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 $\boldsymbol{A}$ is the vertex-edge incidence matrix of a graph. In this case we will always think of the sought solution $\boldsymbol{x}$ as a flow on the graph's edges. The corresponding dual object is a set of potentials $\boldsymbol{\phi}$ 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 $\boldsymbol{A} \in \mathbb{R}^{n\times m}$ where

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

One can easily verify that given a vector $\boldsymbol{x} \in \mathbb{R}^m$ defined on the arcs of the graph (which we will think of as a **flow**), after applying the operator $\boldsymbol{A}$ we obtain the demand routed by this flow $\boldsymbol{Ax} \in \mathbb{R}^n$, 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 $\boldsymbol{x}$ with minimum congestion which routes the demand $\boldsymbol{b}$, while the ℓ_1 minimization problem from 2.4 corresponds to finding the minimum cost flow routing the demand $\boldsymbol{b}$.

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 $\boldsymbol{A}$'s columns $\boldsymbol{c} \in \mathbb{R}^m$ which we refer to as **conductances**. We equivalently refer to the reciprocals $\boldsymbol{r} \in \mathbb{R}^m$, with $\boldsymbol{r}_i = 1/\boldsymbol{c}_i$, 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 $\boldsymbol{x} \in \mathbb{R}^m$, along with a vector of resistances $\boldsymbol{r} \in \mathbb{R}^m$, we let the energy of $\boldsymbol{x}$ be

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

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

$$\mathcal{E}_{\boldsymbol{r}}(\boldsymbol{b}) = \min_{\boldsymbol{x}:\boldsymbol{Ax}=\boldsymbol{b}} \mathcal{E}_{\boldsymbol{r}}(\boldsymbol{x})\,, \tag{A.1}$$

in other words this is the minimum energy over all flows satisfying $\boldsymbol{Ax} = \boldsymbol{b}$. We drop the argument whenever $\boldsymbol{b}$ 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 $\boldsymbol{r} \in \mathbb{R}^m$, we have the following equivalent characterizations for the electrical energy.*

$$\mathcal{E}_{\boldsymbol{r}}(\boldsymbol{b}) = \boldsymbol{b}^\top \left(\boldsymbol{A}\mathbf{D}(\boldsymbol{r})^{-1}\boldsymbol{A}^\top\right)^+ \boldsymbol{b} \tag{A.2}$$

$$= \max_{\boldsymbol{\phi}} 2\cdot \boldsymbol{b}^\top\boldsymbol{\phi} - \sum_{i=1}^m \frac{\left(\boldsymbol{A}^\top\boldsymbol{\phi}\right)_i^2}{r_i} \tag{A.3}$$

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

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

$$x_i = (\boldsymbol{A}^\top\boldsymbol{\phi})_i / r_i \,. \tag{A.5}$$

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

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

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

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

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,* $\boldsymbol{r}/\|\boldsymbol{r}\|_1$ *is a correct approximate infeasibility certificate in the sense that*

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

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

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

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

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

For every coordinate of $\boldsymbol{r}^{(t)}$ 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

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

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

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

and thus

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

which implies what we needed. □

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 $\boldsymbol{s}^{(t')}$ for some $t' \leq T$. Since the algorithm did not return on line 11, we know that $\|\boldsymbol{s}^{(t')}\|_\infty / t' \geq (1+\varepsilon)M$. Therefore there exists a coordinate $i \in [m]$ for which $\boldsymbol{s}_i^{(t')} \geq (1+\varepsilon)Mt'$. In other words, letting I be the set of iterates that have contributed to $\boldsymbol{s}^{(t')}$, one definitely has that

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

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

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

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

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

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

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

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 $\|\boldsymbol{x}^{(t)}\|_\infty \geq m^{1/3} \cdot M$, 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. □

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 $1/\mathcal{E}_{1/\boldsymbol{c}}$.

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*

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

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

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 $1/\mathcal{E}_{1/\boldsymbol{c}}$. More precisely, we show that every iteration satisfies the invariant

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

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

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

For every coordinate of $\boldsymbol{c}^{(t)}$ 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

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

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

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

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

and thus,

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

$$\mathcal{E}_{\|\boldsymbol{c}^{(t)}\|_1/\boldsymbol{c}^{(t)}}(\boldsymbol{d}) = \mathcal{E}(1/\boldsymbol{c}^{(t)})(\boldsymbol{d}) \cdot \|\boldsymbol{c}^{(t)}\|_1 \leq M^2 \cdot \frac{\|\boldsymbol{c}^{(t)}\|_1}{\|\boldsymbol{c}^{(t)}\|_1 - 1} \leq M^2(1+\varepsilon)^2\,,$$

which is what we needed. □

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 $\boldsymbol{s}^{(t')}$ for some $t' \leq T$. Since the algorithm did not return on line 11, we know that $\|\boldsymbol{s}^{(t')}\|_\infty / t' \geq \frac{1}{(1-\varepsilon)M}$, 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 $\boldsymbol{s}^{(t')}$, one has that

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

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

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

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

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

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 $\left\|\frac{A^\top \phi^{(t)}}{b^\top \phi^{(t)}}\right\|_\infty \geq m^{1/3}/M$. 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. □

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:

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

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

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

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

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

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

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

□

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 $x = \arg\min_{Ax=b}\langle 1/c, x^2\rangle$, i.e. the electrical flow corresponding to conductances c. We therefore have:

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

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

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

which is what we wanted. □

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 $\boldsymbol{r}$. We certify a lower bound on $\mathcal{E}_{\boldsymbol{r}'}(\boldsymbol{b})$ using $\boldsymbol{\phi}$ as follows:

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

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

B.4. Lower Bound Lemma

Lemma B.1. *Let a set of nonnegative reals $\beta_1, \dots, \beta_k$ 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. □

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 $\boldsymbol{r}_0$, $\|\boldsymbol{r}_0\|_1 = 1$, for which the corresponding energy is already large enough, $\mathcal{E}_{\boldsymbol{r}_0} \geq (1-2\varepsilon)^2 M^2$, we only need to iterate until we obtain a set of resistances $\boldsymbol{r}$ such that $\|\boldsymbol{r}\|_1 = 3$ (rather than $1/\varepsilon$) in order to certify that the current energy/resistance ratio is as large as desired, i.e. $\mathcal{E}_{\boldsymbol{r}}/\|\boldsymbol{r}\|_1 \geq (1-\varepsilon)^2 M^2$. 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 $\boldsymbol{r}_0$, such that $\|\boldsymbol{r}_0\|_1 = 1$ and $\mathcal{E}_{\boldsymbol{r}_0} \geq (1-2\varepsilon)^2 M^2$. Let us analyze the number of iterations of the method described in Section A.3 that we require before we can return $\boldsymbol{r}$ such that $\mathcal{E}_{\boldsymbol{r}}/\|\boldsymbol{r}\|_1 \geq (1-\varepsilon)^2 M^2$ or a solution $\boldsymbol{x}$ such that $\|\boldsymbol{x}\|_\infty \leq (1+\varepsilon)M$.

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

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

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

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 $(1-1.75\varepsilon)^2M^2$, and $.25\varepsilon$ accuracy. In case of failure, this returns a vector $\boldsymbol{x}$ which certainly satisfies $\|\boldsymbol{x}\|_\infty \leq M$, 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 $(1-2\varepsilon)^2M^2$, 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 $\|\boldsymbol{x}\|_\infty \leq (1+\varepsilon)M$, or an infeasibility certificate certifying that $\|\boldsymbol{x}^*\|_\infty \geq (1-\varepsilon)M$. 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 $[\|\boldsymbol{x}_0\|_2/\sqrt{m}, \|\boldsymbol{x}_0\|_\infty]$ where $\boldsymbol{x}_0$ is the initial iterate obtained with uniform resistances. Using Lemma 2.1 we easily verify that $\|\boldsymbol{x}_0\|_2/\sqrt{m}$ is indeed a lower bound on $\|\boldsymbol{x}^*\|_\infty$, 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 $\log U - \log L$ be 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 $O\left(\mathcal{T}(1/2)\log\log m + \mathcal{T}(\varepsilon)\right)$ iterations.

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