate infeasibility certificate. Now we proceed to prove that throughout the iterations of the algorithm, energy increases at the right rate.

We show that every iteration satisfies the invariant

$$\frac{\mathcal{E}_{\mathbf{r}^{(t+1)}}(\mathbf{d}) - \mathcal{E}_{\mathbf{r}^{(t)}}(\mathbf{d})}{\|\mathbf{r}^{(t+1)} - \mathbf{r}^{(t)}\|_1} \geq M^2. \quad (\text{A.6})$$

This is easy to verify using Lemma A.3, which lower bounds the increase in energy after perturbing resistances. We see that using the perturbation rule defined on line 13 of the algorithm, energy increases as follows

$$\mathcal{E}_{\mathbf{r}^{(t+1)}}(\mathbf{d}) \geq \mathcal{E}_{\mathbf{r}^{(t)}}(\mathbf{d}) + \sum_{i=1}^m r_i^{(t)} (x_i^{(t)})^2 \cdot \left(1 - \frac{1}{\alpha_i^{(t)}}\right).$$

For every coordinate of $\mathbf{r}^{(t)}$ that has changed we see that the ratio between the contribution to above lower bound of that specific coordinate, and the increase in resistance is

$$\frac{r_i^{(t)} (x_i^{(t)})^2 \left(1 - \frac{1}{\alpha_i^{(t)}}\right)}{r_i (\alpha_i^{(t)} - 1)} = \frac{(x_i^{(t)})^2}{\alpha_i^{(t)}} = M^2.$$

Therefore, summing up over all coordinates we obtain the desired inequality. Finally, we notice that initially $\mathcal{E}_{\mathbf{r}^{(0)}}(\mathbf{d}) \geq 0$, and $\|\mathbf{r}^{(0)}\|_1 = 1$. So once $\|\mathbf{r}^{(t)}\|_1 \geq \frac{1}{\varepsilon}$, one has that, using (A.6),

$$\frac{\mathcal{E}_{\mathbf{r}^{(t)}}(\mathbf{d}) - \mathcal{E}_{\mathbf{r}^{(0)}}(\mathbf{d})}{\|\mathbf{r}^{(t)}\|_1 - 1} \geq M^2,$$

and thus

$$\mathcal{E}_{\mathbf{r}^{(t)}}(\mathbf{d}) \geq M^2 (\|\mathbf{r}^{(t)}\|_1 - 1), \text{ and equivalently:}$$

$$\mathcal{E}_{\mathbf{r}^{(t)}/\|\mathbf{r}^{(t)}\|_1}(\mathbf{d}) \geq M^2 \left(1 - \frac{1}{\|\mathbf{r}^{(t)}\|_1}\right) \geq M^2 (1 - \varepsilon),$$

which implies what we needed. \square

Knowing that the algorithm is correct, we can now proceed and prove that it converges fast (convergence rate can be slightly improved by using a more careful schedule for M and ε ; we defer this improvement to Section C).

Lemma A.6. *The algorithm ℓ_∞ -MINIMIZATION returns a solution after $O(m^{1/3} \log(1/\varepsilon)/\varepsilon + \log(m/\varepsilon)/\varepsilon^2)$ iterations.*

Proof. We show that unless the algorithm returns an approximately feasible solution on lines 11 or 15, then there exists a coordinate $i \in [m]$ for which r_i increases very fast.

Suppose the algorithm has run for T iterations without returning an approximately feasible solution. Consider the partial sum of iterates obtained so far $\mathbf{s}^{(t')}$ for some $t' \leq T$. Since the algorithm did not return on line 11, we know that $\|\mathbf{s}^{(t')}\|_\infty/t' \geq (1 + \varepsilon)M$. Therefore there exists a coordinate $i \in [m]$ for which $\mathbf{s}_i^{(t')} \geq (1 + \varepsilon)Mt'$. In other words, letting I be the set of iterates that have contributed to $\mathbf{s}^{(t')}$, one definitely has that

$$\sum_{t \in I} |\mathbf{x}^{(t)}| \geq t' \cdot (1 + \varepsilon)M,$$

and thus

$$\sum_{t \in I} \sqrt{\alpha_i^{(t)}} \geq t' \cdot (1 + \varepsilon),$$

where we used the fact that for each iteration $t \in I$ one has that $\sqrt{\alpha_i^{(t)}} = |x_i^{(t)}|/M$ due to the perturbation rule defined on line 13. This implies that restricting ourselves only to iterations where α_i increased the corresponding resistance r_i , we have that

$$\sum_{t \in I, \alpha_i^{(t)} > 1} \sqrt{\alpha_i^{(t)}} \geq t' \varepsilon, \quad (\text{A.7})$$

By the condition on line 7 we see that for all iterations $t \in I$, one has

$$\sqrt{\alpha_i^{(t)}} \leq m^{1/3}. \quad (\text{A.8})$$

Also since we only consider the iterations $t \in I$ with $\alpha_i^{(t)} > 1$, the rule from line 13 also enforces that for all these iterations

$$\sqrt{\alpha_i^{(t)}} \geq 1 + \varepsilon. \quad (\text{A.9})$$

Equations (A.7), (A.8) and (A.9) suggest that the product $\prod_{t \in I, \alpha_i^{(t)} > 1} \sqrt{\alpha_i^{(t)}}$ increases very fast: intuitively the worst case should occur either when all the factors contribute equally, either all of them are as small as possible (i.e. $1 + \varepsilon$, or as large as possible, i.e. $m^{1/3}$). We formalize this intuition in Lemma B.1, which implies that

$$\prod_{t \in I, \alpha_i^{(t)} > 1} \sqrt{\alpha_i^{(t)}} \geq \min \left\{ \left(m^{1/3}\right)^{\frac{t' \varepsilon}{m^{1/3}}}, (1 + \varepsilon)^{\frac{t' \varepsilon}{1 + \varepsilon}} \right\}.$$

Hence setting

$$t' \geq 10 \left(\frac{m^{1/3} \log(1/\varepsilon)}{\varepsilon} + \frac{\log(m/\varepsilon)}{\varepsilon^2} \right)$$

suffices to lower bound this product by $\sqrt{m/\varepsilon}$. Since each iteration a resistance $r_i^{(t)}$ gets multiplied by the corresponding $\alpha_i^{(t)}$, and all resistances are initially $1/m$, this lower

bound implies that $r_i^{(t)} \geq 1/\varepsilon$. But this means that the algorithm will finish execution after the current iteration, according to the condition on line 5.

Finally, we need to upper bound the number of iterations not in I ; these correspond to those iterations where $\|\mathbf{x}^{(t)}\|_\infty \geq m^{1/3} \cdot M$, so there exists some index i for which $\alpha_i^{(t)} \geq m^{2/3}$. Therefore some resistance gets multiplied by $m^{2/3}$. Since all resistances are initially $1/m$, in the worst case, each such iteration increases one resistance from $1/m$ to $m^{-1/3}$. Therefore this can happen at most $m^{1/3} \log(1/\varepsilon)/\varepsilon$ times, before the sum of resistances becomes at least $1/\varepsilon$, and the algorithm finishes.

Combining these two cases, we obtain our bound. \square

We can prove the convergence bound for ℓ_1 -MINIMIZATION similarly. The main difference is that this time we maintain conductances, and the potential function that enables us to prove convergence is $1/\mathcal{E}_{1/c}$.

A.4. Convergence Proof for ℓ_1 -Minimization

Lemma A.7. *Whenever ℓ_1 -MINIMIZATION returns on line 20, $\mathbf{c}/\|\mathbf{c}\|_1$ is a correct approximate feasibility certificate in the sense that*

$$\frac{1}{\mathcal{E}_{\|\mathbf{c}\|_1/c}} \geq \frac{1/(1+\varepsilon)^2}{M^2}.$$

Proof. By Lemma 2.2, this also yields a solution \mathbf{x} such that $\mathbf{A}\mathbf{x} = \mathbf{b}$ and $\|\mathbf{x}\|_1 \leq \sqrt{\mathcal{E}_{\|\mathbf{x}\|_1/c}} \leq M(1+\varepsilon)$.

In order to prove that at the end of the execution the ℓ_1 norm of this solution is small enough, this time we track as potential function the inverse energy $1/\mathcal{E}_{1/c}$. More precisely, we show that every iteration satisfies the invariant

$$\frac{\frac{1}{\mathcal{E}_{\mathbf{c}^{(t+1)}}(\mathbf{d})} - \frac{1}{\mathcal{E}_{\mathbf{c}^{(t)}}(\mathbf{d})}}{\|\mathbf{c}^{(t+1)} - \mathbf{c}^{(t)}\|_1} \geq \frac{1}{M^2}. \quad (\text{A.10})$$

This is easy to verify using Lemma A.4, which lower bounds the increase in inverse energy after perturbing conductances. We see that using the perturbation rule defined on line 13 of the algorithm, inverse energy increases as follows

$$\begin{aligned} & \frac{1}{\mathcal{E}_{1/c^{(t+1)}}(\mathbf{d})} \geq \frac{1}{\mathcal{E}_{1/c^{(t)}}(\mathbf{d})} \\ & + \frac{1}{\mathcal{E}_{1/c^{(t)}}^2} \cdot \sum_{i=1}^m c_i^{(t)} (\mathbf{A}^\top \boldsymbol{\phi}^{(t)})_i^2 \cdot \left(1 - \frac{1}{\alpha_i^{(t)}}\right). \end{aligned}$$

For every coordinate of $\mathbf{c}^{(t)}$ that has changed we see that the ratio between the contribution to above lower bound of

that specific coordinate, and the increase in conductance is

$$\begin{aligned} & \frac{1}{\mathcal{E}_{1/c^{(t)}}^2} \cdot \frac{c_i^{(t)} (\mathbf{A}^\top \boldsymbol{\phi}^{(t)})_i^2 \left(1 - \frac{1}{\alpha_i^{(t)}}\right)}{c_i (\alpha_i^{(t)} - 1)} = \frac{(\mathbf{A}^\top \boldsymbol{\phi}^{(t)})_i^2}{\mathcal{E}_{1/c^{(t)}}^2} \cdot \frac{1}{\alpha_i^{(t)}} \\ & = \left(\frac{\mathbf{A}^\top \boldsymbol{\phi}^{(t)}}{\mathbf{b}^\top \boldsymbol{\phi}^{(t)}} \right)_i^2 \cdot \frac{1}{\alpha_i^{(t)}} = \frac{1}{M^2}, \end{aligned}$$

where we used the fact that $\mathbf{b}^\top \boldsymbol{\phi}^{(t)} = \mathcal{E}_{1/c^{(t)}}$ (Lemma A.2).

Therefore, summing up over all coordinates we obtain the desired inequality. Since $\mathcal{E}_{1/c^{(0)}} \geq 0$ and $\|\mathbf{c}^{(0)}\|_1 = 1$, we know that once $\|\mathbf{c}^{(t)}\|_1 \geq 1 + \frac{1}{(1+\varepsilon)^2-1} = O(\frac{1}{\varepsilon})$, one has that, using (A.10),

$$\frac{\frac{1}{\mathcal{E}_{1/c^{(t)}}(\mathbf{d})} - \frac{1}{\mathcal{E}_{1/c^{(0)}}(\mathbf{d})}}{\|\mathbf{c}^{(t)}\|_1 - 1} \geq \frac{1}{M^2}$$

and thus,

$$\begin{aligned} \frac{1}{\mathcal{E}_{1/c^{(t)}}(\mathbf{d})} & \geq \frac{\|\mathbf{c}^{(t)}\|_1 - 1}{M^2}, \text{ and equivalently:} \\ \mathcal{E}_{\|\mathbf{c}^{(t)}\|_1/c^{(t)}}(\mathbf{d}) & = \mathcal{E}(1/\mathbf{c}^{(t)})(\mathbf{d}) \cdot \|\mathbf{c}^{(t)}\|_1 \\ & \leq M^2 \cdot \frac{\|\mathbf{c}^{(t)}\|_1}{\|\mathbf{c}^{(t)}\|_1 - 1} \leq M^2(1+\varepsilon)^2, \end{aligned}$$

which is what we needed. \square

Next we prove that the algorithm converges fast. Convergence rate can be slightly improved by using a more careful schedule for M and ε , which we defer to Section C.

Lemma A.8. *The algorithm ℓ_1 -MINIMIZATION returns a solution after $O(m^{1/3} \log(1/\varepsilon)/\varepsilon + \log(m/\varepsilon)/\varepsilon^2)$ iterations.*

Proof. The proof follows the lines of the proof we used for Lemma A.6: unless the algorithm returns an approximate infeasibility certificate on lines 11 or 15, then there exists a coordinate $i \in [m]$ for which c_i increases very fast.

Suppose the algorithm has run for T iterations without returning an approximate infeasibility certificate. Consider the partial sum of iterates obtained so far $\mathbf{s}^{(t')}$ for some $t' \leq T$. Since the algorithm did not return on line 11, we know that $\|\mathbf{s}^{(t')}\|_\infty/t' \geq \frac{1}{(1-\varepsilon)M}$, therefore there exists a coordinate $i \in [m]$ for which $s_i^{(t')} \geq t' \cdot \frac{1}{(1-\varepsilon)M}$. In other words, letting I be the set of iterates that contributed to $\mathbf{s}^{(t')}$, one has that

$$s_i^{(t')} = \sum_{t \in I} \left| \frac{(\mathbf{A}^\top \boldsymbol{\phi}^{(t)})_i}{(\mathbf{b}, \boldsymbol{\phi}^{(t)})} \right| \geq t' \cdot \frac{1}{(1-\varepsilon)M}$$

and thus, since by definition $\alpha_i^{(t)} = \left(\frac{(\mathbf{A}^\top \boldsymbol{\phi}^{(t)})_i}{\mathbf{b}^\top \boldsymbol{\phi}^{(t)}} \right)^2 \cdot M^2$,

$$\sum_{t \in I} \sqrt{\alpha_i^{(t)}} \geq t' \cdot \frac{1}{1 - \varepsilon}.$$

Therefore, restricting ourselves only to iterations where α_i increased the corresponding conductance c_i , we have that

$$\sum_{t \in I, \alpha_i^{(t)} > 1} \sqrt{\alpha_i^{(t)}} \geq t' \cdot \left(\frac{1}{1 - \varepsilon} - 1 \right) \geq t' \cdot \varepsilon. \quad (\text{A.11})$$

By the condition on line 7 we see that for all iterations $t \in I$, one has

$$\sqrt{\alpha_i^{(t)}} \leq m^{1/3}. \quad (\text{A.12})$$

So considering only the iterations $t \in I$ with $\alpha_i^{(t)} > 1$, the rule from line 13 also enforces that for all these iterations

$$\sqrt{\alpha_i^{(t)}} \geq \frac{1}{1 - \varepsilon}. \quad (\text{A.13})$$

Combining Equations (A.11), (A.12), and (A.13), and applying Lemma B.1, exactly the same way we did in the proof of Lemma A.6 implies that

$$\prod_{t \in I, \alpha_i^{(t)} > 1} \sqrt{\alpha_i^{(t)}} \geq \min \left\{ \left(m^{1/3} \right)^{\frac{t' \varepsilon}{m^{1/3}}}, \left(\frac{1}{1 - \varepsilon} \right)^{\frac{t' \varepsilon}{1 - \varepsilon}} \right\}$$

So if

$$t' \geq 10 \left(\frac{m^{1/3} \log(1/\varepsilon)}{\varepsilon} + \frac{\log(m/\varepsilon)}{\varepsilon^2} \right)$$

once again we have that this product is lower bounded by

$\sqrt{m} \cdot \left(1 + \frac{1}{(1+\varepsilon)^2 - 1} \right)$, therefore the corresponding conductance $c_i^{(T)} \geq 1 + \frac{1}{(1+\varepsilon)^2 - 1}$, since its initial value was $1/m$. Since we can only control the total number of iterations T , we can lower bound t' by showing that the number of iterations not in I can not be too large. Just as before, we lower bound the number of iterations where $\left\| \frac{\mathbf{A}^\top \boldsymbol{\phi}^{(t)}}{\mathbf{b}^\top \boldsymbol{\phi}^{(t)}} \right\|_\infty \geq m^{1/3}/M$. Note that whenever this happens, there exists an index i for which $\alpha_i^{(t)} \geq m^{2/3}$. Therefore some conductance gets multiplied by $m^{2/3}$. Again, using an identical argument to the one from the proof of Lemma A.6, we see that this can not happen more than $O(m^{1/3} \log(1/\varepsilon)/\varepsilon)$ times. Combining this with the sufficient number of iterations required by the other case, we obtain our bound. \square

B. Deferred Proofs

B.1. Proof of Lemma A.2

Proof. We can write the formulation from (A.1) as an unconstrained optimization problem using Lagrange multipliers:

$$\begin{aligned} \mathcal{E}_r(\mathbf{b}) &= \min_{\mathbf{A}\mathbf{x}=\mathbf{b}} \langle \mathbf{r}, \mathbf{x}^2 \rangle = \min_{\mathbf{x}} \max_{\boldsymbol{\phi}} \langle \mathbf{r}, \mathbf{x}^2 \rangle + 2\langle \boldsymbol{\phi}, \mathbf{b} - \mathbf{A}\mathbf{x} \rangle \\ &= \max_{\boldsymbol{\phi}} \min_{\mathbf{x}} \langle \mathbf{r}, \mathbf{x}^2 \rangle + 2\langle \boldsymbol{\phi}, \mathbf{b} - \mathbf{A}\mathbf{x} \rangle. \end{aligned}$$

By making the gradient with respect to \mathbf{x} equal to 0, we see that the inner minimization problem is optimized at $2r_i \cdot x_i = 2(\mathbf{A}^\top \boldsymbol{\phi})_i$ for all i , and equivalently $x_i = (\mathbf{A}^\top \boldsymbol{\phi})_i / r_i$. Plugging this back into the maximization objective w.r.t. $\boldsymbol{\phi}$ we obtain

$$\begin{aligned} \mathcal{E}_r(\mathbf{b}) &= \max_{\boldsymbol{\phi}} \left\langle \mathbf{r}, \left(\mathbf{D}(\mathbf{r})^{-1} \mathbf{A}^\top \boldsymbol{\phi} \right)^2 \right\rangle \\ &\quad + 2 \left\langle \boldsymbol{\phi}, \mathbf{b} - \mathbf{A} \mathbf{D}(\mathbf{r})^{-1} \mathbf{A}^\top \boldsymbol{\phi} \right\rangle \\ &= \max_{\boldsymbol{\phi}} 2\langle \boldsymbol{\phi}, \mathbf{b} \rangle - \langle \boldsymbol{\phi}, \mathbf{A} \mathbf{D}(\mathbf{r})^{-1} \mathbf{A}^\top \boldsymbol{\phi} \rangle \\ &= \mathbf{b}^\top \left(\mathbf{A} \mathbf{D}(\mathbf{r})^{-1} \mathbf{A}^\top \right)^+ \mathbf{b}, \end{aligned}$$

where for the last equality we used that by optimality conditions one must have $(\mathbf{A} \mathbf{D}(\mathbf{r})^{-1} \mathbf{A}^\top) \boldsymbol{\phi} = \mathbf{b}$.

Finally, we prove (A.4) by using the fact that for any symmetric matrix \mathbf{L} and vector \mathbf{b} one has that

$$\frac{1}{\max_{\boldsymbol{\phi}} 2\mathbf{b}^\top \boldsymbol{\phi} - \boldsymbol{\phi}^\top \mathbf{L} \boldsymbol{\phi}} = \min_{\boldsymbol{\phi}: \mathbf{b}^\top \boldsymbol{\phi} = 1} \boldsymbol{\phi}^\top \mathbf{L} \boldsymbol{\phi},$$

which can be seen by observing that both expressions are optimized at $\boldsymbol{\phi} = \mathbf{L}^+ \mathbf{b}$, then applying (A.3). \square

B.2. Proof of Lemma A.4

Proof. We use the following basic inequality: for $x, x' > 0$ one has $\frac{1}{x'} \geq \frac{1}{x} + \frac{x-x'}{x^2}$, which follows from $(x-x')^2 \geq 0$. Also, from the definition of energy in (A.1), we obtain an upper bound on the new energy, after perturbing conductances. Let $\mathbf{x} = \arg \min_{\mathbf{A}\mathbf{x}=\mathbf{b}} \langle 1/\mathbf{c}, \mathbf{x}^2 \rangle$, i.e. the electrical flow corresponding to conductances \mathbf{c} . We therefore have:

$$\begin{aligned} \mathcal{E}_{1/\mathbf{c}'}(\mathbf{b}) &\leq \sum_{i=1}^m \frac{1}{\mathbf{c}'_i} x_i^2 = \sum_{i=1}^m \frac{1}{c_i} x_i^2 + \sum_{i=1}^m \frac{1}{c_i} x_i^2 \cdot \left(\frac{c_i}{\mathbf{c}'_i} - 1 \right) \\ &= \mathcal{E}_{1/\mathbf{c}}(\mathbf{b}) + \sum_{i=1}^m \frac{1}{c_i} x_i^2 \left(\frac{c_i}{\mathbf{c}'_i} - 1 \right). \end{aligned}$$

Using the fact that by optimality, $x_i = c_i (\mathbf{A}^\top \boldsymbol{\phi})_i$ (per Lemma A.2), and combining with the previous inequality

we obtain

$$\begin{aligned} \frac{1}{\mathcal{E}_{1/c'}(\mathbf{b})} &\geq \frac{1}{\mathcal{E}_{1/c}(\mathbf{b})} + \frac{1}{\mathcal{E}_{1/c}(\mathbf{b})^2} \cdot (\mathcal{E}_{1/c}(\mathbf{b}) - \mathcal{E}_{1/c'}(\mathbf{b})^2) \\ &\geq \frac{1}{\mathcal{E}_{1/c}(\mathbf{b})} + \frac{1}{\mathcal{E}_{1/c}(\mathbf{b})^2} \cdot \sum_{i=1}^m c_i (\mathbf{A}^\top \boldsymbol{\phi})_i^2 \left(1 - \frac{c_i}{c'_i}\right), \end{aligned}$$

which is what we wanted. \square

B.3. Proof of Lemma A.3

Proof. We use the characterization from Equation A.3 for characterizing electrical energy. Let $\boldsymbol{\phi}$ be the argument that maximizes (A.3) for resistances \mathbf{r} . We certify a lower bound on $\mathcal{E}_{r'}(\mathbf{b})$ using $\boldsymbol{\phi}$ as follows:

$$\begin{aligned} \mathcal{E}_{r'}(\mathbf{b}) &\geq 2 \cdot \mathbf{b}^\top \boldsymbol{\phi} - \sum_{i=1}^m \frac{(\mathbf{A}^\top \boldsymbol{\phi})_i^2}{r'_i} \\ &= 2 \cdot \mathbf{b}^\top \boldsymbol{\phi} - \sum_{i=1}^m \frac{(\mathbf{A}^\top \boldsymbol{\phi})_i^2}{r_i} \\ &\quad + \sum_{i=1}^m \frac{(\mathbf{A}^\top \boldsymbol{\phi})_i^2}{r_i} \cdot \left(1 - \frac{r_i}{r'_i}\right) \\ &= \mathcal{E}_r(\mathbf{b}) + \sum_{i=1}^m \frac{(\mathbf{A}^\top \boldsymbol{\phi})_i^2}{r_i} \cdot \left(1 - \frac{r_i}{r'_i}\right). \end{aligned}$$

Finally substituting the relation between flows and potentials from Lemma A.2, Equation (A.5), we obtain the desired claim. \square

B.4. Lower Bound Lemma

Lemma B.1. *Let a set of nonnegative reals β_1, \dots, β_k such that $1 + \varepsilon \leq \beta_i \leq \rho$ for all i , and $\sum_{i=1}^k \beta_i \geq S$. Then, for any k , one has that*

$$\prod_{i=1}^k \beta_i \geq \min\{\rho^{S/\rho}, (1 + \varepsilon)^{S/(1+\varepsilon)}\}.$$

Proof. Consider a fixed k , and let us attempt to minimize the product of β_i 's subject to the constraints. Equivalently we want to minimize $\sum_{i=1}^k \log(\beta_i)$, which is a concave function. Therefore its minimizer is attained on the boundary of the feasible domain. This means that for some $0 \leq k' \leq k-1$, there are k' elements equal to $1 + \varepsilon$, $k-1-k'$ equal to ρ , and one which is exactly equal to the remaining budget, i.e. $S - k'(1 + \varepsilon) - (k-1-k')\rho$, which yields the product $(1 + \varepsilon)^{k'} \rho^{k-k'-1} (S - k'(1 + \varepsilon) - (k-1-k')\rho)$. This can be relaxed by allowing k and k' to be non-integral. Hence we aim to minimize the product $(1 + \varepsilon)^{k'} \rho^{k-k'}$, subject to $(1 + \varepsilon)k' + \rho(k - k') = S$. Finally, we observe

that we can always obtain a better solution by placing all the available mass on a single one of the factors, i.e. we lower bound either by $(1 + \varepsilon)^{S/(1+\varepsilon)}$, or $\rho^{S/\rho}$, whichever is lowest. \square

C. Using Phases to Improve the Iteration Count

In this section, we show that via minor modifications to our algorithms, we can improve the number of iterations to $O\left(\frac{m^{1/3} \log(1/\varepsilon)}{\varepsilon^{2/3}} + \frac{\log m}{\varepsilon^2}\right)$ thus obtaining the bound promised by Theorem 1.1. This relies on the observation that the entire difficulty of the problem is concentrated on improving the quality of a solution from $(1 - 2\varepsilon)M$ to $(1 - \varepsilon)M$. For conciseness, let us focus on the ℓ_∞ case, and consider the convergence argument described in Sections A.3. Our goal there is to increase the sum of resistances to $1/\varepsilon$, since our argument assumes that the initial energy could be arbitrarily small.

However, if we assume that we warm start the method with a set of resistances \mathbf{r}_0 , $\|\mathbf{r}_0\|_1 = 1$, for which the corresponding energy is already large enough, $\mathcal{E}_{r_0} \geq (1 - 2\varepsilon)^2 M^2$, we only need to iterate until we obtain a set of resistances \mathbf{r} such that $\|\mathbf{r}\|_1 = 3$ (rather than $1/\varepsilon$) in order to certify that the current energy/resistance ratio is as large as desired, i.e. $\mathcal{E}_r / \|\mathbf{r}\|_1 \geq (1 - \varepsilon)^2 M^2$. This in turn improves the number of iterations the algorithm needs before it returns. We expand these ideas in what follows.

Now, suppose we have a set of resistances \mathbf{r}_0 , such that $\|\mathbf{r}_0\|_1 = 1$ and $\mathcal{E}_{r_0} \geq (1 - 2\varepsilon)^2 M^2$. Let us analyze the number of iterations of the method described in Section A.3 that we require before we can return \mathbf{r} such that $\mathcal{E}_r / \|\mathbf{r}\|_1 \geq (1 - \varepsilon)^2 M^2$ or a solution \mathbf{x} such that $\|\mathbf{x}\|_\infty \leq (1 + \varepsilon)M$.

First, we claim that if each update satisfies the invariant from Equation (A.6), then we can stop iteration once $\|\mathbf{r}\|_1 = 3$. Indeed, in this case, one has that

$$\begin{aligned} \frac{\mathcal{E}_r}{\|\mathbf{r}\|_1} &= \frac{\mathcal{E}_{r_0} + (\mathcal{E}_r - \mathcal{E}_{r_0})}{\|\mathbf{r}_0\|_1 + \|\mathbf{r} - \mathbf{r}_0\|_1} \\ &\geq \frac{(1 - 2\varepsilon)^2 M^2 + \|\mathbf{r} - \mathbf{r}_0\|_1 M^2}{1 + \|\mathbf{r} - \mathbf{r}_0\|_1} \\ &\geq M^2 \left(1 - \frac{4\varepsilon}{1 + \|\mathbf{r} - \mathbf{r}_0\|_1}\right) \\ &\geq M^2 (1 - \varepsilon)^2, \end{aligned}$$

whenever $\|\mathbf{r} - \mathbf{r}_0\|_1 \geq 3$.

The remaining analysis is carried over almost identically, except that the threshold set on line 7 is changed to $\rho = \varepsilon m$, and our goal is to get $\prod_{t \in I, \alpha_i^{(t)} > 1} \alpha_i \geq \sqrt{3m}$.

For the iterations that satisfy this threshold, by applying Lemma B.1 we see that it is sufficient to witness a small

number t' of such iterations such that

$$\min \left\{ \rho^{\frac{t'\varepsilon}{\rho}}, (1 + \varepsilon)^{\frac{t'\varepsilon}{1+\varepsilon}} \right\} \geq \sqrt{3m},$$

so $t' = \Theta \left(\frac{\rho}{\varepsilon} \cdot \frac{\log m}{\log \rho} + \frac{\log m}{\varepsilon^2} \right)$ suffices.

For the iterations that do not satisfy the threshold, in the worst case, each of them increases one resistance from $1/m$ to ρ^2/m so this can happen at most $O(m/\rho^2)$ times.

Setting $\rho = (\varepsilon m)^{1/3}$ we get that the total number of iterations is at most $O \left(\frac{m^{1/3}}{\varepsilon^{2/3}} \log(1/\varepsilon) + \frac{\log m}{\varepsilon^2} \right)$.

All of this holds assuming we have a good warm start for resistances. We obtain it by recursively invoking the same method for target value $(1 - 1.75\varepsilon)^2 M^2$, and $.25\varepsilon$ accuracy. In case of failure, this returns a vector \mathbf{x} which certainly satisfies $\|\mathbf{x}\|_\infty \leq M$, so this concludes the entire run on the algorithm; otherwise, it returns a certificate consisting of resistances for which the ratio between the corresponding energy and ℓ_1 norm is at least $(1 - 2\varepsilon)^2 M^2$, so they can be used as a warm start.

Recursion ends once $\varepsilon \geq 1/2$. We note that since the desired accuracy gets increased by a constant factor after each level of recursion, the total number of iterations is dominated by those performed at the top level (i.e. for the lowest ε). Hence our result.

Note that this method can also be implemented slightly more naturally by running Algorithm 1 with a varying schedule for M and ε .

Improving the number of iterations for ℓ_1 minimization is done analogously.

D. From Approximate Decision to Approximate Optimization

Our algorithms are designed to solve an approximate decision problem, given a guess for the value of the optimal solution. While this follows from a standard reduction, for the sake of completeness we prove here that this is sufficient to optimize the problem approximately without paying more than an additional constant overhead in running time.

To be more specific, let us first focus on ℓ_∞ minimization. Theorem 1.1 states that given a guess M and accuracy ε , the algorithm either returns an approximately feasible solution with value $\|\mathbf{x}\|_\infty \leq (1 + \varepsilon)M$, or an infeasibility certificate certifying that $\|\mathbf{x}^*\|_\infty \geq (1 - \varepsilon)M$. Hence this restricts the search interval for the true value either within the interval $[0, (1 + \varepsilon)M]$ or $[(1 - \varepsilon)M, \infty)$.

We initialize our search interval to $[\|\mathbf{x}_0\|_2/\sqrt{m}, \|\mathbf{x}_0\|_\infty]$ where \mathbf{x}_0 is the initial iterate obtained with uniform resistances. Using Lemma 2.1 we easily verify that $\|\mathbf{x}_0\|_2/\sqrt{m}$

is indeed a lower bound on $\|\mathbf{x}^*\|_\infty$, since energy lower bounds the squared optimal value.

Given a search interval $[L, U]$, we let $M = \sqrt{LU}$, $\tilde{\varepsilon} = \min \left\{ \frac{1}{2}, \left(\frac{U}{L} \right)^{1/6} - 1 \right\}$. We invoke Theorem 1.1 for target value M and accuracy $\tilde{\varepsilon}$. Depending on the outcome we update the search interval to $[L, (1 + \tilde{\varepsilon})M]$ or $[(1 - \tilde{\varepsilon})M, U]$.

When $U/L \leq 1 + \varepsilon/4$ we stop the search, call the algorithm for target value $U(1 + \varepsilon/4)$ and accuracy $\frac{\varepsilon/4}{1 + \varepsilon/4}$, then output the approximately feasible iterate returned by the algorithm. The fact that this call indeed returns an approximately feasible iterate follows from the fact that U is certainly feasible, since this is an invariant maintained by our search, and that if the algorithm were to return an infeasibility certificate it must have needed that $U(1 + \varepsilon/4)(1 - \frac{\varepsilon/4}{1 + \varepsilon/4}) < U$, which is false. Thus we know that the returned solution has value at most $U(1 + \varepsilon/4)(1 + \frac{\varepsilon/4}{1 + \varepsilon/4}) \leq L(1 + \varepsilon)$, so it satisfies the desired approximation guarantee.

Finally, we analyze the cost of the search. Note that each iteration of the search reduces $\log U - \log L$ by a constant factor, and it stops whenever it becomes at most $\log(1 + \varepsilon/4) = \Theta(\varepsilon)$. For as long as $U/L > (3/2)^6$, the algorithm is invoked with accuracy $1/2$, and $\log U - \log L$ gets reduced by a constant factor, so this happens at most $O(\log \log m)$ times. Once U/L becomes small enough, i.e. $\log U - \log L < 6 \log(3/2)$, we use accuracy $\exp((\log U - \log L)/6) - 1 = \Theta(\log(U/L))$. Note that from Theorem 1.1 we know that the number of iterations of the algorithm for a single invocation depends on $1/\tilde{\varepsilon}^c$, where c is a fixed constant; due to our schedule for choosing $\tilde{\varepsilon}$, the total cost of this sequence of invocations is dominated by the cost of the final one, where $\tilde{\varepsilon} = \Theta(\varepsilon)$.

So letting $\mathcal{T}(\varepsilon)$ be the number of iterations required by the algorithm from Theorem 1.1 to solve the approximate decision problem to accuracy ε , we have that solving the approximate optimization problem requires $O(\mathcal{T}(1/2) \log \log m + \mathcal{T}(\varepsilon))$ iterations.

The ℓ_1 minimization problem is treated similarly, so we omit its description.