

Matrix Balancing in L_p Norms: A New Analysis of Osborne’s Iteration

Rafail Ostrovsky
UCLA
rafail@cs.ucla.edu*

Yuval Rabani
The Hebrew University of Jerusalem
yrabani@cs.huji.ac.il[†]

Arman Yousefi
UCLA
armany@cs.ucla.edu*

June 28, 2016

Abstract

We study an iterative matrix conditioning algorithm due to Osborne (1960). The goal of the algorithm is to convert a square matrix into a *balanced* matrix where every row and corresponding column have the same norm. The original algorithm was proposed for balancing rows and columns in the L_2 norm, and it works by iterating over balancing a row-column pair in fixed round-robin order. Variants of the algorithm for other norms have been heavily studied and are implemented as standard preconditioners in many numerical linear algebra packages. Recently, Schulman and Sinclair (2015), in a first result of its kind for any norm, analyzed the rate of convergence of a variant of Osborne’s algorithm that uses the L_∞ norm and a different order of choosing row-column pairs. In this paper we study matrix balancing in the L_1 norm and other L_p norms. We show the following results for any matrix $A = (a_{ij})_{i,j=1}^n$, resolving in particular a main open problem mentioned by Schulman and Sinclair.

1. We analyze the iteration for the L_1 norm under a greedy order of balancing. We show that it converges to an ϵ -balanced matrix in $K = O(\min\{\epsilon^{-2} \log w, \epsilon^{-1} n^{3/2} \log(w/\epsilon)\})$ iterations that cost a total of $O(m + Kn \log n)$ arithmetic operations over $O(n \log w)$ -bit numbers. Here m is the number of non-zero entries of A , and $w = \sum_{i,j} |a_{ij}| / a_{\min}$ with $a_{\min} = \min\{|a_{ij}| : a_{ij} \neq 0\}$.
2. We show that the original round-robin implementation converges to an ϵ -balanced matrix in $O(\epsilon^{-2} n^2 \log w)$ iterations totalling $O(\epsilon^{-2} mn \log w)$ arithmetic operations over $O(n \log w)$ -bit numbers.
3. We show that a random implementation of the iteration converges to an ϵ -balanced matrix in $O(\epsilon^{-2} \log w)$ iterations using $O(m + \epsilon^{-2} n \log w)$ arithmetic operations over $O(\log(wn/\epsilon))$ -bit numbers.
4. We demonstrate a lower bound of $\Omega(1/\sqrt{\epsilon})$ on the convergence rate of any implementation of the iteration.
5. We observe, through a known trivial reduction, that our results for L_1 balancing apply to any L_p norm for all finite p , at the cost of increasing the number of iterations by only a factor of p .

We note that our techniques are very different from those used by Schulman and Sinclair.

*Research supported in part by NSF grants 1065276, 1118126 and 1136174, US-Israel BSF grants, OKAWA Foundation Research Award, IBM Faculty Research Award, Xerox Faculty Research Award, B. John Garrick Foundation Award, Teradata Research Award, and Lockheed-Martin Corporation Research Award. This material is also based upon work supported in part by DARPA Safeware program. The views expressed are those of the authors and do not reflect the official policy or position of the Department of Defense or the U.S. Government.

[†]Research supported in part by ISF grant 956-15, by BSF grant 2012333, and by I-CORE Algo.

1 Introduction

Let $A = (a_{ij})_{n \times n}$ be a square matrix with real entries, and let $\|\cdot\|$ be a given norm. For an index $i \in [n]$, let $\|a_{i,\cdot}\|$ and $\|a_{\cdot,i}\|$, respectively, denote the norms of the i th row and the i th column of A , respectively. The matrix A is *balanced* in $\|\cdot\|$ iff $\|a_{\cdot,i}\| = \|a_{i,\cdot}\|$ for all i . An invertible diagonal matrix $D = \text{diag}(d_1, \dots, d_n)$ is said to *balance* a matrix A iff DAD^{-1} is balanced. A matrix A is *balanceable* in $\|\cdot\|$ iff there exists a diagonal matrix D that balances it.

Osborne [8] studied the above problem in the L_2 norm and considered its application in preconditioning a given matrix in order to increase the accuracy of the computation of its eigenvalues. The motivation is that standard linear algebra algorithms that are used to compute eigenvalues are numerically unstable for unbalanced matrices; diagonal balancing addresses this issue by obtaining a balanced matrix that has the same eigenvalues as the original matrix, as DAD^{-1} and A have the same eigenvalues. Osborne suggested an iterative algorithm for finding a diagonal matrix D that balances a matrix A , and also proved that his algorithm converges in the limit. He also observed that if a diagonal matrix $D = \text{diag}(d_1, \dots, d_n)$ balances a matrix A , then the diagonal vector $\mathbf{d} = (d_1, \dots, d_n)$ minimizes the Frobenius norm of the matrix DAD^{-1} . Osborne’s classic algorithm is an iteration that at each step balances a row and its corresponding column by scaling them appropriately. More specifically the algorithm balances row-column pairs in a fixed cyclic order. In order to balance row and column i , the algorithm scales the i th row by $\sqrt{\|a_{\cdot,i}\|/\|a_{i,\cdot}\|}$ and the i th column by $\sqrt{\|a_{i,\cdot}\|/\|a_{\cdot,i}\|}$. Osborne’s algorithm converges to a unique balanced matrix, but there have been no upper bounds on the converges rate of Osborne’s algorithm for the L_2 norm prior to our work.

Parlett and Reinsch [9] generalized Osborne’s algorithm to other norms without proving convergence. The L_1 version of the algorithm has been studied extensively. The convergence in the limit of the L_1 version was proved by Grad [4], uniqueness of the balanced matrix by Hartfiel [5], and a characterization of balanceable matrices was given by Eaves et al. [3]. Again, there have been no upper bounds on the running time of the L_1 version of the iteration. The first polynomial time algorithm for balancing a matrix in the L_1 norm was given by Kalantari, Khachiyan, and Shokoufandeh [6]. Their approach is different from the iterative algorithm of Osborne-Parlett-Reinsch. They reduce the balancing problem to a convex optimization problem and then solve that problem approximately using the ellipsoid algorithm. Their algorithm runs in $O(n^4 \log(n \log w/\epsilon))$ arithmetic operations where $w = \sum_{i,j} |a_{i,j}|/a_{\min}$ for $a_{\min} = \min\{|a_{ij}| : a_{ij} \neq 0\}$ and ϵ is the relative imbalance of the output matrix (see Definition 1).

For matrix balancing in the L_∞ norm, Schneider and Schneider [11] gave an $O(n^4)$ -time non-iterative algorithm. This running time was improved to $O(mn + n^2 \log n)$ by Young, Tarjan, and Orlin [14]. Despite the existence of polynomial time algorithms for balancing in the L_1 and L_∞ norms, and the lack of any theoretical bounds on the running time of the Osborne-Parlett-Reinsch (OPR) iterative algorithm, the latter is favored in practice, and the Parlett and Reinsch variant [9] is implemented as a standard in almost all linear algebra packages (see Chen [2, Section 3.1], also the book [10, Chapter 11] and the code in [1]). One reason is that iterative methods usually perform well in practice and run for far fewer iterations than are needed in the worst case. Another advantage of iterative algorithms is that they are simple, they provide steady partial progress, and they can always generate a matrix that is sufficiently balanced for the subsequent linear algebra computation.

Motivated by the impact of the OPR algorithm and the lack of any theoretical bounds on its running time, Schulman and Sinclair [12] recently showed the first bound on the convergence rate of a modified version of this algorithm in the L_∞ norm. They prove that their modified algorithm converges in $O(n^3 \log(\rho n/\epsilon))$ balancing steps where ρ measures the initial imbalance of A and ϵ is the target imbalance of the output matrix. Their algorithm differs from the original algorithm only in the order of choosing row-column pairs to balance (we will use the term *variant* to indicate a deviation from the original round-robin order). Schulman and Sinclair do not prove any bounds on the running time of the algorithm for other L_p norms; this was explicitly mentioned as an open problem. Notice that when changing the norm,

not only the target balancing condition changes but also the iteration itself, so we cannot deduce an upper bound on the rate of convergence in the L_p norm from the rate of convergence in the L_∞ norm.

In this paper we resolve the open question of [12], and upper bound the convergence rate of the OPR iteration in any L_p norm.¹ Specifically, we show the following bounds for the L_1 norm. They imply the same bounds with an extra factor of p for the L_p norm, by using them on the matrix with entries raised to the power of p . (Below, the $\tilde{O}(\cdot)$ notation hides factors that are logarithmic in various parameters of the problem. Exact bounds await the statements of the theorems in the following sections.) We show that the original algorithm (with no modification) converges to an ϵ -balanced matrix in $\tilde{O}(n^2/\epsilon^2)$ balancing steps, using $\tilde{O}(mn/\epsilon^2)$ arithmetic operations. We also show that a greedy variant converges in $\tilde{O}(1/\epsilon^2)$ balancing steps, using $O(m) + \tilde{O}(n/\epsilon^2)$ arithmetic operations; or alternatively in $\tilde{O}(n^{3/2}/\epsilon)$ iterations, using $\tilde{O}(n^{5/2}/\epsilon)$ arithmetic operations. Thus, the number of arithmetic operations needed by our greedy variant is nearly linear in m or nearly linear in $1/\epsilon$. The near linear dependence on m is significantly better than the Kalantari-Khachiyani-Shokoufandeh algorithm that uses $O(n^4 \log(n \log w/\epsilon))$ arithmetic operations (and also the Schulman and Sinclair version with a stricter, yet L_∞ , guarantee). For an accurate comparison we should note that we may need to maintain $\tilde{O}(n)$ bits of precision, so the running time is actually $O(m + n^2 \log n \log w/\epsilon^2)$ (the Kalantari et al. algorithm maintains $O(\log(wn/\epsilon))$ -bit numbers). We improve this with yet another, randomized, variant that has similar convergence rate (nearly linear in m), but needs only $O(\log(wn/\epsilon))$ bits of precision. Finally, we show that the dependence on ϵ given by our analyses is within the right ballpark—we demonstrate a lower bound of $\Omega(1/\sqrt{\epsilon})$ on the convergence rate of any variant of the algorithm to an ϵ -balanced matrix. Notice the contrast with the Schulman-Sinclair upper bound for balancing in the L_∞ norm that has $O(\log(1/\epsilon))$ dependence on ϵ (this lower bound is for the Kalantari et al. notion of balancing so it naturally applies also to strict balancing).

Osborne observed that a diagonal matrix $D = \text{diag}(d_1, \dots, d_n)$ that balances a matrix A in the L_2 norm also minimizes the Frobenius norm of the matrix DAD^{-1} . Thus, the balancing problem can be reduced to minimizing a convex function. Kalantari et al. [6] gave a convex program for balancing in the L_1 norm. Our analysis is based on their convex program. We relate the OPR balancing step to the coordinate descent method in convex programming. We show that each step reduces the value of the objective function. Our various bounds are derived through analyzing the progress made in each step. In particular, one of the main tools in our analysis is an upper bound on the distance to optimality (measured by the convex objective function) in terms of the L_1 norm of the gradient, which we prove using network flow arguments.

For lack of space, many proofs are missing inline. They appear in Section 7.

2 Preliminaries

In this section we introduce notation and definitions, we discuss some previously known facts and results, and we prove a couple of useful lemmas.

The problem. Let $A = (a_{ij})_{n \times n}$ be a square real matrix, and let $\|\cdot\|$ be a norm on \mathbb{R}^n . For an index $i \in [n]$, let $\|a_{i,\cdot}\|$ and $\|a_{\cdot,i}\|$, respectively, denote the norms of the i th row and the i th column of A , respectively. A matrix A is *balanced* in $\|\cdot\|$ iff $\|a_{\cdot,i}\| = \|a_{i,\cdot}\|$ for all i . An invertible diagonal matrix $D = \text{diag}(d_1, \dots, d_n)$ is said to *balance* a matrix A iff DAD^{-1} is balanced. A matrix A is *balanceable* in $\|\cdot\|$ iff there exists a diagonal matrix D that balances it.

For balancing a matrix A in the L_p norm only the absolute values of the entries of A matter, so we may assume without loss of generality that A is non-negative. Furthermore, balancing a matrix does not

¹It should be noted that the definition of target imbalance ϵ in [12] is stricter than the definition used by [6]. We use the definition in [6]. This is justified by the fact that the numerical stability of eigenvalue calculations depends on the Frobenius norm of the balanced matrix, see [9].

change its diagonal entries, so if a diagonal matrix D balances A with its diagonal entries replaced by zeroes, then D balances A too. Thus, for the rest of the paper, we assume without loss of generality that the given $n \times n$ matrix $A = (a_{ij})$ is non-negative and its diagonal entries are all 0.

A diagonal matrix $D = \text{diag}(d_1, \dots, d_n)$ balances $A = (a_{ij})$ in the L_p norm if and only if $D^p = \text{diag}(d_1^p, \dots, d_n^p)$ balances the matrix $A' = (a_{ij}^p)$ in the L_1 norm. Thus, the problem of balancing matrices in the L_p norm (for any finite p) reduces to the problem of balancing matrices in the L_1 norm; for the rest of the paper we focus on balancing matrices in the L_1 norm.

For an $n \times n$ matrix A , we use $G_A = (V, E, w)$ to denote the weighted directed graph whose adjacency matrix is A . More formally, G_A is defined as follows. Put $V = \{1, \dots, n\}$, put $E = \{(i, j) : a_{ij} \neq 0\}$, and put $w(i, j) = a_{ij}$ for every $(i, j) \in E$. We use an index $i \in [n]$ to refer to both the i th row or column of A , and to the node i of the digraph G_A . Thus, the non-zero entries of the i th column (the i th row, respectively) correspond to the arcs into (out of, respectively) node i . In the L_1 norm it is useful to think of the weight of an arc as a flow being carried by that arc. Thus, $\|a_{\cdot, i}\|_1$ is the total flow into vertex i and $\|a_{i, \cdot}\|_1$ is the total flow out of it. Note that if a matrix A is not balanced then for some nodes i , $\|a_{\cdot, i}\|_1 \neq \|a_{i, \cdot}\|_1$, and thus the flow on the arcs does not constitute a valid circulation because flow conservation is not maintained. Thus, the goal of balancing in the L_1 norm can be stated as applying diagonal scaling to find a flow function on the arcs of the graph G_A that forms a valid circulation. We use both views of the graph (with arc weights or flow), and also the matrix terminology, throughout this paper, as convenient.

Without loss of generality we may assume that the undirected graph underlying G_A is connected. Otherwise, after permuting $V = \{1, \dots, n\}$, the given matrix A can be replaced by $\text{diag}(A_1, \dots, A_r)$ where each of A_1, \dots, A_r is a square matrix whose corresponding directed graph is connected. Thus, balancing A is equivalent to balancing each of A_1, \dots, A_r .

The goal of the iterative algorithm is to balance approximately a matrix A , up to an error term ϵ . We define the error here.

Definition 1 (approximate balancing). *Let $\epsilon > 0$.*

1. *A matrix A is ϵ -balanced iff $\frac{\sqrt{\sum_{i=1}^n (\|a_{\cdot, i}\|_1 - \|a_{i, \cdot}\|_1)^2}}{\sum_{i,j} a_{i,j}} \leq \epsilon$.*

2. *A diagonal matrix D with positive diagonal entries is said to ϵ -balance A iff DAD^{-1} is ϵ -balanced.*

The algorithms. Kalantari et al. [6] introduced the above definition of ϵ -balancing, and showed that their algorithm for ϵ -balancing a matrix in the L_1 norm uses $O(n^4 \ln((n/\epsilon) \ln w))$ arithmetic operations. In their recent work, Schulman and Sinclair [12] use, in the context of balancing in the L_∞ norm, a stronger notion of strict balancing (that requires even very low weight row-column pairs to be nearly balanced). Their iterative algorithm strictly ϵ -balances a matrix in the L_∞ norm in $O(n^3 \log(n\rho/\epsilon))$ iterations where ρ measures the initial imbalance of the matrix. In this paper, we prove upper bounds on the convergence rate of the Osborne-Parlett-Reinsch (OPR) balancing.

The OPR iterative algorithm balances indices in a fixed round-robin order. Schulman and Sinclair considered a variant that uses a different rule to choose the next index to balance. We consider in this paper several alternative implementations of OPR balancing (including the original round-robin implementation) that differ only in the rule by which an index to balance is chosen at each step. For all rules that we consider, the iteration generates a sequence $A = A^{(1)}, A^{(2)}, \dots, A^{(t)}, \dots$ of $n \times n$ matrices that converges to a unique balanced matrix A^* (see Grad [4] and Hartfiel [5]). The matrix $A^{(t+1)}$ is obtained by balancing an index of $A^{(t)}$. If the i th index of $A^{(t)}$ is chosen, we get that $A^{(t+1)} = D^{(t)} A^{(t)} D^{(t)-1}$ where $D^{(t)}$ is a diagonal matrix with $d_{ii}^{(t)} = \sqrt{\|a_{\cdot, i}^{(t)}\|_1 / \|a_{i, \cdot}^{(t)}\|_1}$ and $d_{jj}^{(t)} = 1$ for $j \neq i$. Note that $a_{i, \cdot}^{(t)}$ ($a_{\cdot, i}^{(t)}$, respectively) denotes the i th row (i th column, respectively) of $A^{(t)}$. Also, putting $\bar{D}^{(1)} = I_{n \times n}$ and $\bar{D}^{(t)} = D^{(t-1)} \dots D^{(1)}$ for $t > 1$, we get that $A^{(t)} = \bar{D}^{(t)} A (\bar{D}^{(t)})^{-1}$.

The following lemma shows that each balancing step reduces the sum of entries of the matrix.

Lemma 1. *Balancing the i th index of a non-negative matrix $B = (b_{ij})_{n \times n}$ (with $b_{ii} = 0$) decreases the total sum of the entries of B by $(\sqrt{\|b_{\cdot,i}\|_1} - \sqrt{\|b_{i,\cdot}\|_1})^2$.*

Proof. Before balancing, the total sum of entries in the i th row and in the i th column is $\|b_{i,\cdot}\|_1 + \|b_{\cdot,i}\|_1$. Balancing scales the entries of the i th column by $\sqrt{\|b_{i,\cdot}\|_1 / \|b_{\cdot,i}\|_1}$ and entries of the i th row by $\sqrt{\|b_{\cdot,i}\|_1 / \|b_{i,\cdot}\|_1}$. Thus, after balancing the sum of entries in the i th column, which equals the sum of entries in the i th row, is equal to $\sqrt{\|b_{i,\cdot}\|_1 \cdot \|b_{\cdot,i}\|_1}$. The entries that are not in the balanced row and column are not changed. Therefore, keeping in mind that $b_{ii} = 0$, balancing decreases $\sum_{i,j} b_{ij}$ by $\|b_{\cdot,i}\|_1 + \|b_{i,\cdot}\|_1 - 2\sqrt{\|b_{i,\cdot}\|_1 \cdot \|b_{\cdot,i}\|_1} = (\sqrt{\|b_{\cdot,i}\|_1} - \sqrt{\|b_{i,\cdot}\|_1})^2$. \square

A reduction to convex optimization. Kalantari et al. [6], as part of their algorithm, reduce matrix balancing to a convex optimization problem. We overview their reduction here. Our starting point is Osborne's observation that if a diagonal matrix $D = \text{diag}(d_1, \dots, d_n)$ balances a matrix A in the L_2 norm, then the diagonal vector $\mathbf{d} = (d_1, \dots, d_n)$ minimizes the Frobenius norm of the matrix DAD^{-1} . The analogous claim for the L_1 norm is that if a diagonal matrix $D = \text{diag}(d_1, \dots, d_n)$ balances a matrix A in the L_1 norm, then the diagonal vector $\mathbf{d} = (d_1, \dots, d_n)$ minimizes the function $F(\mathbf{d}) = \sum_{i,j} a_{ij} \frac{d_i}{d_j}$. On the other hand, Eaves et al. [3] observed that a matrix A can be balanced if and only if the digraph G_A is strongly connected. The following theorem [6, Theorem 1] summarizes the above discussion.

Theorem 1 (Kalantari et al.). *Let $A = (a_{ij})_{n \times n}$ be a real non-negative matrix, $a_{ii} = 0$, for all $i = 1, \dots, n$, such that the undirected graph underlying G_A is connected. Then, the following statements are equivalent.*

- (i) *A is balanceable (i.e., there exists a diagonal matrix D such that DAD^{-1} is balanced).*
- (ii) *G_A is strongly connected.*
- (iii) *Let $F(\mathbf{d}) = \sum_{(i,j) \in E} a_{ij} \frac{d_i}{d_j}$. There is a point $\mathbf{d}^* \in \Omega = \{d \in \mathbb{R}^n : d_i > 0, i = 1, \dots, n\}$ such that $F(\mathbf{d}^*) = \inf\{F(\mathbf{d}) : \mathbf{d} \in \Omega\}$.*

We refer the reader to [6, Theorem 1] for a proof. We have the following corollary.

Corollary 1. *\mathbf{d}^* minimizes F over Ω if and only if $D^* = \text{diag}(d_1^*, \dots, d_n^*)$ balances A .*

Proof. As F attains its infimum at $\mathbf{d}^* \in \Omega$, its gradient ∇F satisfies $\nabla F(\mathbf{d}^*) = 0$. Also, $\frac{\partial F(\mathbf{d}^*)}{\partial d_i} = 0$ if and only if $\sum_{j=1}^n a_{ij} \cdot (d_i^*/d_j^*) = \sum_{j=1}^n a_{ji} \cdot (d_j^*/d_i^*)$ for all $i \in [n]$. In other words, $\nabla F(\mathbf{d}^*) = 0$ if and only if the matrix D^*AD^{*-1} is balanced where $D^* = \text{diag}(d_1^*, \dots, d_n^*)$. Thus, \mathbf{d}^* minimizes F over Ω if and only if $D^* = \text{diag}(d_1^*, \dots, d_n^*)$ balances A . \square

It can also be shown that under the assumption of Theorem 1, the balancing matrix D^* is unique up to a scalar factor (see Osborne [8] and Eaves et al. [3]). Therefore, the problem of balancing matrix A can be reduced to optimizing the function F . Since we are optimizing over the set Ω of strictly positive vectors, we can apply a change of variables $\mathbf{d} = (e^{x_1}, \dots, e^{x_n}) \in \mathbb{R}^n$ to obtain a convex objective function:

$$f(\mathbf{x}) = f_A(\mathbf{x}) = \sum_{i,j=1}^n a_{ij} e^{x_i - x_j}. \quad (1)$$

Kalantari et al. [6] use the convex function f because it can be minimized using the ellipsoid algorithm. We do not need the convexity of f , and use f instead of F only because it is more convenient to work

with, and it adds some intuition. Notice that the partial derivative of f with respect to x_i is

$$\frac{\partial f(\mathbf{x})}{\partial x_i} = \sum_{j=1}^n a_{ij} \cdot e^{x_i - x_j} - \sum_{j=1}^n a_{ji} \cdot e^{x_j - x_i}, \quad (2)$$

which is precisely the difference between the L_1 norms of the i th row and the i th column of the matrix DAD^{-1} , where $D = \text{diag}(e^{x_1}, \dots, e^{x_n})$. Also, by definition, the diagonal matrix $\text{diag}(e^{x_1}, \dots, e^{x_n})$ ϵ -balances A iff

$$\frac{\|\nabla f(\mathbf{x})\|_2}{f(\mathbf{x})} = \frac{\sqrt{\sum_{i=1}^n \left(\sum_{j=1}^n a_{ij} e^{x_i - x_j} - \sum_{j=1}^n a_{ji} e^{x_j - x_i} \right)^2}}{\sum_{i,j=1}^n a_{ij} e^{x_i - x_j}} \leq \epsilon. \quad (3)$$

We now state and prove a key lemma that our analysis uses. The lemma uses combinatorial flow and circulation arguments to measure progress by bounding $f(\mathbf{x}) - f(\mathbf{x}^*)$ in terms of $\|\nabla f(\mathbf{x})\|_1$ which is a global measure of imbalances of all vertices.

Lemma 2. *Let f be the function defined in Equation (1), and let \mathbf{x}^* be a global minimum of f . Then, for all $\mathbf{x} \in \mathbb{R}^n$, $f(\mathbf{x}) - f(\mathbf{x}^*) \leq \frac{n}{2} \cdot \|\nabla f(\mathbf{x})\|_1$.*

Proof. Recall that $f(\mathbf{x}) = f_A(\mathbf{x})$ is the sum of entries of a matrix $B = (b_{ij})$ defined by $b_{ij} = a_{ij} \cdot e^{x_i - x_j}$. Notice that $f(\mathbf{x}) = f_B(\vec{\mathbf{0}})$, and $f(\mathbf{x}^*) = f_B(\mathbf{x}^{**})$, where $\mathbf{x}^{**} = \mathbf{x}^* - \mathbf{x}$. Alternatively, $f(\mathbf{x})$ is the sum of flows (or weights) of the arcs of G_B , and $f(\mathbf{x}^*)$ is the sum of flows of the arcs of a graph G^* (an arc ij of G^* carries a flow of $a_{ij} \cdot e^{x_i^* - x_j^*}$). Notice that G_B and G^* have the same set of arcs, but with different weights. By Equation (2), $\|\nabla f_A(\mathbf{x})\|_1 = \sum_{i=1}^n \left| \|b_{\cdot,i}\|_1 - \|b_{i,\cdot}\|_1 \right|$, i.e., it is the sum over all the nodes of G_B of the difference between the flow into the node and flow out of it. Also notice that G_B is unbalanced (else the statement of the lemma is trivial), however G^* is balanced. Therefore, the arc flows in G^* , but not those in G_B , form a valid circulation.

Our proof now proceeds in two main steps. In the first step we show a way of reducing the flow on some arcs of G_B , such that the revised flows make every node balanced (and thus form a valid circulation). We also make sure that the total flow reduction is at most $\frac{n}{2} \cdot \|\nabla f_A(\mathbf{x})\|_1$. In the second step we show that sum of revised flows of all the arcs is a lower bound on $f(\mathbf{x}^*)$. These two steps together prove the lemma.

We start with the first step. The nodes of G_B are not balanced. Let S and T be a partition of the unbalanced nodes of G_B , with $S = \{i \in [n] : \|b_{\cdot,i}\|_1 > \|b_{i,\cdot}\|_1\}$ and $T = \{i \in [n] : \|b_{\cdot,i}\|_1 < \|b_{i,\cdot}\|_1\}$. That is, the flow into a node in S exceeds the flow out of it, and the flow into a node in T is less than the flow out of it. We have that

$$\sum_{i \in S} (\|b_{\cdot,i}\|_1 - \|b_{i,\cdot}\|_1) - \sum_{i \in T} (\|b_{i,\cdot}\|_1 - \|b_{\cdot,i}\|_1) = \sum_{i \in [n]} (\|b_{\cdot,i}\|_1 - \|b_{i,\cdot}\|_1) = 0.$$

Thus, we can view each node $i \in S$ as a source with supply $\|b_{\cdot,i}\|_1 - \|b_{i,\cdot}\|_1$, and each node $i \in T$ as a sink with demand $\|b_{i,\cdot}\|_1 - \|b_{\cdot,i}\|_1$, and the total supply equals the total demand. We now add some weighted arcs connecting the nodes in S to the nodes in T . These arcs carry the supply at the nodes in S to the demand at the nodes in T . Note that we may add arcs that are parallel to some existing arcs in G_B . Such arcs can be replaced by adding flow to the parallel existing arcs of G_B . In more detail, to compute the flows of the added arcs (or the added flow to existing arcs), we add arcs inductively as follows. We start with any pair of nodes $i \in S$ and $j \in T$, and add an arc from i to j carrying flow equal to the minimum between the supply at i and the demand at j . Adding this arc will balance one of its endpoints, but in the new graph the sum of supplies at the nodes of S is still equal to the sum of demands at the nodes of T , so we can repeat the process. (Notice that either S or T or both lose one node.) Each additional arc balances

at least one unbalanced node, so G_B gets balanced by adding at most n additional arcs from nodes in S to nodes in T . The total flow on the added arcs is exactly $\sum_{i \in S} (\|b_{\cdot, i}\|_1 - \|b_{i, \cdot}\|_1) = \frac{1}{2} \cdot \|\nabla f(\mathbf{x})\|_1$.

Let E' be the set of newly added arcs, and let $G_{B'}$ be the new graph with arc weights given by $B' = (b'_{ij})$. Since $G_{B'}$ is balanced, the arc flows form a valid circulation. We next decompose the total flow of arcs into cycles. Consider a cycle C in $G_{B'}$ that contains at least one arc from E' (i.e., $C \cap E' \neq \emptyset$). Reduce the flow on all arcs in C by $\alpha = \min_{ij \in C} b'_{ij}$. This can be viewed as peeling off from $G_{B'}$ a circulation carrying flow α . This reduces the flow on at least one arc to zero, and the remaining flow on arcs is still a valid circulation, so we can repeat the process. It can be repeated as long as there is positive flow on some arc in E' . Eliminating the flow on all arcs in E' using cycles reduces the total flow on the arcs by at most n times the total initial flow on the arcs in E' (i.e., $\frac{n}{2} \cdot \|\nabla f(\mathbf{x}^{(1)})\|_1$), because each cycle contains at most n arcs and its flow α that is peeled off reduces the flow on at least one arc in E' by α . After peeling off all the flow on all arcs in E' , all the arcs with positive flow are original arcs of G_B . Let $G_{B''}$ be the graph with the remaining arcs and their flows which are given by $B'' = (b''_{ij})$. The total flow on the arcs of $G_{B''}$ is at least $f(\mathbf{x}) + \frac{1}{2} \cdot \|\nabla f(\mathbf{x})\|_1 - \frac{n}{2} \cdot \|\nabla f(\mathbf{x})\|_1 \geq f(\mathbf{x}) - \frac{n}{2} \cdot \|\nabla f(\mathbf{x})\|_1$.

Next we show that the total flow on the arcs of $G_{B''}$ is a lower bound on $f(\mathbf{x}^*)$. Our key tool for this is the fact that balancing operations preserve the product of arc flows on any cycle in the original graph G_B , because balancing a node i multiplies the flow on the arcs into i by some factor r and the flow on the arcs out of i by $\frac{1}{r}$. Thus, the geometric mean of the flows of the arcs on any cycle is not changed by a balancing operation. The arc flows in $G_{B''}$ form a valid circulation, and thus can be decomposed into flow cycles C_1, \dots, C_q by a similar peeling-off process that was described earlier. Let n_1, \dots, n_q be the lengths of cycles, and let $\alpha_1, \dots, \alpha_q$ be their flows. The total flow on arcs in $G_{B''}$ is, therefore, $\sum_{k=1}^q n_k \cdot \alpha_k$. Notice that, by construction, $b''_{ij} \leq b_{ij}$, and the decomposition into cycles gives that $b''_{ij} = \sum_{k:ij \in C_k} \alpha_k$. Thus, $f(\mathbf{x}^*) = \sum_{i,j=1}^n b_{ij} e^{x_i^* - x_j^*} \geq \sum_{i,j=1}^n b''_{ij} e^{x_i^* - x_j^*} = \sum_{i,j=1}^n \sum_{k:ij \in C_k} \alpha_k e^{x_i^* - x_j^*} = \sum_{k=1}^q \sum_{ij \in C_k} \alpha_k e^{x_i^* - x_j^*} \geq \sum_{k=1}^q n_k \left(\prod_{ij \in C_k} \alpha_k e^{x_i^* - x_j^*} \right)^{1/n_k} = \sum_{k=1}^q n_k \alpha_k = \sum_{i,j=1}^n b''_{ij}$, where the last inequality uses the arithmetic-geometric mean inequality. Notice that the right-hand side is the total flow on the arcs of $G_{B''}$, which is at least $f(\mathbf{x}) - \frac{n}{2} \cdot \|\nabla f(\mathbf{x}^{(1)})\|_1$. Thus, $f(\mathbf{x}^*) \geq f(\mathbf{x}) - \frac{n}{2} \cdot \|\nabla f(\mathbf{x})\|_1$, and this completes the proof of the lemma. \square

3 Greedy Balancing

Here we present and analyze a greedy variant of the OPR iteration. Instead of balancing indices in a fixed round-robin order, the greedy modification chooses at iteration t an index i_t of $A^{(t)}$ such that balancing the chosen index results in the largest decrease in the sum of entries of $A^{(t)}$. In other words, we pick i_t such that the following equation holds.

$$i_t = \arg \max_{i \in [n]} \left(\sqrt{\|a_{\cdot, i}^{(t)}\|_1} - \sqrt{\|a_{i, \cdot}^{(t)}\|_1} \right)^2 \quad (4)$$

We give two analyses of this variant, one that shows that the number of operations is nearly linear in the size of G_A , and another that shows that the number of operations is nearly linear in $1/\epsilon$. More specifically, we prove the following theorem.

Theorem 2. *Given an $n \times n$ matrix A , let $m = |E(G_A)|$, the greedy implementation of the OPR iterative algorithm outputs an ϵ -balanced matrix in K iterations which cost a total of $O(m + Kn \log n)$ arithmetic operations over $O(n \log w)$ -bit numbers, where $K = O(\min\{\epsilon^{-2} \log w, \epsilon^{-1} n^{3/2} \log(w/\epsilon)\})$.*

The proof uses the convex optimization framework introduced in Section 2. Recall that $A^{(t)} = \bar{D}^{(t)} A (\bar{D}^{(t)})^{-1}$. If we let $\bar{D}^{(t)} = \text{diag}(e^{x_1^{(t)}}, \dots, e^{x_n^{(t)}})$, the iterative sequence can be viewed as generating a

sequence of points $\mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \dots, \mathbf{x}^{(t)}, \dots$ in \mathbb{R}^n , where $\mathbf{x}^{(t)} = (x_1^{(t)}, \dots, x_n^{(t)})$ and $A^{(t)} = \bar{D}^{(t)} A (\bar{D}^{(t)})^{-1} = (a_{ij} e^{x_i^{(t)} - x_j^{(t)}})_{n \times n}$. Initially, $\mathbf{x}^{(1)} = (0, \dots, 0)$, and $\mathbf{x}^{(t+1)} = \mathbf{x}^{(t)} + \alpha_t \mathbf{e}_i$, where $\alpha_t = \ln(d_{ii}^{(t)})$ and \mathbf{e}_i is the i th vector of the standard basis for \mathbb{R}^n . By Equation (1), the value $f(\mathbf{x}^{(t)})$ is sum of the entries of the matrix $A^{(t)}$. The following key lemma allows us to lower bound the decrease in the value of $f(\mathbf{x}^{(t)})$ in terms of a value that can be later related to the stopping condition.

Lemma 3. *If index i_t defined in Equation (4) is picked to balance $A^{(t)}$, then $f(\mathbf{x}^{(t)}) - f(\mathbf{x}^{(t+1)}) \geq \frac{\|\nabla f(\mathbf{x}^{(t)})\|_2^2}{4f(\mathbf{x}^{(t)})}$.*

Corollary 2. *If matrix $A^{(t)}$ is not ϵ -balanced, by balancing index i_t at iteration t , we have $f(\mathbf{x}^{(t)}) - f(\mathbf{x}^{(t+1)}) \geq \frac{\epsilon^2}{4} \cdot f(\mathbf{x}^{(t)})$.*

Proof of Theorem 2. By Corollary 2, while $A^{(t)}$ is not ϵ -balanced, there exists an index i_t to balance such that $f(\mathbf{x}^{(t)}) - f(\mathbf{x}^{(t+1)}) \geq \frac{\epsilon^2}{4} \cdot f(\mathbf{x}^{(t)})$. Thus, $f(\mathbf{x}^{(t+1)}) \leq \left(1 - \frac{\epsilon^2}{4}\right) \cdot f(\mathbf{x}^{(t)})$. Iterating for t steps yields $f(\mathbf{x}^{(t+1)}) \leq \left(1 - \frac{\epsilon^2}{4}\right)^t \cdot f(\mathbf{x}^{(1)})$. So, on the one hand, $f(\mathbf{x}^{(1)}) = \sum_{i,j=1}^n a_{ij}$ since $f(\mathbf{x}^{(1)})$ is the sum of entries in $A^{(1)}$. On the other hand, we argue that the value of $f(\mathbf{x}^{(t+1)})$ is at least $\min_{(i,j) \in E} a_{ij}$. To see this, consider a directed cycle in the graph G_A . It's easy to see that balancing operations preserve the product of weights of the arcs on any cycle. Thus, the weight of at least one arc in the cycle is at least its weight in the input matrix A . Therefore, $a_{\min} \leq f(\mathbf{x}^{(t+1)}) \leq \left(1 - \frac{\epsilon^2}{4}\right)^t \cdot f(\mathbf{x}^{(1)}) = \left(1 - \frac{\epsilon^2}{4}\right)^t \cdot \sum_{i,j=1}^n a_{ij}$. Thus, $t \leq \frac{4}{\epsilon^2} \cdot \ln w$ and this is an upper bound on the number of balancing operations before an ϵ -balanced matrix is obtained. The algorithm initially computes $\|a_{\cdot,i}\|_1$ and $\|a_{i,\cdot}\|_1$ for all $i \in [n]$ in $O(m)$ time. Also the algorithm initially computes the value of $(\sqrt{\|a_{i,\cdot}\|_1} - \sqrt{\|a_{\cdot,i}\|_1})^2$ for all i in $O(m)$ time and inserts the values in a priority queue in $O(n \log n)$ time. The values of $\|a_{i,\cdot}^{(t)}\|_1$, $\|a_{\cdot,i}^{(t)}\|_1$ for all i and $(\sqrt{\|a_{i,\cdot}^{(t)}\|_1} - \sqrt{\|a_{\cdot,i}^{(t)}\|_1})^2$ are updated after each balancing operation. In each iteration the weights of at most n arcs change. Updating the values of $\|a_{i,\cdot}^{(t)}\|_1$ and $\|a_{\cdot,i}^{(t)}\|_1$ takes $O(n)$ time and updating the values of $(\sqrt{\|a_{i,\cdot}^{(t)}\|_1} - \sqrt{\|a_{\cdot,i}^{(t)}\|_1})^2$ involves at most n updates of values in the priority queue, each taking time $O(\log n)$. Thus, the first iteration takes $O(m)$ operations and each iteration after that takes $O(n \log n)$ operations, so the total running time of the algorithm in terms of arithmetic operations is $O(m + (n \log n \log w)/\epsilon^2)$.

An alternative analysis completes the proof. Notice that $\|\nabla f(\mathbf{x}^{(t)})\|_2 \leq \|\nabla f(\mathbf{x}^{(t)})\|_1 \leq \sqrt{n} \cdot \|\nabla f(\mathbf{x}^{(t)})\|_2$. Therefore, $f(\mathbf{x}^{(t)}) - f(\mathbf{x}^{(t+1)}) \geq \frac{\|\nabla f(\mathbf{x}^{(t)})\|_2^2}{4f(\mathbf{x}^{(t)})} \geq \frac{\|\nabla f(\mathbf{x}^{(t)})\|_2}{4\sqrt{n} \cdot f(\mathbf{x}^{(t)})} \cdot \|\nabla f(\mathbf{x}^{(t)})\|_1 \geq \frac{1}{2n^{3/2}} \cdot \frac{\|\nabla f(\mathbf{x}^{(t)})\|_2}{f(\mathbf{x}^{(t)})} \cdot (f(\mathbf{x}^{(t)}) - f(\mathbf{x}^*))$, where the first inequality follows from Lemma 3, and the last inequality follows from Lemma 2. Therefore, while A^t is not ϵ -balanced (so $\frac{\|\nabla f(\mathbf{x}^{(t)})\|_2}{f(\mathbf{x}^{(t)})} > \epsilon$), we have that $f(\mathbf{x}^{(t)}) - f(\mathbf{x}^{(t+1)}) \geq \frac{\epsilon}{2n^{3/2}} \cdot (f(\mathbf{x}^{(t)}) - f(\mathbf{x}^*))$. Rearranging the terms, we get $f(\mathbf{x}^{(t+1)}) - f(\mathbf{x}^*) \leq \left(1 - \frac{\epsilon}{2n^{3/2}}\right) \cdot (f(\mathbf{x}^{(t)}) - f(\mathbf{x}^*))$. Therefore, $f(\mathbf{x}^{(t+1)}) - f(\mathbf{x}^*) \leq \left(1 - \frac{\epsilon}{2n^{3/2}}\right)^t \cdot (f(\mathbf{x}^{(1)}) - f(\mathbf{x}^*))$. Notice that by Lemma 3, $f(\mathbf{x}^{(t+1)}) - f(\mathbf{x}^*) \geq f(\mathbf{x}^{(t+1)}) - f(\mathbf{x}^{(t+2)}) \geq \left(\frac{\|\nabla f(\mathbf{x}^{(t+1)})\|_2}{2f(\mathbf{x}^{(t+1)})}\right)^2 \cdot f(\mathbf{x}^{(t+1)}) \geq \left(\frac{\|\nabla f(\mathbf{x}^{(t+1)})\|_2}{2f(\mathbf{x}^{(t+1)})}\right)^2 \cdot a_{\min}$. On the other hand, $f(\mathbf{x}^{(1)}) - f(\mathbf{x}^*) \leq f(\mathbf{x}^{(1)}) \leq \sum_{i,j=1}^n a_{ij}$. Thus, for $t = 2\epsilon^{-1} \cdot n^{3/2} \ln(4w/\epsilon^2)$, we have that $\frac{\|\nabla f(\mathbf{x}^{(t+1)})\|_2}{f(\mathbf{x}^{(t+1)})} \leq \epsilon$, so the matrix is ϵ -balanced. \square

4 Round-Robin Balancing (the original algorithm)

Recall that original Osborne-Parlett-Reinsch algorithm balances indices in a fixed round-robin order. Although the greedy variant of the OPR iteration is a simple modification of the implementation, the convergence rate of the original algorithm (with no change) is interesting. This is important because the original algorithm has a slightly simpler implementation, and also because this is the implementation used in almost all numerical linear algebra software including MATLAB, LAPACK and EISPACK (refer to [13, 7] for further background). We give some answer to this question in the following theorem.

Theorem 3. *Given an $n \times n$ matrix A , the original implementation of the OPR iteration outputs an ϵ -balanced matrix in $O(\epsilon^{-2}n^2 \log w)$ iterations totalling $O(\epsilon^{-2}mn \log w)$ arithmetic operations over $O(n \log w)$ -bit numbers (m is the number of non-zero entries of A).*

5 Randomized Balancing

In Theorem 2 the arithmetic operations were applied to $O(n \ln w)$ -bit numbers. This will cause an additional factor of $O(n \ln w)$ in the running time of the algorithm. In this section we fix this issue by presenting a randomized variant of the algorithm that applies arithmetic operations to numbers of $O(\ln(wn/\epsilon))$ bits. Thus, we obtain a algorithm for balancing that runs in nearly linear time. While the greedy algorithm works by picking the node i that maximizes $(\sqrt{\|a_{i,\cdot}\|} - \sqrt{\|a_{\cdot,i}\|})^2$, the key idea of the randomized algorithm is sampling a node for balancing using sampling probabilities that do not depend on the difference in arc weights (the algorithm uses low-precision rounded weights, so this can affect significantly the difference). Instead, our sampling probabilities depend on the sum of weights of the arcs incident on a node.

We first introduce some notation. We use $O(\ln(wn/\epsilon))$ bits of precision to approximate x_i -s with \hat{x}_i -s. Thus, $x_i - 2^{-O(\ln(nw/\epsilon))} \leq \hat{x}_i \leq x_i$. In addition to maintaining $\hat{\mathbf{x}}^{(t)} = (\hat{x}_1^{(t)}, \hat{x}_2^{(t)}, \dots, \hat{x}_n^{(t)})$ at every time t . The algorithm also maintains for every i and j the value of $\hat{a}_{ij}^{(t)}$ which is $a_{ij}e^{\hat{x}_i^{(t)} - \hat{x}_j^{(t)}}$ truncated to $O(\ln(wn/\epsilon))$ bits of precision. We set the hidden constant to give a truncation error of $r = (\epsilon/wn)^{10} a_{\min}$, so $a_{ij}^{(t)} - r \leq \hat{a}_{ij}^{(t)} \leq a_{ij}^{(t)}$. The algorithm also maintains for every i , $\|\hat{a}_{i,\cdot}^{(t)}\| = \sum_{j=1}^n \hat{a}_{ij}^{(t)}$ and $\|\hat{a}_{\cdot,i}^{(t)}\| = \sum_{j=1}^n \hat{a}_{ji}^{(t)}$. For every i , we use the notation $\|a_{i,\cdot}^{(t)}\| = \sum_{j=1}^n a_{ij}^{(t)}$ and $\|a_{\cdot,i}\| = \sum_{j=1}^n a_{ji}^{(t)}$. Note that the algorithm does not maintain the values $a_{ij}^{(t)}$, $\|a_{\cdot,i}^{(t)}\|$ or $\|a_{i,\cdot}^{(t)}\|$.

The algorithm works as follows (see the pseudo-code of Algorithm 1 that appears in Section 7). In each iteration it samples an index i with probability $p_i = \frac{\|\hat{a}_{i,\cdot}^{(t)}\| + \|\hat{a}_{\cdot,i}^{(t)}\|}{2 \sum_{i,j} \hat{a}_{ij}^{(t)}}$. If i is sampled, a balancing operation is applied to index i only if the arcs incident on i have significant weight, and i 's imbalance is sufficiently large. Put $\hat{M}_i = \max\{\|\hat{a}_{i,\cdot}^{(t)}\|, \|\hat{a}_{\cdot,i}^{(t)}\|\}$ and put $\hat{m}_i = \min\{\|\hat{a}_{i,\cdot}^{(t)}\|, \|\hat{a}_{\cdot,i}^{(t)}\|\}$. The imbalance is considered large if $\hat{m}_i = 0$ (this can happen because of the low precision), or if $\hat{m}_i \neq 0$ and $\frac{\hat{M}_i}{\hat{m}_i} \geq 1 + \frac{\epsilon}{n}$. A balancing operation is done by adding α to $\hat{\mathbf{x}}_i^{(t)}$, where $\alpha = \frac{1}{2} \ln(\|\hat{a}_{\cdot,i}^{(t)}\| / \|\hat{a}_{i,\cdot}^{(t)}\|)$, unless $\hat{m}_i = 0$, in which case we replace the 0 value by nr . This updates the weights of the arcs incident on i . Also, the L_1 norms of changed rows and columns are updated. (For convenience we use in this section $\|\cdot\|$ instead of $\|\cdot\|_1$ to denote the L_1 norm.)

Note that in the pseudo-code, \leftarrow indicates an assignment where the value on the right-hand side is computed to $O(\ln(wn/\epsilon))$ bits of precision. Thus, we have

$$\alpha - (\epsilon/wn)^{10} \leq \hat{x}_i^{(t+1)} - \hat{x}_i^{(t)} \leq \alpha \quad (5)$$

and

$$a_{ij}e^{\hat{x}_i^{(t+1)} - \hat{x}_j^{(t+1)}} - r \leq \hat{a}_{ij}^{(t+1)} \leq a_{ij}e^{\hat{x}_i^{(t+1)} - \hat{x}_j^{(t+1)}},$$

$$a_{ji}e^{\hat{x}_i^{(t+1)} - \hat{x}_j^{(t+1)}} - r \leq \hat{a}_{ji}^{(t+1)} \leq a_{ji}e^{\hat{x}_i^{(t+1)} - \hat{x}_j^{(t+1)}}.$$

Theorem 4. *With probability at least $\frac{9}{10}$, Algorithm 1 returns in time $O(m \ln \sum_{ij} a_{ij} + \epsilon^{-2} n \ln(wn/\epsilon) \ln w)$ an ϵ -balanced matrix.*

The idea of proof is to show that in every iteration of the algorithm we reduce $f(\cdot)$ by at least a factor of $1 - \Omega(\epsilon^2)$. Before we prove the theorem, we state and prove a couple of useful lemmas.

Fix an iteration t , and define three sets of indices as follows: $A = \{i : \|\hat{a}_{i,\cdot}^{(t)}\| + \|\hat{a}_{\cdot,i}^{(t)}\| \geq \epsilon a_{\min}/10wn\}$, $B = \{i : \hat{m}_i \neq 0 \wedge \hat{M}_i/\hat{m}_i \geq 1 + \epsilon/n\}$, and $C = \{i : \hat{m}_i = 0\}$. If the random index i satisfies $i \notin A$ or $i \in A \setminus (B \cup C)$, the algorithm does not perform any balancing operation on i . The following lemma states that the expected decrease due to balancing such indices is small, and thus skipping them does not affect the speed of convergence substantially.

Lemma 4. *For every iteration t , $\sum_{i \notin A \cap (B \cup C)} p_i \cdot \left(\sqrt{\|a_{i,\cdot}^{(t)}\|} - \sqrt{\|a_{\cdot,i}^{(t)}\|} \right)^2 < \frac{2\epsilon^2}{n} \cdot f(\mathbf{x}^{(t)})$, where p is the probability distribution over indices at time t .*

We now show a lower bound on the decrease in $f(\cdot)$, if a node $i \in A \cap (B \cup C)$ is balanced.

Lemma 5. *If $i \in A \cap (B \cup C)$ is balanced in iteration t , then $f(\hat{\mathbf{x}}^{(t+1)}) - f(\hat{\mathbf{x}}^{(t)}) \geq \frac{1}{10} \cdot \left(\sqrt{\|a_{i,\cdot}^{(t)}\|} - \sqrt{\|a_{\cdot,i}^{(t)}\|} \right)^2$.*

Proof of Theorem 4. By Lemma 5, the expected decrease in $f(\cdot)$ in iteration t is lower bounded as follows.

$$\begin{aligned} \mathbb{E}[f(\hat{\mathbf{x}}^{(t)}) - f(\hat{\mathbf{x}}^{(t+1)})] &\geq \sum_{i \in A \cap (B \cup C)} p_i \cdot \frac{1}{10} \left(\sqrt{\|a_{i,\cdot}^{(t)}\|} - \sqrt{\|a_{\cdot,i}^{(t)}\|} \right)^2 \\ &= \frac{1}{10} \cdot \left(\sum_{i=1}^n p_i \cdot \left(\sqrt{\|a_{i,\cdot}^{(t)}\|} - \sqrt{\|a_{\cdot,i}^{(t)}\|} \right)^2 - \sum_{i \notin A \cap (B \cup C)} p_i \cdot \left(\sqrt{\|a_{i,\cdot}^{(t)}\|} - \sqrt{\|a_{\cdot,i}^{(t)}\|} \right)^2 \right). \end{aligned}$$

The second term can be bounded, using Lemma 4, by $\sum_{i \notin A \cap (B \cup C)} \left(\sqrt{\|a_{i,\cdot}^{(t)}\|} - \sqrt{\|a_{\cdot,i}^{(t)}\|} \right)^2 \leq \frac{2\epsilon^2}{n} \cdot f(\hat{\mathbf{x}}^{(t)})$.

For the first term, we can write

$$\begin{aligned} \sum_{i=1}^n p_i \cdot \left(\sqrt{\|a_{i,\cdot}^{(t)}\|} - \sqrt{\|a_{\cdot,i}^{(t)}\|} \right)^2 &\geq \sum_{i=1}^n p_i \cdot \frac{(\|a_{i,\cdot}^{(t)}\| - \|a_{\cdot,i}^{(t)}\|)^2}{2(\|a_{i,\cdot}^{(t)}\| + \|a_{\cdot,i}^{(t)}\|)} \\ &= \sum_{i=1}^n \frac{\|\hat{a}_{i,\cdot}^{(t)}\| + \|\hat{a}_{\cdot,i}^{(t)}\|}{2 \sum_{ij} \hat{a}_{ij}^{(t)}} \cdot \frac{(\|a_{i,\cdot}^{(t)}\| - \|a_{\cdot,i}^{(t)}\|)^2}{2(\|a_{i,\cdot}^{(t)}\| + \|a_{\cdot,i}^{(t)}\|)} \\ &\geq \frac{1}{16} \cdot \sum_{i=1}^n \frac{(\|a_{i,\cdot}^{(t)}\| - \|a_{\cdot,i}^{(t)}\|)^2}{\sum_{ij} a_{ij}^{(t)}} \\ &= \frac{\|\nabla f(\hat{\mathbf{x}}^{(t)})\|_2^2}{16f(\hat{\mathbf{x}}^{(t)})} \geq \frac{\epsilon^2}{16} \cdot f(\hat{\mathbf{x}}^{(t)}), \end{aligned}$$

where the penultimate inequality holds because $\frac{\hat{M}_i}{\hat{M}_i} \geq \frac{1}{2}$, so $\frac{\|\hat{a}_{i,\cdot}^{(t)}\| + \|\hat{a}_{\cdot,i}^{(t)}\|}{\|a_{i,\cdot}^{(t)}\| + \|a_{\cdot,i}^{(t)}\|} \geq \frac{\hat{M}_i}{2\hat{M}_i} \geq \frac{1}{4}$, and the last inequality

holds as long as the matrix is not ϵ -balanced, so $\frac{\|\nabla f(\widehat{\mathbf{x}}^{(t)})\|_2}{f(\widehat{\mathbf{x}}^{(t)})} \geq \epsilon$. Combining everything together, we get

$$\begin{aligned} \mathbb{E}[f(\widehat{\mathbf{x}}^{(t)}) - f(\widehat{\mathbf{x}}^{(t+1)})] &\geq \frac{1}{10} \cdot \left(\sum_{i=1}^n p_i \cdot \left(\sqrt{\|a_{i,\cdot}^{(t)}\|} - \sqrt{\|a_{\cdot,i}^{(t)}\|} \right)^2 - \sum_{i \notin A \cap (B \cup C)} p_i \cdot \left(\sqrt{\|a_{i,\cdot}^{(t)}\|} - \sqrt{\|a_{\cdot,i}^{(t)}\|} \right)^2 \right) \\ &\geq \frac{1}{10} \cdot \left(\frac{\epsilon^2}{16} \cdot f(\widehat{\mathbf{x}}^{(t)}) - \frac{2\epsilon^2}{n} \cdot f(\widehat{\mathbf{x}}^{(t)}) \right) \geq \frac{\epsilon^2}{320} \cdot f(\widehat{\mathbf{x}}^{(t)}), \end{aligned}$$

where the last inequality assumes $n \geq 64$. This implies that the expected number of iterations to obtain an ϵ -balanced matrix is $O(\epsilon^{-2} \ln w)$. Markov's inequality implies that with probability $\frac{9}{10}$ an ϵ -balanced matrix is obtained in $O(\epsilon^{-2} \ln w)$ iterations. It is easy to see that each iteration of the algorithm takes $O(n \ln(wn/\epsilon))$ time. Initializations take $O(m \ln \sum_{ij} a_{ij})$ time. So the total running time of the algorithm is $O(m \ln \sum_{ij} a_{ij} + \epsilon^{-2} n \ln(wn/\epsilon) \ln w)$. \square

6 A Lower Bound on the Rate of Convergence

In this section we prove the following lower bound.

Theorem 5. *There are matrices for which all variants of the Osborne-Parlett-Reinsch iteration (i.e., regardless of the order of indices chosen to balance) require $\Omega(1/\sqrt{\epsilon})$ iterations to balance the matrix to the relative error of ϵ .*

Before proving this theorem, we present the claimed construction. Let A be the following 4×4 matrix, and let A^* denote the corresponding fully-balanced matrix.

$$A = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 1 & 0 & \beta + \epsilon & 0 \\ 0 & \epsilon & 0 & 1 \\ 0 & 0 & 1 & 0 \end{bmatrix}, \quad A^* = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 1 & 0 & \sqrt{\epsilon(\beta + \epsilon)} & 0 \\ 0 & \sqrt{\epsilon(\beta + \epsilon)} & 0 & 1 \\ 0 & 0 & 1 & 0 \end{bmatrix}$$

Here $\epsilon > 0$ is arbitrarily small, and $\beta = 100\epsilon$. It's easy to see that $A^* = D^* A D^{*-1}$ where

$$D = \text{diag} \left(1, 1, \sqrt{\frac{\beta + \epsilon}{\epsilon}}, \sqrt{\frac{\beta + \epsilon}{\epsilon}} \right).$$

To prove Theorem 5, we show that balancing A to the relative error of ϵ requires $\Omega(1/\sqrt{\epsilon})$ iterations, regardless of the order of balancing operations. Notice that in order to fully balance A , we simply need to replace a_{23} and a_{32} by their geometric mean. We measure the rate of convergence using the ratio a_{32}/a_{23} . This ratio is initially $\frac{\epsilon}{\beta + \epsilon} = \frac{1}{101}$. When the matrix is fully balanced, the ratio becomes 1. We show that this ratio increases by a small factor in each iteration, and that it has to increase sufficiently for the matrix to be ϵ -balanced. This is summarized in the following two lemmas.

Lemma 6 (change in ratio). $\frac{a_{32}^{(t+1)}}{a_{23}^{(t+1)}} \leq \left(\frac{1 + 7\sqrt{\beta}}{1 + \epsilon} \right) \cdot \frac{a_{32}^{(t)}}{a_{23}^{(t)}}$.

Lemma 7 (stopping condition). *If $A^{(t)}$ is ϵ -balanced, then $\frac{a_{32}^{(t)}}{a_{23}^{(t)}} > \frac{1}{100}$.*

Before proving the two lemmas we show how they lead to the proof of Theorem 5.

Proof of Theorem 5. By Lemma 6, $\frac{a_{32}^{(t+1)}}{a_{23}^{(t+1)}} \leq \left(\frac{1+7\sqrt{\beta}}{1+\epsilon}\right)^t \cdot \frac{a_{32}}{a_{23}} = \left(\frac{1+7\sqrt{\beta}}{1+\epsilon}\right)^t \cdot \frac{\epsilon}{\beta+\epsilon}$. By Lemma 7, if $A^{(t+1)}$ is ϵ -balanced, then $\frac{1}{100} < \frac{a_{32}^{(t+1)}}{a_{23}^{(t+1)}} \leq \left(\frac{1+7\sqrt{\beta}}{1+\epsilon}\right)^t \cdot \frac{\epsilon}{\beta+\epsilon} \leq (1+7\sqrt{\beta})^t \cdot \frac{\epsilon}{\beta+\epsilon}$. Using $\beta = 100\epsilon$, we get the condition that $(1+7\sqrt{\beta})^t > \frac{101}{100}$, which implies that $t = \Omega(1/\sqrt{\epsilon})$. \square

Proof of Lemma 6. Using the notation we defined earlier, we have that $f(\mathbf{x}^{(1)}) = \sum_{i,j=1}^4 a_{ij} = 4 + 2\epsilon + \beta$ and $f(\mathbf{x}^*) = \sum_{i,j=1}^4 a_{ij}^* = 4 + 2\sqrt{\epsilon(\beta+\epsilon)}$, so $f(\mathbf{x}^{(1)}) - f(\mathbf{x}^*) < \beta$. We observe that at each iteration t , $a_{12}^{(t)}a_{21}^{(t)} = a_{34}^{(t)}a_{43}^{(t)} = 1$ and $a_{23}^{(t)}a_{32}^{(t)} = \epsilon(\beta+\epsilon)$ because the product of weights of arcs on any cycle in G_A is preserved (for instance, arcs (1, 2) and (2, 1) form a cycle and initially $a_{12}a_{21} = 1$).

The ratio $a_{32}^{(t)}/a_{23}^{(t)}$ is only affected in iterations that balance index 2 or 3. Let's assume a balancing operation at index 2, a similar analysis applies to balancing at index 3. By balancing at index 2 at time t we have

$$\frac{a_{32}^{(t+1)}}{a_{32}^{(t)}} = \frac{a_{23}^{(t)}}{a_{23}^{(t+1)}} = \sqrt{\frac{a_{21}^{(t)} + a_{23}^{(t)}}{a_{12}^{(t)} + a_{32}^{(t)}}}. \quad (6)$$

Thus, to prove Lemma 6, it suffices to show that

$$\frac{a_{32}^{(t+1)}}{a_{32}^{(t)}} \cdot \frac{a_{23}^{(t)}}{a_{23}^{(t+1)}} = \frac{a_{21}^{(t)} + a_{23}^{(t)}}{a_{12}^{(t)} + a_{32}^{(t)}} \leq \frac{1 + 7\sqrt{\beta}}{1 + \epsilon}. \quad (7)$$

By our previous observation, $a_{12}^{(t)}a_{21}^{(t)} = 1$, so if $a_{21}^{(t)} = y$, then $a_{12}^{(t)} = 1/y$. Similarly $a_{23}^{(t)}a_{32}^{(t)} = \epsilon(\beta+\epsilon)$ implies that there exists z such that $a_{23}^{(t)} = (\beta+\epsilon)z$ and $a_{32}^{(t)} = \epsilon/z$. Therefore:

$$\frac{a_{21}^{(t)} + a_{23}^{(t)}}{a_{12}^{(t)} + a_{32}^{(t)}} = \frac{y + (\beta+\epsilon)z}{(1/y) + (\epsilon/z)} \quad (8)$$

We bound the right hand side of Equation (8) by proving upper bounds on y and z . We first show that $y < 1 + 2\sqrt{\beta}$. To see this notice that on the one hand,

$$f(\mathbf{x}^{(t)}) = \sum_{i,j=1}^4 a_{ij}^{(t)} = a_{12}^{(t)} + a_{21}^{(t)} + a_{23}^{(t)} + a_{32}^{(t)} + a_{34}^{(t)} + a_{43}^{(t)} \geq y + \frac{1}{y} + 2\sqrt{\epsilon(\beta+\epsilon)} + 2, \quad (9)$$

where we used $a_{34}^{(t)} + a_{43}^{(t)} \geq 2$ and $a_{34}^{(t)}a_{43}^{(t)} \geq 2\sqrt{\epsilon(\beta+\epsilon)}$, both implied by the arithmetic-geometric mean inequality. On the other hand,

$$f(\mathbf{x}^{(t)}) \leq f(\mathbf{x}^{(1)}) \leq f(\mathbf{x}^*) + \beta = 4 + 2\sqrt{\epsilon(\beta+\epsilon)} + \beta. \quad (10)$$

Combining Equations (9) and (10) together, we have $y + (1/y) - 2 \leq \beta$. For sufficiently small ϵ , the last inequality implies, in particular, that $y < 2$. Thus, we have $(y-1)^2 \leq y\beta < 2\beta$, and this implies that $y < 1 + 2\sqrt{\beta}$.

Next we show that $z \leq 1$. Assume for contradiction that $z > 1$. By the arithmetic-geometric mean inequality $a_{12}^{(t)} + a_{21}^{(t)} \geq 2$ and $a_{34}^{(t)} + a_{43}^{(t)} \geq 2$. Thus,

$$f(\mathbf{x}^{(t)}) = \sum_{i,j=1}^4 a_{ij}^{(t)} \geq 2 + (\beta+\epsilon)z + \frac{\epsilon}{z} + 2 = 4 + \beta z + \epsilon \left(z + \frac{1}{z} \right) > 4 + \beta + 2\epsilon = f(\mathbf{x}^{(1)}),$$

where the last inequality follows because $z > 1$, and $z + 1/z > 2$. But this is a contradiction, because each balancing iteration reduces the value of F , so $f(\mathbf{x}^{(t)}) \leq f(\mathbf{x}^{(1)})$.

We can now bound $(a_{21}^{(t)} + a_{23}^{(t)})/(a_{12}^{(t)} + a_{32}^{(t)})$. By Equation (8), and using our bounds for y and z ,

$$\frac{a_{21}^{(t)} + a_{23}^{(t)}}{a_{12}^{(t)} + a_{32}^{(t)}} = \frac{y + (\beta + \epsilon)z}{(1/y) + (\epsilon/z)} \leq \frac{(1 + 2\sqrt{\beta}) + (\beta + \epsilon)}{\frac{1}{1 + 2\sqrt{\beta}} + \epsilon} \leq \frac{1 + 4\sqrt{\beta}}{\frac{1}{1 + 2\sqrt{\beta}} + \frac{\epsilon}{1 + 2\sqrt{\beta}}} \leq \frac{1 + 7\sqrt{\beta}}{1 + \epsilon}.$$

The last line uses the fact that $\sqrt{\beta} \gg \beta = 100\epsilon \geq \epsilon$, which holds if ϵ is sufficiently small. \square

Proof of Lemma 7. Let $t - 1$ be the last iteration before an ϵ -balanced matrix is obtained. We argued that there is $z \leq 1$ such that $a_{23}^{(t)} = (\beta + \epsilon)z$ and $a_{32}^{(t)} = \epsilon/z$. Assume for the sake of contradiction that $a_{32}^{(t)}/a_{23}^{(t)} < 1/100$. This implies that $(\epsilon/z)/((\beta + \epsilon)z) < 1/100$, and thus $z^2 > 100/101$. So, we get

$$\begin{aligned} f(\mathbf{x}^{(t)}) - f(\mathbf{x}^*) &\geq a_{23}^{(t)} + a_{32}^{(t)} - 2\sqrt{a_{23}^{(t)}a_{32}^{(t)}} = \left(\sqrt{a_{23}^{(t)}} - \sqrt{a_{32}^{(t)}}\right)^2 \geq a_{23}^{(t)}\left(1 - \sqrt{\frac{1}{100}}\right)^2 \\ &= 0.81 \cdot (\beta + \epsilon)z \geq 0.81 \cdot (\beta + \epsilon) \cdot \sqrt{\frac{100}{101}} \geq 81 \cdot \epsilon. \end{aligned} \quad (11)$$

By Lemma 2, the left hand side the of above can be bounded as follows.

$$f(\mathbf{x}^{(t)}) - f(\mathbf{x}^*) \leq n\|\nabla f(\mathbf{x}^{(t)})\|_1 \leq n^2\|\nabla f(\mathbf{x}^{(t)})\|_2 \quad (12)$$

Note that for sufficiently small ϵ , $f(\mathbf{x}^{(t)}) \leq f(\mathbf{x}^{(1)}) \leq 5$. Combining Equations (11) and (12), and using $n = 4$ and $f(\mathbf{x}^{(t)}) \geq 5$, we get that

$$\frac{\|\nabla f(\mathbf{x}^{(t)})\|_2}{f(\mathbf{x}^{(t)})} > \frac{81}{80} \cdot \epsilon > \epsilon. \quad (13)$$

By Equation (3), this contradicts our assumption that $t - 1$ is the last iteration. \square

7 Proofs

Proof of Lemma 3. The value $f(\mathbf{x}^{(t)})$ is the sum of the entries of $A^{(t)}$. By Lemma 1, balancing the i -th index of $A^{(t)}$ reduces the value of $f(\mathbf{x}^{(t)})$ by $\left(\sqrt{\|a_{\cdot,i}^{(t)}\|_1} - \sqrt{\|a_{i,\cdot}^{(t)}\|_1}\right)^2$. To simplify notation, we drop the superscript t in the following equations. We have

$$\left(\sqrt{\|a_{\cdot,i}\|_1} - \sqrt{\|a_{i,\cdot}\|_1}\right)^2 = \frac{(\|a_{\cdot,i}\|_1 - \|a_{i,\cdot}\|_1)^2}{(\sqrt{\|a_{\cdot,i}\|_1} + \sqrt{\|a_{i,\cdot}\|_1})^2} \geq \frac{(\|a_{\cdot,i}\|_1 - \|a_{i,\cdot}\|_1)^2}{2(\|a_{\cdot,i}\|_1 + \|a_{i,\cdot}\|_1)}. \quad (14)$$

It is easy to see that

$$\max_{i \in [n]} \frac{(\|a_{\cdot,i}\|_1 - \|a_{i,\cdot}\|_1)^2}{(\|a_{\cdot,i}\|_1 + \|a_{i,\cdot}\|_1)} \geq \frac{\sum_{i=1}^n (\|a_{\cdot,i}\|_1 - \|a_{i,\cdot}\|_1)^2}{\sum_{i=1}^n (\|a_{\cdot,i}\|_1 + \|a_{i,\cdot}\|_1)}. \quad (15)$$

But the right hand side of the above inequality (after resuming the use of the superscript t) equals $\frac{\|\nabla f(\mathbf{x}^{(t)})\|_2^2}{2f(\mathbf{x}^{(t)})}$. This is because for all i , $(\|a_{\cdot,i}^{(t)}\|_1 - \|a_{i,\cdot}^{(t)}\|_1)$ is by Equation (2) the i -th coordinate of $\nabla f(\mathbf{x}^{(t)})$, and in the denominator $\sum_{i=1}^n (\|a_{\cdot,i}^{(t)}\|_1 + \|a_{i,\cdot}^{(t)}\|_1) = 2f(\mathbf{x}^{(t)})$. Together with Equations (14) and (15),

this implies that balancing $i_t = \arg \max_{i \in [n]} \left\{ \left(\sqrt{\|a_{\cdot,i}^{(t)}\|_1} - \sqrt{\|a_{i,\cdot}^{(t)}\|_1}\right)^2 \right\}$ decreases $f(\mathbf{x}^{(t)})$ by the claimed value. \square

Proof of Corollary 2. From Equation (3), we know that the diagonal matrix $\text{diag}(e^{x_1}, \dots, e^{x_n})$ balances A with relative error ϵ if and only if $\frac{\|\nabla f(\mathbf{x})\|_2}{f(\mathbf{x})} \leq \epsilon$. Thus, if $A^{(t)}$ is not ϵ -balanced, $\frac{\|\nabla f(\mathbf{x}^{(t)})\|_2}{f(\mathbf{x}^{(t)})} > \epsilon$. By Lemma 3, $f(\mathbf{x}^{(t)}) - f(\mathbf{x}^{(t+1)}) \geq \frac{\|\nabla f(\mathbf{x}^{(t)})\|_2^2}{4f(\mathbf{x}^{(t)})} = \frac{1}{4} \cdot \left(\frac{\|\nabla f(\mathbf{x}^{(t)})\|_2}{f(\mathbf{x}^{(t)})}\right)^2 \cdot f(\mathbf{x}^{(t)}) \geq \frac{\epsilon^2}{4} \cdot f(\mathbf{x}^{(t)})$. \square

Proof of Theorem 3. In the original Osborne-Parlett-Reinsch algorithm, the indices are balanced in a fixed round-robin order. A *round* of balancing is a sequence of n balancing operations where each index is balanced exactly once. Thus, in the OPR algorithm all n indices are balanced in the same order every round. We prove a more general statement that any algorithm that balances indices in rounds (even if the indices are not balanced in the same order every round) obtains an ϵ -balanced matrix in at most $O((n \log w)/\epsilon^2)$ rounds. To this end, we show that applying a round of balancing to a matrix that is not ϵ -balanced reduces the value of function f at least by a factor of $1 - \epsilon^2/16n$.

To simplify notation, we consider applying a round of balancing to the initial matrix $A^{(1)} = A$. The argument clearly holds for any time- t matrix $A^{(t)}$. If A is not ϵ -balanced, by Lemma 3 and Corollary 2, there exists an index i such that by balancing i the value of f is reduced by:

$$f(\mathbf{x}^{(1)}) - f(\mathbf{x}^{(2)}) = \left(\sqrt{\|a_{\cdot,i}\|_1} - \sqrt{\|a_{i,\cdot}\|_1} \right)^2 \geq \frac{\epsilon^2}{4} f(\mathbf{x}^{(1)}). \quad (16)$$

If i is the first index to balance in the next round of balancing, then in that round the value of f is reduced at least by a factor of $1 - \epsilon^2/4 \geq 1 - \epsilon^2/16n$, and we are done. Consider the graph G_A corresponding to the matrix A . If node i is not the first node in G_A to be balanced, then some of its neighbors in the graph G_A might be balanced before i . The main problem is that balancing neighbors of i before i may reduce the imbalance of i significantly, so we cannot argue that when we reach i and balance it the value of f reduces significantly. Nevertheless, we show that balancing i and its neighbors in this round will reduce the value of f by at least the desired amount. Let t denote the time that i is balanced in the round. For every arc (j, i) into i , let $\delta_j = |a_{ji} - a_{ji}^{(t)}|$, and for every arc (i, j) out of i let $\sigma_j = |a_{ij} - a_{ij}^{(t)}|$. These values measure the weight change of these arcs due to balancing a neighbor of i at any time since the beginning of the round. The next lemma shows if the weight of an arc incident on i has changed since the beginning of the round, it must have reduced the value of f .

Claim 1. *If balancing node j changes a_{ji} to $a_{ji} + \delta$, then the balancing reduces the value of f by at least δ^2/a_{ji} . Similarly if balancing node j changes a_{ij} to $a_{ij} + \delta$, then the balancing reduces the value of f by at least δ^2/a_{ij} .*

Proof. To simplify notation we assume that j is balanced in the first iteration of the round. If balancing j changes a_{ji} to $a_{ji} + \delta$, then by the definition of balancing,

$$\frac{a_{ji} + \delta}{a_{ji}} = \sqrt{\frac{\|a_{\cdot,j}\|_1}{\|a_{j,\cdot}\|_1}}. \quad (17)$$

Thus, by Lemma 1 the value of f reduces by

$$\left(\sqrt{\|a_{\cdot,j}\|_1} - \sqrt{\|a_{j,\cdot}\|_1} \right)^2 = \left(\sqrt{\frac{\|a_{\cdot,j}\|_1}{\|a_{j,\cdot}\|_1}} - 1 \right)^2 \|a_{j,\cdot}\|_1 = \left(\frac{a_{ji} + \delta}{a_{ji}} - 1 \right)^2 \|a_{j,\cdot}\|_1 = \left(\frac{\delta}{a_{ji}} \right)^2 \|a_{j,\cdot}\|_1 \geq \frac{\delta^2}{a_{ji}}$$

The proof for the second part of the claim is similar. \square

Going back to the proof of Theorem 3, let t denote the iteration in the round that i is balanced. By Claim 1, balancing neighbors of i has already reduced the value of f by

$$\sum_{j:(j,i) \in E} \frac{\delta_j^2}{a_{ji}} + \sum_{j:(i,j) \in E} \frac{\sigma_j^2}{a_{ij}}. \quad (18)$$

Balancing i reduces value of f by an additional $\left(\sqrt{\|a_{\cdot,i}^{(t)}\|_1} - \sqrt{\|a_{i,\cdot}^{(t)}\|_1}\right)^2$, so the value of f in the current round is reduced by at least:

$$R = \sum_{j:(j,i) \in E} \frac{\delta_j^2}{a_{ji}} + \sum_{j:(i,j) \in E} \frac{\sigma_j^2}{a_{ij}} + \left(\sqrt{\|a_{\cdot,i}^{(t)}\|_1} - \sqrt{\|a_{i,\cdot}^{(t)}\|_1}\right)^2$$

Assume without loss of generality that $\|a_{i,\cdot}\|_1 > \|a_{\cdot,i}\|_1$. To lower bound R , we consider two cases:

case (i) $\sum_{j:(j,i) \in E} \delta_j + \sum_{j:(i,j) \in E} \sigma_j \geq \frac{1}{2}(\|a_{i,\cdot}\|_1 - \|a_{\cdot,i}\|_1)$. In this case,

$$\begin{aligned} R &\geq \sum_{j:(j,i) \in E} \frac{\delta_j^2}{a_{ji}} + \sum_{j:(i,j) \in E} \frac{\sigma_j^2}{a_{ij}} \geq \frac{1}{\|a_{\cdot,i}\|_1} \sum_{j:(j,i) \in E} \delta_j^2 + \frac{1}{\|a_{i,\cdot}\|_1} \sum_{j:(i,j) \in E} \sigma_j^2 \\ &\geq \frac{1}{n\|a_{\cdot,i}\|_1} \left(\sum_{j:(j,i) \in E} \delta_j\right)^2 + \frac{1}{n\|a_{i,\cdot}\|_1} \left(\sum_{j:(i,j) \in E} \sigma_j\right)^2, \end{aligned} \quad (19)$$

where the last inequality follows by Cauchy-Schwarz inequality. By assumption of case (i),

$$\max\left(\sum_{j:(j,i) \in E} \delta_j, \sum_{j:(i,j) \in E} \sigma_j\right) \geq \frac{1}{4}(\|a_{i,\cdot}\|_1 - \|a_{\cdot,i}\|_1) \quad (20)$$

Equations (19) and (20) together imply that

$$\begin{aligned} R &\geq \frac{(\sum_{j:(j,i) \in E} \delta_j)^2 + (\sum_{j:(i,j) \in E} \sigma_j)^2}{n \max(\|a_{\cdot,i}\|_1, \|a_{i,\cdot}\|_1)} \geq \frac{1}{16n} \frac{(\|a_{i,\cdot}\|_1 - \|a_{\cdot,i}\|_1)^2}{\max(\|a_{\cdot,i}\|_1, \|a_{i,\cdot}\|_1)} \\ &= \frac{(\sqrt{\|a_{\cdot,i}\|_1} - \sqrt{\|a_{i,\cdot}\|_1})^2 (\sqrt{\|a_{\cdot,i}\|_1} + \sqrt{\|a_{i,\cdot}\|_1})^2}{16n \max(\|a_{\cdot,i}\|_1, \|a_{i,\cdot}\|_1)} \geq \frac{1}{16n} \left(\sqrt{\|a_{\cdot,i}\|_1} - \sqrt{\|a_{i,\cdot}\|_1}\right)^2. \end{aligned}$$

case (ii) $\sum_{j:(j,i) \in E} \delta_j + \sum_{j:(i,j) \in E} \sigma_j < \frac{1}{2}(\|a_{i,\cdot}\|_1 - \|a_{\cdot,i}\|_1)$. By definition of δ_j 's and σ_j 's:

$$\|a_{\cdot,i}\|_1 - \sum_{j:(j,i) \in E} \delta_j \leq \|a_{\cdot,i}^{(t)}\|_1 \leq \|a_{\cdot,i}\|_1 + \sum_{j:(j,i) \in E} \delta_j \quad (21)$$

$$\|a_{i,\cdot}\|_1 - \sum_{j:(i,j) \in E} \sigma_j \leq \|a_{i,\cdot}^{(t)}\|_1 \leq \|a_{i,\cdot}\|_1 + \sum_{j:(i,j) \in E} \sigma_j. \quad (22)$$

Combining Equations (21) and (22), and the assumption of case (ii) gives:

$$\|a_{i,\cdot}^{(t)}\|_1 + \|a_{\cdot,i}^{(t)}\|_1 \leq \|a_{i,\cdot}\|_1 + \|a_{\cdot,i}\|_1 + \sum_{j:(i,j) \in E} \sigma_j + \sum_{j:(j,i) \in E} \delta_j \leq 2(\|a_{i,\cdot}\|_1 + \|a_{\cdot,i}\|_1) \quad (23)$$

$$\|a_{i,\cdot}^{(t)}\|_1 - \|a_{\cdot,i}^{(t)}\|_1 \geq \|a_{i,\cdot}\|_1 - \|a_{\cdot,i}\|_1 - \sum_{j:(i,j) \in E} \sigma_j - \sum_{j:(j,i) \in E} \delta_j \geq \frac{1}{2}(\|a_{i,\cdot}\|_1 - \|a_{\cdot,i}\|_1). \quad (24)$$

Using Equations (23) and (24), we can write:

$$\begin{aligned} R &\geq \left(\sqrt{\|a_{i,\cdot}^{(t)}\|_1} - \sqrt{\|a_{\cdot,i}^{(t)}\|_1} \right)^2 = \frac{\left(\|a_{i,\cdot}^{(t)}\|_1 - \|a_{\cdot,i}^{(t)}\|_1 \right)^2}{\left(\sqrt{\|a_{i,\cdot}^{(t)}\|_1} + \sqrt{\|a_{\cdot,i}^{(t)}\|_1} \right)^2} \geq \frac{\left(\|a_{i,\cdot}^{(t)}\|_1 - \|a_{\cdot,i}^{(t)}\|_1 \right)^2}{8 \left(\|a_{i,\cdot}^{(t)}\|_1 + \|a_{\cdot,i}^{(t)}\|_1 \right)} \\ &\geq \frac{\left(\|a_{i,\cdot}^{(t)}\|_1 - \|a_{\cdot,i}^{(t)}\|_1 \right)^2}{16 \left(\|a_{i,\cdot}^{(t)}\|_1 + \|a_{\cdot,i}^{(t)}\|_1 \right)} \geq \frac{1}{16} \left(\sqrt{\|a_{i,\cdot}^{(t)}\|_1} - \sqrt{\|a_{\cdot,i}^{(t)}\|_1} \right)^2. \end{aligned}$$

Thus, we have shown in both cases that in one round the balancing operations on node i and its neighbors reduces the value of f by at least

$$\frac{1}{16n} \left(\sqrt{\|a_{\cdot,i}\|_1} - \sqrt{\|a_{i,\cdot}\|_1} \right)^2, \quad (25)$$

which in turn is at least $\Omega\left(\frac{\epsilon^2}{n} f(\mathbf{x}^{(1)})\right)$ by Equation (16). Thus, we have shown that if A is not ϵ -balanced, one round of balancing (where each index is balanced exactly once) reduces the objective function f by a factor of at least $1 - \Omega\left(\frac{\epsilon^2}{n} f(\mathbf{x}^{(1)})\right)$. By an argument similar to the one in the proof of Theorem 2, we get that the algorithm obtains an ϵ -balanced matrix in at most $O(\epsilon^{-2} n \log w)$ rounds. The number of balancing iterations in each round is n , and the number of arithmetic operations in each round is $O(m)$, so the original OPR algorithm obtains an ϵ -balanced matrix using $O(\epsilon^{-2} mn \log w)$ arithmetic operations. \square

Proof of Lemma 4. Notice that for every i ,

$$\left(\|\widehat{a}_{i,\cdot}^{(t)}\| + \|\widehat{a}_{\cdot,i}^{(t)}\| \right) \cdot \left(\sqrt{\|a_{i,\cdot}^{(t)}\|} - \sqrt{\|a_{\cdot,i}^{(t)}\|} \right)^2 \leq \left(\|\widehat{a}_{i,\cdot}^{(t)}\| + \|\widehat{a}_{\cdot,i}^{(t)}\| \right) \cdot \frac{\left(\|a_{i,\cdot}^{(t)}\| - \|a_{\cdot,i}^{(t)}\| \right)^2}{\|a_{i,\cdot}^{(t)}\| + \|a_{\cdot,i}^{(t)}\|} \leq \left(\|a_{i,\cdot}^{(t)}\| - \|a_{\cdot,i}^{(t)}\| \right)^2,$$

because $\|\widehat{a}_{i,\cdot}^{(t)}\| + \|\widehat{a}_{\cdot,i}^{(t)}\| \leq \|a_{i,\cdot}^{(t)}\| + \|a_{\cdot,i}^{(t)}\|$. We first bound the sum over $i \notin A$.

$$\begin{aligned} \sum_{i \notin A} p_i \cdot \left(\sqrt{\|a_{i,\cdot}^{(t)}\|} - \sqrt{\|a_{\cdot,i}^{(t)}\|} \right)^2 &= \sum_{i \notin A} \frac{\|\widehat{a}_{i,\cdot}^{(t)}\| + \|\widehat{a}_{\cdot,i}^{(t)}\|}{2 \sum_{i,j} \widehat{a}_{ij}^{(t)}} \cdot \left(\sqrt{\|a_{i,\cdot}^{(t)}\|} - \sqrt{\|a_{\cdot,i}^{(t)}\|} \right)^2 \\ &\leq \frac{1}{2 \sum_{i,j} \widehat{a}_{ij}^{(t)}} \cdot \sum_{i \notin A} \left(\|a_{i,\cdot}^{(t)}\| - \|a_{\cdot,i}^{(t)}\| \right)^2 \\ &\leq \frac{1}{2 \sum_{i,j} \widehat{a}_{ij}^{(t)}} \cdot \sum_{i \notin A} \left(\|\widehat{a}_{i,\cdot}^{(t)}\| + \|\widehat{a}_{\cdot,i}^{(t)}\| + 2nr \right)^2 \\ &\leq \frac{1}{2 \sum_{i,j} \widehat{a}_{ij}^{(t)}} \cdot \sum_{i \notin A} (2\epsilon a_{\min}/10wn)^2 \\ &\leq \frac{1}{2 \sum_{i,j} \widehat{a}_{ij}^{(t)}} \cdot n \cdot (\epsilon a_{\min}/5wn)^2 \leq \frac{\epsilon^2}{25n} \cdot a_{\min} \leq \frac{\epsilon^2}{25n} \cdot f(\mathbf{x}^{(t)}) \end{aligned}$$

where the second inequality follows because, for every j , $a_{ij}^{(t)} \leq \widehat{a}_{ij}^{(t)} + r$ and $a_{ji}^{(t)} \leq \widehat{a}_{ji}^{(t)} + r$, and the third inequality follows because $\|\widehat{a}_{i,\cdot}^{(t)}\| + \|\widehat{a}_{\cdot,i}^{(t)}\| < \epsilon a_{\min}/10wn$ and $nr < \epsilon a_{\min}/10wn$.

Next, we bound the sum over $i \in A \setminus (BUC)$. Recall $\widehat{M}_i = \max\{\|\widehat{a}_{i,\cdot}^{(t)}\|, \|\widehat{a}_{\cdot,i}^{(t)}\|\}$ and $\widehat{m}_i = \min\{\|\widehat{a}_{i,\cdot}^{(t)}\|, \|\widehat{a}_{\cdot,i}^{(t)}\|\}$. Put $M_i = \max\{\|a_{i,\cdot}^{(t)}\|, \|a_{\cdot,i}^{(t)}\|\}$ and $m_i = \min\{\|a_{i,\cdot}^{(t)}\|, \|a_{\cdot,i}^{(t)}\|\}$. Let $k = \arg \max_{i \in A \setminus (BUC)} (M_i - m_i)^2$. We

Algorithm 1 RandomBalance(A, ϵ)

Input: Matrix $A \in \mathbb{R}^{n \times n}, \epsilon$
Output: An ϵ -balanced matrix

```

1:  $r = a_{\min} \cdot (\epsilon/wn)^{10}$ 
2: Let  $\hat{a}_{ij}^{(1)} = a_{ij}^{(1)}$  for all  $i$  and  $j$ 
3: Let  $\|\hat{a}_{i,\cdot}^{(t)}\| = \|a_{i,\cdot}^{(t)}\|$  and  $\|\hat{a}_{\cdot,i}^{(t)}\| = \|a_{\cdot,i}^{(t)}\|$  for all  $i$ 
4: for  $t = 1$  to  $O(\epsilon^{-2} \ln w)$  do
5:   Pick  $i$  randomly with probability  $p_i = \frac{\|\hat{a}_{i,\cdot}^{(t)}\| + \|\hat{a}_{\cdot,i}^{(t)}\|}{2 \sum_{i,j} \hat{a}_{ij}^{(t)}}$ 
6:   if  $\|\hat{a}_{i,\cdot}^{(t)}\| + \|\hat{a}_{\cdot,i}^{(t)}\| \geq \epsilon a_{\min}/10wn$  then
7:      $\hat{M}_i = \max\{\|\hat{a}_{i,\cdot}^{(t)}\|, \|\hat{a}_{\cdot,i}^{(t)}\|\}$ ,  $\hat{m}_i = \min\{\|\hat{a}_{i,\cdot}^{(t)}\|, \|\hat{a}_{\cdot,i}^{(t)}\|\}$ 
8:     if  $\hat{m}_i = 0$  or  $\hat{M}_i/\hat{m}_i \geq 1 + \epsilon/n$  then
9:       if  $\hat{m}_i \neq 0$  then  $\alpha = \frac{1}{2} \ln(\|\hat{a}_{\cdot,i}^{(t)}\|/\|\hat{a}_{i,\cdot}^{(t)}\|)$ 
10:      else if  $\hat{m}_i = \|\hat{a}_{i,\cdot}^{(t)}\| = 0$  then  $\alpha = \frac{1}{2} \ln(nr/\|\hat{a}_{i,\cdot}^{(t)}\|)$ 
11:      else if  $\hat{m}_i = \|\hat{a}_{\cdot,i}^{(t)}\| = 0$  then  $\alpha = \frac{1}{2} \ln(\|\hat{a}_{\cdot,i}^{(t)}\|/nr)$ 
12:      end if
13:      Let  $\hat{\mathbf{x}}^{(t+1)} \leftarrow \hat{\mathbf{x}}^{(t)} + \alpha \mathbf{e}_i$  (truncated to  $O(\ln(wn/\epsilon))$  bits of precision)
14:      for  $j = 1$  to  $n$  do
15:        if  $j$  is a neighbor of  $i$  then
16:           $\hat{a}_{ij}^{(t+1)} \leftarrow a_{ij} e^{\hat{x}_i^{(t+1)} - \hat{x}_j^{(t+1)}}$  and  $\hat{a}_{ji}^{(t+1)} \leftarrow a_{ji} e^{\hat{x}_j^{(t+1)} - \hat{x}_i^{(t+1)}}$ , (truncated to  $O(\ln(wn/\epsilon))$ 
bits)
17:           $\|\hat{a}_{j,\cdot}^{(t+1)}\| = \|\hat{a}_{j,\cdot}^{(t)}\| - \hat{a}_{ji}^{(t)} + \hat{a}_{ji}^{(t+1)}$  and  $\|\hat{a}_{\cdot,j}^{(t+1)}\| = \|\hat{a}_{\cdot,j}^{(t)}\| - \hat{a}_{ij}^{(t)} + \hat{a}_{ij}^{(t+1)}$ 
18:          end if
19:        end for
20:         $\|\hat{a}_{i,\cdot}^{(t+1)}\| = \sum_{j=1}^n \hat{a}_{ij}^{(t+1)}$  and  $\|\hat{a}_{\cdot,i}^{(t+1)}\| = \sum_{j=1}^n \hat{a}_{ji}^{(t+1)}$ 
21:      end if
22:    end if
23:  end for
24:  return the resulting matrix

```

have

$$\begin{aligned}
\sum_{i \in A \setminus (BUC)} p_i \cdot \left(\sqrt{\|a_{i,\cdot}^{(t)}\|} - \sqrt{\|a_{\cdot,i}^{(t)}\|} \right)^2 &\leq \frac{1}{2 \sum_{i,j} \hat{a}_{ij}^{(t)}} \cdot \sum_{i \in A \setminus (BUC)} \left(\|a_{i,\cdot}^{(t)}\| - \|a_{\cdot,i}^{(t)}\| \right)^2 \\
&\leq \frac{1}{2 \sum_{i,j} \hat{a}_{ij}^{(t)}} \cdot \sum_{i \in A \setminus (BUC)} (M_i - m_i)^2 \\
&\leq \frac{1}{2 \sum_{i,j} \hat{a}_{ij}^{(t)}} \cdot n \cdot m_k^2 \left(\frac{M_k}{m_k} - 1 \right)^2.
\end{aligned}$$

To bound the last quantity, we prove an upper bound on $\frac{\hat{M}_k}{m_k}$ using the fact that $\frac{\hat{M}_k}{m_k} < 1 + \frac{\epsilon}{n}$. As $k \in A$, we have $\hat{M}_k + \hat{m}_k = \|\hat{a}_{k,\cdot}^{(t)}\| + \|\hat{a}_{\cdot,k}^{(t)}\| \geq \frac{\epsilon a_{\min}}{10wn}$. Thus, $\hat{M}_k \geq \frac{\epsilon a_{\min}}{20wn}$. Combining this with $\frac{\hat{M}_k}{m_k} < 1 + \frac{\epsilon}{n}$ implies

that $\hat{m}_k > \frac{1}{2}\hat{M}_k \geq \frac{\epsilon a_{\min}}{40wn}$. Hence,

$$\frac{M_k}{m_k} \leq \frac{M_k}{\hat{m}_k} \leq \frac{\hat{M}_k + nr}{\hat{m}_k} \leq \frac{\hat{M}_k}{\hat{m}_k} + \frac{nr}{\epsilon a_{\min}/40wn} = \frac{\hat{M}_k}{\hat{m}_k} + 40n \cdot \left(\frac{\epsilon}{wn}\right)^9 \leq \frac{\hat{M}_k}{\hat{m}_k} + \frac{40\epsilon^9}{n^8} \leq 1 + \frac{2\epsilon}{n}.$$

(Notice that $w \geq 1$.) Using the upper bound on $\frac{M_k}{m_k}$, we obtain

$$\begin{aligned} \sum_{i \in A \setminus (B \cup C)} p_i \cdot \left(\sqrt{\|a_{i,\cdot}^{(t)}\|} - \sqrt{\|\hat{a}_{i,\cdot}^{(t)}\|} \right)^2 &\leq \frac{1}{2 \sum_{i,j} \hat{a}_{ij}^{(t)}} \cdot n \cdot m_k^2 \left(\frac{M_k}{m_k} - 1 \right)^2 \\ &\leq \frac{1}{2 \sum_{i,j} \hat{a}_{ij}^{(t)}} \cdot n \cdot m_k^2 \left(\frac{2\epsilon}{n} \right)^2 \\ &\leq \frac{2\epsilon^2}{n} \cdot m_k \leq \frac{\epsilon^2}{n} \cdot f(\mathbf{x}^{(t)}), \end{aligned}$$

where the penultimate inequality uses the fact that $m_k \leq \hat{m}_k + nr \leq \hat{m}_k + \frac{\epsilon a_{\min}}{40wn} < 2\hat{m}_k \leq \hat{M}_k + \hat{m}_k \leq \sum_{i,j} \hat{a}_{ij}^{(t)}$. \square

Proof of Lemma 5. We will assume that $\epsilon < \frac{1}{10}$. We first consider the case that $i \in A \cap B$ (notice that $B \cap C = \emptyset$). The update using $O(\ln(wn/\epsilon))$ bits of precision gives $\hat{x}_i^{(t)} + \alpha - (\epsilon/nw)^{10} \leq \hat{x}_i^{(t+1)} \leq \hat{x}_i^{(t)} + \alpha$, so

$$\sqrt{\frac{\|\hat{a}_{i,\cdot}^{(t)}\|}{\|\hat{a}_{i,\cdot}^{(t+1)}\|}} \cdot e^{\hat{x}_i^{(t)} - (\epsilon/nw)^{10}} \leq e^{\hat{x}_i^{(t+1)}} \leq \sqrt{\frac{\|\hat{a}_{i,\cdot}^{(t)}\|}{\|\hat{a}_{i,\cdot}^{(t+1)}\|}} \cdot e^{\hat{x}_i^{(t)}}.$$

Therefore,

$$\|a_{i,\cdot}^{(t+1)}\| = \sum_{j=1}^n a_{ij} e^{\hat{x}_i^{(t+1)} - \hat{x}_j^{(t+1)}} \leq \sqrt{\frac{\|\hat{a}_{i,\cdot}^{(t)}\|}{\|\hat{a}_{i,\cdot}^{(t+1)}\|}} \cdot \sum_{j=1}^n a_{ij} e^{\hat{x}_i^{(t)} - \hat{x}_j^{(t)}} = \sqrt{\frac{\|\hat{a}_{i,\cdot}^{(t)}\|}{\|\hat{a}_{i,\cdot}^{(t+1)}\|}} \cdot \|a_{i,\cdot}^{(t)}\| \sqrt{\frac{\|\hat{a}_{i,\cdot}^{(t)}\|}{\|\hat{a}_{i,\cdot}^{(t+1)}\|}},$$

and

$$\|a_{i,\cdot}^{(t+1)}\| = \sum_{j=1}^n a_{ji} e^{\hat{x}_j^{(t+1)} - \hat{x}_i^{(t+1)}} \leq e^{(\epsilon/nw)^{10}} \cdot \sqrt{\frac{\|\hat{a}_{i,\cdot}^{(t)}\|}{\|\hat{a}_{i,\cdot}^{(t+1)}\|}} \cdot \sum_{j=1}^n a_{ji} e^{\hat{x}_j^{(t)} - \hat{x}_i^{(t)}} \leq (1 + 2(\epsilon/nw)^{10}) \cdot \sqrt{\frac{\|\hat{a}_{i,\cdot}^{(t)}\|}{\|\hat{a}_{i,\cdot}^{(t+1)}\|}} \cdot \|a_{i,\cdot}^{(t)}\|.$$

We used the fact that $e^x \leq 1 + 2x$ for $x \leq \frac{1}{2}$. We will now use the notation \hat{M}_i , \hat{m}_i , M_i , and m_i (the reader can recall the definitions from the proof of Lemma 4). We also put $\delta = 2(\epsilon/nw)^{10}$, and $\sigma = \frac{\hat{M}_i/\hat{m}_i}{M_i/m_i}$. Thus, decrease of function $f(\cdot)$ due to balancing i is $f(\hat{\mathbf{x}}^{(t)}) - f(\hat{\mathbf{x}}^{(t+1)}) = M_i + m_i - \|a_{i,\cdot}^{(t+1)}\| - \|a_{i,\cdot}^{(t+1)}\| \geq M_i + m_i - (1 + \delta) \left(M_i \sqrt{\hat{m}_i/\hat{M}_i} + m_i \sqrt{\hat{M}_i/\hat{m}_i} \right) = M_i + m_i - (1 + \delta) \left(\sqrt{1/\sigma} + \sqrt{\sigma} \right) \cdot \sqrt{M_i m_i} = (\sqrt{M_i} - \sqrt{m_i})^2 - ((1 + \delta)/\sqrt{\sigma} + (1 + \delta)\sqrt{\sigma} - 2) \cdot \sqrt{M_i m_i}$. We now consider three cases, and in each case show that

$$((1 + \delta)/\sqrt{\sigma} + (1 + \delta)\sqrt{\sigma} - 2) \cdot \sqrt{M_i m_i} \leq \frac{9}{10} \cdot \left(\sqrt{M_i} - \sqrt{m_i} \right)^2.$$

case (i): $1 \leq \sigma < 1 + \frac{\epsilon^4}{n^2}$. We first note that $M_i \geq \hat{M}_i \geq \frac{\hat{M}_i + \hat{m}_i}{2} > \frac{\epsilon a_{\min}}{20wn}$. Also, $m_i \leq \hat{m}_i + nr$, so $\frac{m_i}{M_i} \leq \frac{\hat{m}_i + nr}{M_i} \leq \frac{1}{1+\epsilon/n} + \frac{nr}{M_i} \leq 1 - \frac{\epsilon}{2n}$. Since $\epsilon < \frac{1}{10}$, we have

$$\begin{aligned} \frac{9}{10} \left(1 - \sqrt{\frac{m_i}{M_i}}\right)^2 &\geq \frac{9}{10} \left(1 - \sqrt{1 - \frac{\epsilon}{2n}}\right)^2 \\ &\geq \frac{4\epsilon^4}{n^2} \\ &\geq \left((1+\delta) + (1+\delta) \cdot \left(1 + \frac{\epsilon^4}{n^2}\right) - 2\right) \\ &\geq \left(\frac{1+\delta}{\sqrt{\sigma}} + (1+\delta)\sqrt{\sigma} - 2\right) \cdot \sqrt{\frac{m_i}{M_i}}, \end{aligned}$$

where the third inequality holds by definition of δ , and the last inequality holds because $m_i/M_i \leq 1$ and $\sigma \in [1, 1 + \epsilon^4/n^2]$. By multiplying both sides of the inequality by M_i we obtain the desired bound.

case (ii): $\sigma < 1$. We first prove a lower bound on the value of σ , as follows: $\frac{\hat{M}_i}{\hat{m}_i} \geq \frac{\hat{M}_i}{m_i} \geq \frac{M_i - nr}{m_i} \geq \frac{M_i}{m_i} \left(1 - \frac{nr}{M_i}\right) \geq \frac{M_i}{m_i} \cdot \left(1 - \frac{20\epsilon^9}{n^8}\right)$, and thus $\sigma \geq 1 - \frac{20\epsilon^9}{n^8}$. So we have

$$\begin{aligned} \left(\frac{1+\delta}{\sqrt{\sigma}} + (1+\delta)\sqrt{\sigma} - 2\right) \cdot \sqrt{\frac{m_i}{M_i}} &\leq \frac{1+\delta}{\sqrt{1 - \frac{20\epsilon^9}{n^8}}} + (1+\delta) - 2 \\ &\leq (1+\delta) \left(1 + \frac{20\epsilon^9}{n^8}\right) + (1+\delta) - 2 \\ &\leq \frac{24\epsilon^9}{n^8} < \frac{4\epsilon^4}{n^2} \leq \frac{9}{10} \cdot \left(1 - \sqrt{\frac{m_i}{M_i}}\right)^2, \end{aligned}$$

proving the desired inequality in this case. The first inequality holds because $\frac{m_i}{M_i} \leq 1$ and $1 - \frac{20\epsilon^9}{n^8} \leq \sigma \leq 1$.

case (iii): $\sigma > 1 + \frac{\epsilon^4}{n^2}$. The idea is to show that M_i/m_i is large so the desired inequality follows. We know that $\frac{\sigma M_i}{m_i} = \frac{\hat{M}_i}{\hat{m}_i} \leq \frac{M_i}{m_i}$ and therefore $\hat{m}_i \leq \frac{m_i}{\sigma}$. On the other hand, $\hat{m}_i \geq m_i - nr$, so $m_i \leq \frac{nr}{1-1/\sigma}$. Clearly, $1/\sigma < 1 - \frac{\epsilon^4}{2n^2}$, so $m_i < \frac{nr}{\epsilon^4/2n^2}$. Also, $M_i \geq \epsilon a_{\min}/20wn$. Therefore, $\frac{M_i}{m_i} \geq \frac{\epsilon a_{\min}/20wn}{2n^3r/\epsilon^4} \geq \frac{n^6}{40\epsilon^5}$. Next, notice that since $\hat{m}_i > 0$ it must be that $\hat{m}_i \geq r$. Therefore, $m_i \leq \hat{m}_i + nr \leq 2n\hat{m}_i$. This implies that $\frac{\hat{M}_i}{\hat{m}_i} \leq \frac{M_i}{m_i} \leq 2n \cdot \frac{M_i}{m_i}$, so $\sigma \leq 2n$. Finally,

$$\left(\frac{1+\delta}{\sqrt{\sigma}} + (1+\delta)\sqrt{\sigma} - 2\right) \leq (1+\delta) \cdot \sqrt{2n} \leq \frac{1}{10} \cdot \sqrt{\frac{M_i}{m_i}},$$

with room to spare (using the lower bound on $\frac{M_i}{m_i}$). Multiplying both sides by $\sqrt{\frac{M_i}{m_i}}$ gives

$$\left(\frac{1+\delta}{\sqrt{\sigma}} + (1+\delta)\sqrt{\sigma} - 2\right) \cdot \sqrt{\frac{M_i}{m_i}} \leq \frac{1}{10} \frac{M_i}{m_i} \leq \frac{9}{10} \left(\sqrt{\frac{M_i}{m_i}} - 1\right)^2,$$

with more room to spare. This completes proof of the case $i \in A \cap B$.

We now move on to the case $i \in A \cap C$, so $\hat{M}_i + \hat{m}_i \geq \frac{\epsilon a_{\min}}{10wn}$ and $\hat{m}_i = 0$. In the algorithm, $\alpha = \frac{1}{2} \ln(nr/\|\hat{a}_{i,\cdot}^{(t)}\|)$ or $\alpha = \frac{1}{2} \ln(\|\hat{a}_{\cdot,i}^{(t)}\|/nr)$. The idea is that we therefore replace \hat{m}_i (which is 0) by nr in some of the equations. In particular, $f(\hat{\mathbf{x}}^{(t)}) - f(\hat{\mathbf{x}}^{(t+1)}) \geq M_i + m_i - (1 + \delta) \left(M_i \sqrt{\frac{nr}{\hat{M}_i}} + m_i \sqrt{\frac{\hat{M}_i}{nr}} \right)$. Note that since $\hat{m}_i = 0$ then $m_i \leq nr$. Therefore, $\frac{\hat{M}_i}{nr} \leq \frac{M_i}{nr} \leq \frac{M_i}{m_i}$. On the other hand, since $i \in A$, $\hat{M}_i \geq \epsilon a_{\min}/20wn$, so $\frac{\hat{M}_i}{nr} \geq \frac{\epsilon a_{\min}/20wn}{n(\epsilon/wn)^{10} a_{\min}} \geq \frac{n^8}{20\epsilon^9}$. Thus we get

$$\begin{aligned}
f(\hat{\mathbf{x}}^{(t)}) - f(\hat{\mathbf{x}}^{(t+1)}) &\geq M_i + m_i - (1 + \delta) \left(M_i \sqrt{\frac{nr}{\hat{M}_i}} + m_i \sqrt{\frac{\hat{M}_i}{nr}} \right) \\
&\geq M_i + m_i - (1 + \delta) \left(M_i \sqrt{\frac{20\epsilon^9}{n^8}} + m_i \sqrt{\frac{M_i}{m_i}} \right) \\
&\geq M_i + m_i - 2(1 + \delta) M_i \sqrt{\frac{20\epsilon^9}{n^8}} \\
&\geq M_i \left(1 - \frac{20\epsilon^4}{n^4} \right) \geq \frac{1}{10} M_i \geq \frac{1}{10} (\sqrt{M_i} - \sqrt{m_i})^2,
\end{aligned}$$

where the third inequality holds because $m_i \sqrt{\frac{M_i}{m_i}} = M_i \sqrt{\frac{m_i}{M_i}} \leq M_i \sqrt{\frac{nr}{\hat{M}_i}}$. □

References

- [1] EISPACK implementation. <http://www.netlib.org/eispack/balanc.f>.
- [2] T.-Y. Chen. Balancing sparse matrices for computing eigenvalues. Master's thesis, UC Berkeley, May 1998.
- [3] B. C. Eaves, A. J. Hoffman, U. G. Rothblum, and H. Schneider. Line-sum-symmetric scalings of square nonnegative matrices. In *Mathematical Programming Essays in Honor of George B. Dantzig Part II*, pages 124–141. Springer, 1985.
- [4] J. Grad. Matrix balancing. *The Computer Journal*, 14(3):280–284, 1971.
- [5] D. J. Hartfiel. Concerning diagonal similarity of irreducible matrices. In *Proceedings of the American Mathematical Society*, pages 419–425, 1971.
- [6] B. Kalantari, L. Khachiyan, and A. Shokoufandeh. On the complexity of matrix balancing. *SIAM Journal on Matrix Analysis and Applications*, 118(2):450–463, 1997.
- [7] D. Kressner. *Numerical methods for general and structured eigenvalue problems*. Princeton University Press, 2005.
- [8] E. E. Osborne. On pre-conditioning of matrices. *Journal of the ACM (JACM)*, 7(4):338–345, 1960.
- [9] B. N. Parlett and C. Reinsch. Balancing a matrix for calculation of eigenvalues and eigenvectors. *Numerische Mathematik*, 13(4):293–304, 1969.
- [10] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes: The Art of Scientific Computing, 3rd Edition*. Cambridge University Press, 2007.

- [11] H. Schneider and M. H. Schneider. Max-balancing weighted directed graphs and matrix scaling. *Mathematics of Operations Research*, 16(1):208–222, February 1991.
- [12] L. J. Schulman and A. Sinclair. Analysis of a classical matrix preconditioning algorithm. In *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing*, pages 831–840, 2015.
- [13] L. N. Trefethen and M. Embree. *Spectra and pseudospectra: The behavior of nonnormal matrices and operators*. Springer, 2005.
- [14] N. E. Young, R. E. Tarjan, and J. B. Orlin. Faster parametric shortest path and minimum-balance algorithms. *Networks*, 21(2):205–221, 1991.