



Improved Quantum Lower and Upper Bounds for Matrix Scaling

Sander Gribling, Harold Nieuwboer

► To cite this version:

Sander Gribling, Harold Nieuwboer. Improved Quantum Lower and Upper Bounds for Matrix Scaling. 39th International Symposium on Theoretical Aspects of Computer Science (STACS 2022), Laboratoire d'Informatique & Systèmes (UMR 7020), Institut des sciences de l'information et de leurs interactions (INS2I), Université d'Aix-Marseille (AMU), Mar 2022, Marseille, France. pp.35:1-35:23, 10.4230/LIPIcs.STACS.2022.35 . hal-03673451

HAL Id: hal-03673451

<https://u-paris.hal.science/hal-03673451>

Submitted on 20 May 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Improved Quantum Lower and Upper Bounds for Matrix Scaling

Sander Gribling ✉

IRIF, Université de Paris, CNRS, France

Harold Nieuwboer ✉

Korteweg-de Vries Institute for Mathematics and QuSoft,
University of Amsterdam, The Netherlands

Abstract

Matrix scaling is a simple to state, yet widely applicable linear-algebraic problem: the goal is to scale the rows and columns of a given non-negative matrix such that the rescaled matrix has prescribed row and column sums. Motivated by recent results on first-order quantum algorithms for matrix scaling, we investigate the possibilities for quantum speedups for classical second-order algorithms, which comprise the state-of-the-art in the classical setting.

We first show that there can be essentially no quantum speedup in terms of the input size in the high-precision regime: any quantum algorithm that solves the matrix scaling problem for $n \times n$ matrices with at most m non-zero entries and with ℓ_2 -error $\varepsilon = \tilde{\Theta}(1/m)$ must make $\tilde{\Omega}(m)$ queries to the matrix, even when the success probability is exponentially small in n . Additionally, we show that for $\varepsilon \in [1/n, 1/2]$, any quantum algorithm capable of producing $\frac{\varepsilon}{100}$ - ℓ_1 -approximations of the row-sum vector of a (dense) normalized matrix uses $\Omega(n/\varepsilon)$ queries, and that there exists a constant $\varepsilon_0 > 0$ for which this problem takes $\Omega(n^{1.5})$ queries.

To complement these results we give improved quantum algorithms in the low-precision regime: with quantum graph sparsification and amplitude estimation, a box-constrained Newton method can be sped up in the large- ε regime, and outperforms previous quantum algorithms. For entrywise-positive matrices, we find an ε - ℓ_1 -scaling in time $\tilde{O}(n^{1.5}/\varepsilon^2)$, whereas the best previously known bounds were $\tilde{O}(n^2 \text{polylog}(1/\varepsilon))$ (classical) and $\tilde{O}(n^{1.5}/\varepsilon^3)$ (quantum).

2012 ACM Subject Classification Theory of computation \rightarrow Design and analysis of algorithms; Theory of computation \rightarrow Quantum computation theory

Keywords and phrases Matrix scaling, quantum algorithms, lower bounds

Digital Object Identifier 10.4230/LIPIcs.STACS.2022.35

Related Version *Full Version:* <https://arxiv.org/abs/2109.15282v1> [17]

Funding *Sander Gribling:* Partially supported by SIRTEQ-grant QuIPP.

Harold Nieuwboer: Supported by NWO grant OCENW.KLEIN.267.

Acknowledgements We thank Joran van Apeldoorn, Michael Walter and Ronald de Wolf for interesting and helpful discussions, and the latter two for giving feedback on a first version of this paper. Moreover, we thank Ronald de Wolf for pointing us to [26], which allows for an exponentially small success probability in Theorem 2.1.

1 Introduction

The matrix scaling problem asks to scale each row and column of a given matrix $\mathbf{A} \in [0, 1]^{n \times n}$ by a positive number in such a way that the resulting matrix has marginals (i.e., row- and column-sums) that are close to some prescribed marginals. For example, one could ask to scale the matrix in such a way that it becomes doubly stochastic.



© Sander Gribling and Harold Nieuwboer;

licensed under Creative Commons License CC-BY 4.0

39th International Symposium on Theoretical Aspects of Computer Science (STACS 2022).

Editors: Petra Berenbrink and Benjamin Monmege; Article No. 35; pp. 35:1–35:23

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



Matrix scaling has applications in a wide variety of areas including numerical linear algebra [4], optimal transport in machine learning [13], statistics [23, 14, 9, 8], and also in more theoretical settings, e.g. for approximating the permanent [28]. For a survey, we refer the reader to [19]. Furthermore, the matrix scaling problem is a special (commutative) instance of a more general (non-commutative) class of problems, which includes operator and tensor scaling; these problems have many more applications and are a topic of much recent interest [16, 10].

Formally, the matrix scaling problem is defined for the ℓ_p -norm as follows. Given a matrix $\mathbf{A} \in [0, 1]^{n \times n}$ with at most m non-zero entries, entrywise-positive target marginals $\mathbf{r}, \mathbf{c} \in \mathbb{R}^n$ with $\|\mathbf{r}\|_1 = 1 = \|\mathbf{c}\|_1$, and a parameter $\varepsilon \geq 0$, find vectors $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$ such that the (rescaled) matrix $\mathbf{A}(\mathbf{x}, \mathbf{y}) := (A_{ij}e^{x_i+y_j})_{i,j \in [n]}$ satisfies

$$\|\mathbf{r}(\mathbf{A}(\mathbf{x}, \mathbf{y})) - \mathbf{r}\|_p \leq \varepsilon, \quad \|\mathbf{c}(\mathbf{A}(\mathbf{x}, \mathbf{y})) - \mathbf{c}\|_p \leq \varepsilon. \quad (1.1)$$

Here $\mathbf{r}(\mathbf{A}(\mathbf{x}, \mathbf{y})) = (\sum_{j=1}^n A_{ij}e^{x_i+y_j})_{i \in [n]}$ is the vector of row-marginals of the matrix $\mathbf{A}(\mathbf{x}, \mathbf{y})$ and similarly $\mathbf{c}(\mathbf{A}(\mathbf{x}, \mathbf{y})) = (\sum_{i=1}^n A_{ij}e^{x_i+y_j})_{j \in [n]}$ is the vector of column-marginals. We refer to \mathbf{x} and \mathbf{y} as the scaling vectors, whereas e^{x_i} and e^{y_j} are called scaling factors. A common choice of target marginals is $(\mathbf{r}, \mathbf{c}) = (\frac{1}{n}, \frac{1}{n})$, i.e., every row and column sum target is $1/n$, and we refer to these as the *uniform* target marginals. As is standard in the matrix scaling literature, we will henceforth assume that \mathbf{A} is *asymptotically* (\mathbf{r}, \mathbf{c}) -scalable: for every $\varepsilon > 0$, there exist \mathbf{x}, \mathbf{y} such that $\mathbf{A}(\mathbf{x}, \mathbf{y})$ satisfies Equation (1.1). This depends only on the support of \mathbf{A} [30, Thm. 3], and is the case if and only if (\mathbf{r}, \mathbf{c}) is in the convex hull of the points $(\mathbf{e}_i, \mathbf{e}_j) \in \mathbb{R}^{2n}$ such that $A_{ij} > 0$, where the \mathbf{e}_i are the standard basis vectors for \mathbb{R}^n . We will also always assume the smallest non-zero entry of each of \mathbf{A} , \mathbf{r} and \mathbf{c} is at least $1/\text{poly}(n)$.

Many classical algorithms for the matrix scaling problem can be viewed from the perspective of convex optimization. For example, one can solve the matrix scaling problem by minimizing the convex (potential) function

$$f(\mathbf{x}, \mathbf{y}) = \sum_{i,j=1}^n A_{ij}e^{x_i+y_j} - \langle \mathbf{r}, \mathbf{x} \rangle - \langle \mathbf{c}, \mathbf{y} \rangle, \quad (1.2)$$

where $\langle \cdot, \cdot \rangle$ denotes the standard inner product on \mathbb{R}^n . The popular and practical Sinkhorn algorithm [31] – which alternates between rescaling the rows and columns to the desired marginals – can be viewed as a (block-)coordinate descent algorithm on f , i.e., a first-order method. Given its simplicity, it is no wonder that it has been rediscovered in many settings, and is known by many names, such as the RAS algorithm, iterative proportional fitting, or raking.

It is known that the iterates in the Sinkhorn algorithm converge to a (\mathbf{r}, \mathbf{c}) -scaled matrix whenever \mathbf{A} is asymptotically (\mathbf{r}, \mathbf{c}) -scalable. The convergence rate of Sinkhorn's algorithm is known in various settings, and we give a brief overview of the (classical) time complexity of finding an ε - ℓ_1 -scaling, noting that a single iteration can be implemented in time $\tilde{O}(m)$. When \mathbf{A} is entrywise positive then one can scale in time $\tilde{O}(n^2/\varepsilon)$ [15]; in the ℓ_2 -setting for uniform target marginals a similar result can be found in [21, 20]. In the general setting where \mathbf{A} has at most $m \leq n^2$ non-zero entries the complexity becomes $\tilde{O}(m/\varepsilon^2)$ (for arbitrary target marginals (\mathbf{r}, \mathbf{c})); a proof may be found in [2] for the entrywise-positive case, [11] for exactly scalable matrices (i.e., where the problem can be solved for $\varepsilon = 0$) and [5] for asymptotically scalable matrices.

While simple, the Sinkhorn algorithm is by no means the fastest when the parameter ε is small. The classical state-of-the-art algorithms are based on second-order methods such as (traditional) interior point methods or so-called *box-constrained Newton methods* [12, 1],

the latter of which we describe in more detail below. We note that these algorithms depend on fast algorithms for graph sparsification and Laplacian system solving, so are rather complicated compared to Sinkhorn's algorithm. The box-constrained Newton methods can find ε - ℓ_1 -scaling vectors in time $\tilde{O}(mR_\infty)$, where the \tilde{O} hides polylogarithmic factors in n and $1/\varepsilon$, and R_∞ is a certain diameter bound (made precise later in the introduction). For entrywise-positive matrices, R_∞ is of size $\tilde{O}(1)$, and in general it is known to be $\tilde{O}(n)$ [1, Lem. 3.3]. Alternatively, the interior-point method of [12] has a time complexity of $\tilde{O}(m^{3/2})$, which is better than the box-constrained Newton method for general inputs, but worse for entrywise-positive matrices.

Recently, a quantum algorithm for matrix scaling was developed based on Sinkhorn's algorithm [5]. It uses quantum approximate counting for computing marginals, and finds ε - ℓ_1 -scaling vectors in time $\tilde{O}(\sqrt{mn}/\varepsilon^4)$ for general matrices or $\tilde{O}(n^{1.5}/\varepsilon^3)$ for entrywise-positive matrices. This improves the dependence on m and n at the cost of a higher dependence on $1/\varepsilon$ when compared to the classical Sinkhorn algorithm (which we recall runs in $\tilde{O}(m/\varepsilon^2)$, or $\tilde{O}(n^2/\varepsilon)$ for entrywise-positive matrices). Furthermore, it was shown that this quantum algorithm is optimal for (sufficiently small) constant ε : there exists an $\varepsilon_0 > 0$ (independent of n) such that every quantum algorithm that ε_0 - ℓ_1 -scales to uniform target marginals with probability at least $2/3$ must make at least $\Omega(\sqrt{mn})$ queries. It was left as an open problem whether one can also obtain quantum speedups (in terms of n or m) using second-order methods. In this work we give improved quantum lower and upper bounds on the complexity of matrix scaling. We first prove a lower bound: we show that every quantum algorithm that solves the matrix scaling problem for small enough ε must make a number of queries proportional to the number of non-zero entries in the matrix, even when the success probability of the algorithm is only assumed to be exponentially small. This shows that one cannot hope to get a quantum algorithm for matrix scaling with a polylogarithmic $1/\varepsilon$ -dependence and sublinear dependence on m . However, this does not rule out that second-order methods can be useful in the quantum setting. Indeed, we give a quantum box-constrained Newton method which has a better $1/\varepsilon$ -dependence than the previously mentioned quantum Sinkhorn algorithm, and in certain settings is strictly better, such as for entrywise-positive instances.

1.1 Lower bounds

As previously mentioned, we show for entrywise-positive instances that a polynomial $1/\varepsilon$ -dependence is necessary for a scaling algorithm whose n -dependence is $n^{2-\gamma}$ for a constant $\gamma > 0$. More precisely, we prove the following theorem (which we extend to an $\tilde{\Omega}(m)$ -lower bound in the general setting of $m \leq n^2$ non-zero entries in Corollary 2.11):

► **Theorem 1.1.** *There exists a constant $C > 0$ such that every matrix scaling algorithm that, with probability $\geq \frac{3}{2} \exp(-n/100)$, finds scaling vectors for entrywise-positive $n \times n$ -matrices with ℓ_2 -error $C/(n^2 \sqrt{\ln n})$ must make at least $\Omega(n^2)$ queries to the matrix. This even holds for uniform targets and matrices with smallest entry $\Omega(1/n^2)$.*

The proof of this lower bound is based on a reduction from deciding whether bit strings have Hamming weight $n/2 + 1$ or $n/2 - 1$. Specifically, given k bit strings $z^1, \dots, z^k \in \{\pm 1\}^n$ for $k = \Theta(n)$, each with Hamming weight $|z^i| = n/2 + a_i$ where $a_i \in \{\pm 1\}$, we show that any matrix scaling algorithm can be used to determine all the a_i . One can show that every quantum algorithm that computes all the a_i 's needs to make $\Omega(nk)$ quantum queries to the bit string z^1, \dots, z^k , even if the algorithm has only exponentially small success probability: to determine a single a_i with success probability at least $2/3$, one needs to make $\Omega(n)$ quantum

queries to the bit string z^i [7, 29, 3], and one can use the strong direct product theorem of Lee and Roland [26] to prove the lower bound for computing all k a_i 's simultaneously. To convert the problem of computing the a_i to an instance of matrix scaling, one constructs a $2k \times n$ matrix \mathbf{A} whose first k rows are (roughly) given by the vectors $1 + z^i/b$ for some $b \geq 2$, and whose last k rows are given by $1 - z^i/b$. For such an \mathbf{A} , the column sums are all $2k$, and the row sums are determined by the a_i . If the matrix \mathbf{A}' obtained by a single Sinkhorn step from \mathbf{A} (i.e., rescaling all the rows) were exactly column scaled, then the *optimal* scaling factors encode the a_i . We show that, if one randomly (independently for each i) permutes the z^i beforehand, this is approximately the case: the column sums of this \mathbf{A}' will be close to the desired column sums with high probability, and hence the first step of Sinkhorn gives approximately optimal scaling factors (which encode the a_i). Then, we give a lower bound on the strong convexity parameter of the potential f , to show that *all* sufficiently precise minimizers of f also encode the a_i . In other words, from sufficiently precise scaling factors, we can recover the a_i , yielding the reduction to matrix scaling, and consequently a lower bound for the matrix scaling problem.

We additionally study the problem of computing an ε - ℓ_1 -approximation of the vector of row sums of an ℓ_1 -normalized $n \times n$ matrix \mathbf{A} . This is a common subroutine for matrix scaling algorithms; for instance, the gradient of the potential function f from (1.2) that we optimize for the upper bound can be determined from the row and column sums by subtracting the desired row and column sums, so the complexity of this subroutine directly relates to the complexity of each iteration in our algorithm. We give the following lower bound for this problem.

► **Theorem 1.2 (Informal).** *For $\varepsilon \in [1/n, 1/2]$ and an ℓ_1 -normalized matrix $\mathbf{A} \in [0, 1]^{n \times n}$, computing an $\frac{\varepsilon}{100}$ - ℓ_1 -approximation of $\mathbf{r}(\mathbf{A})$ takes $\Omega(n/\varepsilon)$ queries to \mathbf{A} . Moreover, there exists a constant $\varepsilon_0 > 0$ such that computing an ε_0 - ℓ_1 -approximation of $\mathbf{r}(\mathbf{A})$ takes $\Omega(n^{1.5})$ queries to \mathbf{A} .*

The first lower bound in the theorem is proven in Theorem 2.12. Its proof is based on a reduction from $\Theta(n)$ independent instances of the majority problem, as for the lower bound for matrix scaling. The second lower bound can be derived from the lower bound for matrix scaling given in [5]: using a constant number of calls to a subroutine that provides constant-precision approximations to the row- and column-sum vectors, one can implement Sinkhorn's algorithm to find a constant-precision ℓ_1 -scaling, which for a small enough constant takes $\Omega(n^{1.5})$ queries. Hence, there exists a constant $\varepsilon_0 > 0$ (independent of n) such that computing an ε_0 - ℓ_1 -approximation of $\mathbf{r}(\mathbf{A})$ takes at least $\Omega(n^{1.5})$ queries to the matrix entries.

1.2 Upper bounds

While the first lower bound (Theorem 1.1) shows that a (quantum) algorithm for matrix scaling cannot have both an $m^{1-\gamma}$ -dependence for $\gamma > 0$ and a polylogarithmic $1/\varepsilon$ -dependence, one can still hope to obtain a second-order $\tilde{O}(\sqrt{mn}/\text{poly}(\varepsilon))$ -time algorithm with a better $1/\varepsilon$ -dependence than the quantum Sinkhorn algorithm of [5] (which we recall is based on quantum approximate counting). We show that one can build on a box-constrained Newton method [12, 1] to obtain a quantum algorithm which achieves this, at the cost of depending quadratically on a certain diameter bound R_∞ ; recall for comparison that the classical box-constrained Newton methods run in time $\tilde{O}(mR_\infty)$. For general matrices, one has the bound $R_\infty = \tilde{O}(n)$ [1, Lem. 3.3]. The performance of the resulting quantum box-constrained Newton method is summarized in the following theorem:

► **Theorem 1.3** (Informal version of Corollaries 3.14 and 3.15). *For asymptotically-scalable matrices $\mathbf{A} \in \mathbb{R}_{\geq 0}^{n \times n}$ with m non-zero entries and target marginals (\mathbf{r}, \mathbf{c}) , one can find (\mathbf{x}, \mathbf{y}) such that $\mathbf{A}(\mathbf{x}, \mathbf{y})$ is $O(\varepsilon)$ - ℓ_1 -scaled to (\mathbf{r}, \mathbf{c}) in quantum time $\tilde{O}(R_\infty^2 \sqrt{mn}/\varepsilon^2)$ where R_∞ is the ℓ_∞ -norm of at least one ε^2 -minimizer of f . When \mathbf{A} is entrywise positive we have $R_\infty = \tilde{O}(1)$, so the algorithm runs in quantum time $\tilde{O}(n^{1.5}/\varepsilon^2)$.*

We emphasize that the diameter bound R_∞ does not need to be provided as an input to the algorithm. Note that for entrywise-positive matrices, the algorithm improves over the quantum Sinkhorn method, which runs in time $\tilde{O}(n^{1.5}/\varepsilon^3)$.

Let us give a sketch of the box-constrained method that we use, see Section 3.1 for details. The algorithm aims to minimize the (highly structured) convex potential function f from Equation (1.2). A natural iterative method for minimizing convex functions f is to minimize in each iteration i the quadratic Taylor expansion $\frac{1}{2}\mathbf{x}^T \nabla^2 f(\mathbf{x}^{(i)}) \mathbf{x} + \mathbf{x}^T \nabla f(\mathbf{x}^{(i)}) + f(\mathbf{x}^{(i)})$ of the function at the current iterate. A box-constrained method constrains the minimization of the quadratic Taylor expansion to those \mathbf{x} that lie in an ℓ_∞ -ball of radius c around the current iterate (hence the name):

$$\mathbf{x}^{(i)} = \underset{\|\mathbf{x} - \mathbf{x}^{(i)}\|_\infty \leq c}{\operatorname{argmin}} \frac{1}{2}\mathbf{x}^T \nabla^2 f(\mathbf{x}^{(i)}) \mathbf{x} + \mathbf{x}^T \nabla f(\mathbf{x}^{(i)}).$$

This is guaranteed to decrease a convex function f whenever it is *second-order robust*, i.e., whenever the Hessian of f at a point is a good multiplicative approximation of the Hessian at every other point in a constant-radius ℓ_∞ -ball. One can show that the steps taken decrease the potential gap by a multiplicative factor which depends on the distance to the minimizer.

One then observes that the function f from Equation (1.2) is second-order robust. Moreover, its Hessian has an exceptionally nice structure: given by

$$\nabla^2 f(\mathbf{x}, \mathbf{y}) = \begin{bmatrix} \operatorname{diag}(\mathbf{r}(\mathbf{A}(\mathbf{x}, \mathbf{y}))) & \mathbf{A}(\mathbf{x}, \mathbf{y}) \\ \mathbf{A}(\mathbf{x}, \mathbf{y})^T & \operatorname{diag}(\mathbf{c}(\mathbf{A}(\mathbf{x}, \mathbf{y}))) \end{bmatrix},$$

it is similar to a *Laplacian* matrix. This means that the key subroutine in this method (approximately) minimizes quadratic forms $\frac{1}{2}\mathbf{z}^T \mathbf{H} \mathbf{z} + \mathbf{z}^T \mathbf{b}$ over ℓ_∞ -balls, where \mathbf{H} is a Laplacian matrix; without the ℓ_∞ -constraint, this amounts to solving the Laplacian system $\mathbf{H} \mathbf{z} = \mathbf{b}$. Such a subroutine can be implemented for the more general class of symmetric diagonally-dominant matrices (with non-positive off-diagonal entries) on a classical computer in (almost) linear time in the number of non-zero entries of \mathbf{H} [12]. For technical reasons, one has to add a regularization term to f , and the regularized potential instead has a symmetric diagonally-dominant Hessian structure. Given the recent quantum algorithm for graph sparsification and Laplacian system solving of Apers and de Wolf [6], one would therefore hope to obtain a quantum speedup for the box-constrained Newton method. We show that one can indeed achieve this by first using the quantum algorithm for graph sparsification, and then using the classical method for the minimization procedure. We note, however, that in order to achieve a quantum speedup in terms of m and n , we incur a polynomial dependence in the time complexity on the precision with which we can approximate \mathbf{H} and \mathbf{b} (as opposed to only a *polylogarithmic* dependence classically). Such a speedup with respect to one parameter (dimension) at the cost of a slowdown with respect to another (precision) is more common in recent quantum algorithms for optimization problems and typically requires a more careful analysis of the impact of approximation errors. Interestingly, for the classical box-constrained Newton method, the minimization subroutine is the bottleneck, whereas in our quantum algorithm, the cost of a single iteration is dominated by the time it takes

to approximate the vector \mathbf{b} . Using quantum approximate counting (carefully) as in [5], one can obtain an additive $\delta \cdot \|\mathbf{A}(\mathbf{x}, \mathbf{y})\|_1$ -approximation of \mathbf{b} in time roughly \sqrt{mn}/δ . To obtain an efficient quantum algorithm we therefore need to control $\|\mathbf{A}(\mathbf{x}, \mathbf{y})\|_1$ throughout the run of the algorithm. We do so efficiently by testing in each iteration whether the 1-norm of $\mathbf{A}(\mathbf{x}, \mathbf{y})$ is too large, if it is, we divide the matrix by 2 (by shifting \mathbf{x} by an appropriate multiple of the all-ones vector), which reduces the potential.

1.3 Open problems

Our lower bound on matrix scaling shows that it is not possible to provide significant quantum speedups for scaling of entrywise-positive matrices in the high-precision scaling regime. However, the best classical upper bound for ε -scaling when no assumptions are made on the support of the matrices is $\tilde{O}(m^{3/2})$, where m is the number of non-zero entries [12] (recall that this hides a polylogarithmic dependence on $1/\varepsilon$). The algorithm that achieves this bound is an interior-point method, rather than a box-constrained Newton method. It is an interesting open problem whether such an algorithm also admits a quantum speedup in terms of m while retaining a polylogarithmic $1/\varepsilon$ -dependence. Note that while the interior-point method relies on fast Laplacian system solvers, it is not enough to merely replace this by a quantum Laplacian system solver, as the dimension of the linear system in question is $m + n$ rather than $\Theta(n)$. More generally, the possibility of obtaining quantum advantages in high-precision regimes for optimization problems is still a topic of ongoing investigation.

A second natural question is whether the lower bounds from Theorem 1.2 for computing an approximation of the row sums are tight. The best upper bound for the row-sum vector approximation that we are aware of is the one we use in our scaling algorithm: we can compute an ε - ℓ_1 -approximation of the row sums in time $\tilde{O}(n^{1.5}/\varepsilon)$. For constant $\varepsilon_0 \geq \varepsilon > 0$ this matches the lower bound $\Omega(n^{1.5})$ (up to log-factors), but for non-constant $\varepsilon > \frac{1}{100n}$ it remains an interesting open problem to close the gap between $\tilde{O}(n^{1.5}/\varepsilon)$ and $\Omega(n/\varepsilon)$.

2 Lower bounds for matrix scaling and marginal approximation

In this section we prove two lower bounds: an $\tilde{\Omega}(m)$ -lower bound for $1/\text{poly}(n)$ - ℓ_2 -scaling $n \times n$ matrices with at most m non-zero entries, and for $\varepsilon \in [1/n, 1/2]$ an $\Omega(n/\varepsilon)$ -lower bound for ε - ℓ_1 -approximation of the row-sum vector of a normalized $n \times n$ matrix (with non-negative entries). The proofs for both lower bounds are based on a reduction from the lower bound given below in Theorem 2.1. In Section 2.1 we construct the associated instances for matrix scaling, and in Section 2.2 we analyze their column marginals after a single iteration of the Sinkhorn algorithm. Afterwards, in Section 2.3 we show that these column marginals are close enough to the target marginals for the reduction to matrix scaling to work, and in Section 2.4 we put the ingredients together, with the main theorem being Theorem 2.10. Finally, in Section 2.5 we prove the lower bound for computing approximations to the row marginals. The lower bound we reduce from is the following:

► **Theorem 2.1.** *Let n be even, $\tau \in [1/n, 1/2]$ such that $n\tau$ is an integer, and let $k \geq 1$ be an integer. Given k binary strings $z^1, \dots, z^k \in \{\pm 1\}^n$, where z^i has Hamming weight $n/2 + a_i\tau n$ for $a_i \in \{-1, 1\}$, computing with probability $\geq \exp(-k/100)$ a string $\tilde{a} \in \{-1, 1\}^k$ that agrees with a in $\geq 99\%$ of the positions requires $\Omega(k/\tau)$ quantum queries.*

Proof. Let $\mathcal{D} = \{z \in \{\pm 1\}^n : |z| = n/2 + \tau n \text{ or } |z| = n/2 - \tau n\}$ and define the partial Boolean function $f: \mathcal{D} \rightarrow \{\pm 1\}$ by $f(z) = 1$ if $|z| = n/2 + \tau n$, and $f(z) = -1$ otherwise. It is known that computing f with probability at least $2/3$ takes $\Theta(1/\tau)$ quantum queries to z [29, Cor. 1.2], i.e., the bounded-error quantum query complexity $Q_{1/3}(f)$ is $\Theta(1/\tau)$.

We now proceed with bounding the query complexity of computing 99% of the entries of $f^{(k)}: \mathcal{D}^k \rightarrow \{\pm 1\}^k$ defined by $f^{(k)}(z^1, \dots, z^k) = (f(z^1), \dots, f(z^k))$. We will make use of the general adversary bound $\text{Adv}^\pm(f)$ [18] which is known to satisfy $\text{Adv}^\pm(f) = \Theta(Q_{1/3}(f))$ [25, Thm. 1.1]. The strong direct product theorem of Lee and Roland [26, Thm. 5.5] says that for every $0 \leq \delta < 1$, $\mu \in [\frac{1+\sqrt{\delta}}{2}, 1]$ and integers k, K , every quantum algorithm that outputs a bit string $\tilde{a} \in \{\pm 1\}^k$, and makes T quantum queries to the bit strings z^1, \dots, z^k with $T \leq \frac{k\delta}{K(1-\delta)} \text{Adv}^\pm(f)$ has the property that \tilde{a} agrees with $f^{(k)}(z^1, \dots, z^k)$ on at least a μ -fraction of the entries with probability at most $\exp(k(\frac{1}{K} - D(\mu \parallel \frac{1+\sqrt{\delta}}{2})))$.¹ Here $D(\mu \parallel \frac{1+\sqrt{\delta}}{2})$ is the Kullback–Leibler divergence between the distributions $(\mu, 1-\mu)$ and $(\frac{1+\sqrt{\delta}}{2}, \frac{1-\sqrt{\delta}}{2})$. For $\mu = 0.99$, $\delta = 0.1$ and $K = 3$, one has $\frac{1}{K} - D(\mu \parallel \frac{1+\sqrt{\delta}}{2}) \approx -0.03 \leq -1/100$. Therefore, the strong direct product theorem shows that computing 99% of the entries of $f^{(k)}(z^1, \dots, z^k) = a$ correctly, with success probability at least $\exp(-k/100)$, takes $\Omega(k \text{Adv}^\pm(f)) = \Omega(k Q_{1/3}(f)) = \Omega(k/\tau)$ quantum queries. ◀

We will use this lower bound with $k = n/2$ and $\tau = 1/n$. The following intuition is useful to keep in mind. For a fixed $b \geq 2$, define the $2k \times n$ matrix \mathbf{A} whose $(2i-1)$ -th row equals $1 + z^i/b$ and whose $(2i)$ -th row equals $1 - z^i/b$. Then \mathbf{A} has the property that the row-marginals encode the Hamming weights of the z^i , and are all very close to n . (This implies that the first row-rescaling step of Sinkhorn’s algorithm encodes the a_i .) Moreover, the column-marginals are exactly uniform. Hence, one may hope that all sufficiently precise scalings of \mathbf{A} to uniform targets have scaling factors that are close to those given by the first row-rescaling step of Sinkhorn’s algorithm (and hence learn most of the a_i).

Below we formalize this approach. We show that if one randomly permutes the coordinates of each z^i (independently over i), then with high probability, all ε -scalings of the resulting matrix \mathbf{A}^σ are close to the first step of Sinkhorn’s algorithm; here we need to choose b sufficiently large ($\sim \sqrt{\ln(n)}$) and ε sufficiently small ($\sim \frac{1}{n^{2b}}$). The section is organized as follows. In Section 2.1 we formally define our matrix scaling instances and we analyse the first row-rescaling step of Sinkhorn’s algorithm. In Section 2.2 we show that after the row-rescaling step, with high probability (over the choice of permutations), the column-marginals are close to uniform. In Sections 2.3 and 2.4 we use the strong convexity of the potential f from Equation (1.2) to show that if the above event holds, then all approximate minimizers of f can be used to solve the counting problem.

2.1 Definition of the scaling instances and analysis of row marginals

Let $n \geq 4$ be even. Let $k = n/2$ and let $z^1, \dots, z^k \in \{\pm 1\}^n$ have Hamming weight $|z^i| = |\{j : z_j^i = 1\}| = n/2 + a_i$ for $a_i \in \{\pm 1\}$. Sample uniformly random permutations $\sigma^1, \dots, \sigma^k \in S_n$ and define w^i by $w_j^i = z_{(\sigma^i)^{-1}(j)}^i$. Let $b \geq 2$ be some number depending on n , and consider the $2k \times n$ matrix \mathbf{A}^σ whose entries are $\mathbf{A}_{2i-1,j}^\sigma = 1 + \frac{w_j^i}{b}$ and $\mathbf{A}_{2i,j}^\sigma = 1 - \frac{w_j^i}{b}$. Then each column sum $c_j(\mathbf{A}^\sigma)$ is $2k$, and the row sums of \mathbf{A}^σ are given by

$$r_{2i-1}(\mathbf{A}^\sigma) = n + \frac{1}{b} \sum_{j=1}^n w_j^i = n + \frac{2}{b} a_i, \quad r_{2i}(\mathbf{A}^\sigma) = n - \frac{2}{b} a_i.$$

¹ In [26] the upper bound on T is stated in terms of $\text{Adv}^*(F)$ where $F = (\delta_{f(x), f(y)})_{x, y \in \mathcal{D}}$ is the Gram matrix of f . For Boolean functions f one has $\text{Adv}^*(F) = \text{Adv}^\pm(f)$ [25, Thm. 3.4].

Let

$$X_{2i-1} = \frac{1}{2k} \cdot \frac{1}{n + \frac{2}{b}a_i} \text{ and } X_{2i} = \frac{1}{2k} \cdot \frac{1}{n - \frac{2}{b}a_i} \quad \text{for all } i \in [k] \quad (2.1)$$

be the row scaling factors obtained from a single Sinkhorn step. We first observe that the difference between $x_{2i-1} := \ln(X_{2i-1})$ and $x_{2i} := \ln(X_{2i})$ permits to recover a_i .

► **Lemma 2.2.** *For the specific row-scaling factors \mathbf{X} for \mathbf{A}^σ given in (2.1), for every $i \in [k]$ it holds that $|\ln(X_{2i-1}/X_{2i})| \geq \frac{4}{nb}$, and $\text{sign}(\ln(X_{2i}/X_{2i-1})) = a_i$.*

Proof. Using $nb > 2$, we have $|\ln(X_{2i-1}/X_{2i})| = \left| \ln \left(\frac{n + \frac{2}{b}}{n - \frac{2}{b}} \right) \right| = \ln \left(\frac{nb+2}{nb-2} \right) \geq \frac{4}{nb}$. ◀

2.2 Concentration of column marginals

We record here an explicit expression for the j -th column marginal of \mathbf{XA}^σ for the \mathbf{X} from (2.1), which follows from straightforward algebraic manipulations.

► **Lemma 2.3.** *We have $c_j(\mathbf{XA}^\sigma) = \frac{1}{2k(n^2-4/b^2)} \left(2kn - \frac{4}{b^2} \sum_{i=1}^k w_j^i a_i \right)$ for $j \in [n]$.*

We now show that with high probability (over the choice of permutations) the column marginals are close to uniform. To do so, we first compute the expectation of $\sum_{i=1}^k w_j^i a_i$ (Corollary 2.5). This quantity allows us to obtain the desired concentration of the column marginals via Hoeffding's inequality (Lemma 2.6).

► **Lemma 2.4.** *Let $I = \{i \in [k] : a_i = 1\}$ and $I^c = [k] \setminus I$. Define random variables W_j, W_j^c by $W_j = \sum_{i \in I} w_j^i$ and $W_j^c = \sum_{i \in I^c} w_j^i$. Then $\mathbb{E}[W_j] = \frac{2|I|}{n}$ and $\mathbb{E}[W_j^c] = -\frac{2|I^c|}{n}$.*

Proof. Observe that each w_j^i is 1 with probability $\frac{1}{2} + \frac{a_i}{n}$ because σ^i is chosen uniformly randomly from S_n , and is -1 with probability $\frac{1}{2} - \frac{a_i}{n}$. Therefore $\mathbb{E}[w_j^i] = \frac{2a_i}{n}$. By linearity of expectation, the result follows. ◀

► **Corollary 2.5.** *We have $\mathbb{E} \left[\sum_{i=1}^k w_j^i a_i \right] = \mathbb{E}[W_j] - \mathbb{E}[W_j^c] = \frac{2(|I|+|I^c|)}{n} = \frac{2k}{n}$.*

► **Lemma 2.6.** *For $t \geq 0$ and $j \in [n]$, with probability at least $1 - 2e^{-t^2/2}$, we have $|c_j(\mathbf{XA}^\sigma) - \frac{1}{n}| = O \left(\frac{t}{b^2 n^2 \sqrt{k}} \right)$.*

Proof. One can verify that $|c_j(\mathbf{XA}^\sigma) - \frac{1}{n}| = \frac{4}{2kn(n^2-4/b^2)b^2} \left| 2k - n \sum_{i=1}^k w_j^i a_i \right|$. For fixed j and distinct $i, i' \in [k]$, w_j^i and $w_j^{i'}$ are independently distributed random variables because σ^i and $\sigma^{i'}$ are independent. Therefore, $V_j := W_j - W_j^c = \sum_{i=1}^k w_j^i a_i$ is a sum of k independent random variables, with each $a_i w_j^i \in [-1, 1]$, and Hoeffding's inequality yields for any $t \geq 0$ that $\Pr[|V_j - \mathbb{E}[V_j]| \geq t \cdot \sqrt{k}] \leq 2 \exp(-t^2/2)$. Assuming that $|V_j - \mathbb{E}[V_j]| \leq t\sqrt{k}$, we have $|2k - n \sum_{i=1}^k a_i w_j^i| = n|\mathbb{E}[V_j] - V_j| \leq nt\sqrt{k}$. With this estimate, we see that

$$\left| c_j(\mathbf{XA}^\sigma) - \frac{1}{n} \right| \leq \frac{4}{2kn(n^2-4/b^2)b^2} \cdot nt\sqrt{k} = \frac{2t}{b^2(n^2-4/b^2)\sqrt{k}}. \quad \blacktriangleleft$$

► **Corollary 2.7.** *For any $t \geq 0$, with probability $\geq 1 - 2ne^{-t^2/2}$, we have $\|c(\mathbf{XA}^\sigma) - \frac{1}{n}\|_2 \leq \frac{2\sqrt{nt}}{b^2(n^2-4/b^2)\sqrt{k}} = O \left(\frac{t}{b^2 n^2} \right)$.*

2.3 Strong convexity properties of the potential

For a λ -strongly convex function f , the set $\{\mathbf{z} : \|\nabla f(\mathbf{z})\|_2 \leq \varepsilon\}$ has a diameter that is bounded by a function of λ (we make this well-known fact precise in Lemma A.4). We show that our potential is strongly convex when viewed as a function from (a suitable subset of) the linear subspace $V = \{(\mathbf{x}, \mathbf{y}) \in \mathbb{R}^n \times \mathbb{R}^n : \langle (\mathbf{x}, \mathbf{y}), (\mathbf{1}_n, -\mathbf{1}_n) \rangle = 0\}$ to \mathbb{R} (note that f is invariant under translation by multiples of $(\mathbf{1}_n, -\mathbf{1}_n)$). We use this to prove the following lemma, which shows that whenever $\nabla f(\mathbf{x}, \mathbf{y})$ is small, (\mathbf{x}, \mathbf{y}) is close to the minimizer of f on V . It is easy to verify that Corollary 2.7 in fact gives an upper bound on the ℓ_2 -norm of the gradient at $(\ln(\mathbf{X}), \mathbf{0})$ (with \mathbf{X} as in (2.1)). This implies that $(\ln(\mathbf{X}), \mathbf{0})$ is close to the minimizer of f on V , and by the triangle inequality, is also close to any other (\mathbf{x}, \mathbf{y}) for which $\|\nabla f(\mathbf{x}, \mathbf{y})\|_2$ is small. The full proof is given in Appendix A.

► **Lemma 2.8.** *Let $f: V \subset \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}$ be the standard potential for the matrix \mathbf{A}^σ , where V is the orthogonal complement of $(\mathbf{1}_n, -\mathbf{1}_n)$. Then for every $(\mathbf{x}, \mathbf{y}) \in V$ and $\delta \in (0, 1)$, if $\|\nabla f(\mathbf{x}, \mathbf{y})\|_2 \leq \frac{\delta}{27ne^2}$, then $\|(\mathbf{x}, \mathbf{y}) - (\mathbf{x}^*, \mathbf{y}^*)\|_\infty \leq \delta$.*

2.4 Concluding the lower bound for matrix scaling

Let $(\bar{\mathbf{x}}, \bar{\mathbf{y}}) \in V$ be the unique vector such that $(\bar{\mathbf{x}}, \bar{\mathbf{y}}) - (\mathbf{x}, \mathbf{y})$ is a multiple of $(\mathbf{1}_n, -\mathbf{1}_n)$, where (\mathbf{x}, \mathbf{y}) are the scaling vectors of the first step of Sinkhorn. By choosing t and b appropriately we obtain, with high probability over the choice of permutations, a bound on the distance between $(\bar{\mathbf{x}}, \bar{\mathbf{y}})$ and the unique scaling vectors $(\mathbf{x}^*, \mathbf{y}^*) \in V$ of an exact scaling of \mathbf{A}^σ . This allows us to conclude that, with high probability, all sufficiently precise scalings of \mathbf{A}^σ encode the Hamming weights a_i .

► **Corollary 2.9.** *There exists a constant $C > 0$ such that for $b = C\sqrt{\ln n}$ the following holds. With probability $\geq 2/3$ (over the choice of σ) we have for the exact scaling vectors $(\mathbf{x}^*, \mathbf{y}^*) \in V$ of \mathbf{A}^σ that $a_i = \text{sign}(x_{2i}^* - x_{2i-1}^*)$ for all i . Furthermore, there exists a constant $C' > 0$ such that for any (x', y') that yield a (C'/n^2b) - ℓ_2 -scaling of \mathbf{A}^σ , a_i can be recovered from x' as $a_i = \text{sign}(x_{2i}' - x_{2i-1}') = \text{sign}(x_{2i}' - x_{2i-1}')$.*

Proof. Applying Corollary 2.7 with $t = 10\sqrt{\ln n}$ shows that with probability at least $2/3$ we have $\|\nabla f(\bar{\mathbf{x}}, \bar{\mathbf{y}})\|_2 = \|\nabla f(\mathbf{x}, \mathbf{y})\|_2 = \frac{t}{b} \frac{2\sqrt{n}}{b(n^2 - 4/b^2)\sqrt{k}}$. Hence, there exists a constant $C > 0$ such that for $b = Ct$ we have $\|\nabla f(\bar{\mathbf{x}}, \bar{\mathbf{y}})\|_2 \leq \frac{1}{nb} \frac{1}{27ne^2}$. Lemma 2.8 then implies that $\|(\bar{\mathbf{x}}, \bar{\mathbf{y}}) - (\mathbf{x}^*, \mathbf{y}^*)\|_\infty \leq \frac{1}{nb}$ and hence $|(x_{2i-1}^* - x_{2i}^*) - (x_{2i-1}' - x_{2i}')| \leq \frac{2}{nb}$. Together with Lemma 2.2 (which shows that $|x_{2i-1} - x_{2i}| \geq \frac{4}{nb}$) this means that $a_i = \text{sign}(x_{2i}^* - x_{2i-1}^*)$. Moreover, $|x_{2i-1}^* - x_{2i}^*| \geq \frac{2}{nb}$.

Now consider approximate scalings of \mathbf{A}^σ . Without loss of generality we may assume that the (x', y') that yield a $(\frac{1}{2nb} \frac{1}{27ne^2})$ - ℓ_2 -scaling of \mathbf{A}^σ belong to V (otherwise we shift it by an appropriate multiple of $(\mathbf{1}_n, -\mathbf{1}_n)$). Then, again due to Lemma 2.8, we obtain that $\|(\mathbf{x}', \mathbf{y}') - (\mathbf{x}^*, \mathbf{y}^*)\|_\infty \leq \frac{1}{2nb} \leq \frac{1}{4}|x_{2i-1}^* - x_{2i}^*|$ and hence $|(x_{2i-1}' - x_{2i}') - (x_{2i-1}^* - x_{2i}^*)| \leq \frac{1}{2}|x_{2i-1}^* - x_{2i}^*|$ which means that $\text{sign}(x_{2i}' - x_{2i-1}') = \text{sign}(x_{2i-1}^* - x_{2i}^*) = a_i$. ◀

► **Theorem 2.10.** *There exists a constant $C > 0$ such that any matrix scaling algorithm that, with probability $\geq \frac{3}{2} \exp(-n/100)$, finds scalings for $n \times n$ -matrices with ℓ_2 -error $C/(n^2\sqrt{\ln n})$ must make at least $\Omega(n^2)$ queries to the matrix. This even holds for uniform targets and entrywise-positive matrices with smallest entry $\Omega(1/n^2)$.*

Proof. We construct a set of hard instances as in Section 2.1. Let $n \geq 4$ be even. Let $k = n/2$ and let $z^1, \dots, z^k \in \{\pm 1\}^n$ have Hamming weight $|z^i| = |\{j : z_j^i = 1\}| = n/2 + a_i$ for $a_i \in \{\pm 1\}$. By Theorem 2.1, finding at least 99% of the a_i 's with probability $\geq \exp(-n/100)$ takes $\Omega(n^2)$ -queries to the z_j^i . One can recover the a_i 's with probability $\geq 2/3$ as follows.

First, sample the $\sigma^1, \dots, \sigma^{n/2}$ uniformly from S_n . A single query to \mathbf{A}^σ takes a single query to some w^i , which takes a single query to z^i . Using Corollary 2.9, there exists a constant $C > 0$ such that, with probability $\geq 2/3$, any scaling of \mathbf{A}^σ with ℓ_2 -error $C/(n^2\sqrt{\ln n})$ recovers all a_i 's. Therefore any matrix scaling algorithm finding such a scaling with probability $\geq \exp(-n/100)$ allows us to find all a_i 's with probability $\geq \exp(-n/100)$. \blacktriangleleft

► **Corollary 2.11.** *There exist constants $C_0, C_1 > 0$ such that every matrix scaling algorithm that, with probability $\geq \exp(-C_0 n / \ln(n))$, finds scalings for $n \times n$ -matrices with at most m non-zero entries and ℓ_2 -error $C_1/(m\sqrt{\ln(m/n)})$ must make at least $\tilde{\Omega}(m)$ queries. This even holds for uniform targets and matrices with smallest non-zero entry $\Omega(1/m)$.*

2.5 Lower bound for computing the row marginals

In Theorem 2.12 we show that computing an ε - ℓ_1 -approximation of the row (or column marginals) of an entrywise-positive $n \times n$ matrix takes $\Omega(n/\varepsilon)$ queries to its entries (for $\varepsilon = \Omega(1/n)$). As a consequence, the same holds for computing an approximation of the gradient of common (convex) potential functions used for matrix scaling – among which is the potential we use in Section 3 – takes as many queries. Although the bound does not imply that testing whether a matrix is ε - ℓ_1 -scaled takes at least $\Omega(n/\varepsilon)$ queries, it gives reasonable evidence that this should be the case. The proof can be found in the full version [17].

► **Theorem 2.12.** *Let $\tau \in [1/n, 1/2]$. Suppose we have a quantum algorithm that, given query access to a positive $n \times n$ matrix \mathbf{A} with row-sums $\mathbf{r} = (r_1, \dots, r_n)$ and column-sums $\mathbf{c} = (1/n, \dots, 1/n)$, outputs (with probability $\geq \exp(-n/100)$) a vector $\tilde{\mathbf{r}} \in \mathbb{R}_+^n$ such that $\|\tilde{\mathbf{r}} - \mathbf{r}\|_1 < \tau/100$. Then this algorithm uses $\Omega(n/\tau)$ queries.*

3 Quantum box-constrained Newton method for matrix scaling

In this section, we show how to obtain a quantum speedup based on the box-constrained Newton method for matrix scaling from [12], with the main result being Theorem 3.13, and its consequences for matrix scaling given in Corollaries 3.14 and 3.15. We first recall some of the concepts that are used in the algorithm, including the definition of second-order robust convex functions, the notion of a k -oracle, and a theorem regarding efficient (classical) implementation of a k -oracle for the class of symmetric diagonally-dominant matrices with non-positive off-diagonal entries. We then show that for a second-order robust function $g: \mathbb{R}^n \rightarrow \mathbb{R}$ and a given $\mathbf{x} \in \mathbb{R}^n$ such that the sublevel set $\{\mathbf{x}' : g(\mathbf{x}') \leq g(\mathbf{x})\}$ is bounded, one can use a k -oracle and approximations to the gradient and Hessian of g to find a vector \mathbf{x}' such that the potential gap $g(\mathbf{x}') - g(\mathbf{x}^*)$ is smaller than $g(\mathbf{x}) - g(\mathbf{x}^*)$ where \mathbf{x}^* is a minimizer of g . This result extends [12, Thm. 3.4] to a setting where one can only obtain rough approximations of the gradient and Hessian of g . We then show that this applies to a regularized version \tilde{f} of the potential f discussed in the introduction; to approximate the Hessian of \tilde{f} , we use a quantum algorithm for graph sparsification, whereas we approximate the gradient of \tilde{f} using quantum approximate summing. One challenge is that the quality of the gradient approximation is directly related to the 1-norm of the matrix $\mathbf{A}(\mathbf{x}, \mathbf{y})$, so we must control this throughout the algorithm, which we achieve by manually shifting \mathbf{x} when the norm becomes too large, and showing that this does not increase the regularized potential under suitable circumstances.

3.1 Minimizing second-order robust convex functions

In what follows we will minimize a convex function (potential) that satisfies a certain regularity condition: its Hessian can be approximated well on an infinity-norm ball.

► **Definition 3.1** ([12, Def. 3.1]). *A convex function $g: \mathbb{R}^n \rightarrow \mathbb{R}$ is called second-order robust with respect to ℓ_∞ if for any $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$ with $\|\mathbf{x} - \mathbf{y}\|_\infty \leq 1$, $\frac{1}{e^2} \nabla^2 g(\mathbf{x}) \preceq \nabla^2 g(\mathbf{y}) \preceq e^2 \nabla^2 g(\mathbf{x})$.*

This implies that the local quadratic approximation to g has a good quality on a small ℓ_∞ -norm ball. It is therefore natural to consider the problem of minimizing a convex quadratic function over an ℓ_∞ -norm ball. We will use the following notion.

► **Definition 3.2** (k -oracle). *An algorithm \mathcal{A} is called a k -oracle for a class of matrices $\mathcal{M} \subseteq \mathbb{R}^{n \times n}$ if for input (\mathbf{H}, \mathbf{b}) with $\mathbf{H} \in \mathcal{M}$, $\mathbf{b} \in \mathbb{R}^n$, it returns a vector $\mathbf{x} \in \mathbb{R}^n$ such that $\|\mathbf{x}\|_\infty \leq k$ and $\frac{1}{2} \mathbf{x}^T \mathbf{H} \mathbf{x} + \langle \mathbf{b}, \mathbf{x} \rangle \leq \frac{1}{2} \cdot \min_{\|\mathbf{z}\|_\infty \leq 1} (\frac{1}{2} \mathbf{z}^T \mathbf{H} \mathbf{z} + \langle \mathbf{b}, \mathbf{z} \rangle)$.*

► **Definition 3.3** (SDD matrix). *A matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$ is called symmetric diagonally-dominant if it is symmetric, and for every $i \in [n]$, one has $A_{ii} \geq \sum_{j \neq i} |A_{ij}|$.*

In [12] it is shown how to efficiently implement an $O(\log(n))$ -oracle for the class of SDD matrices \mathbf{H} whose off-diagonal entries are non-positive. Their algorithm uses an efficient construction of a *vertex sparsifier chain* of \mathbf{H} due to [27, 24].

► **Theorem 3.4** ([12, Thm. 5.11]). *Given a classical description of an SDD matrix $\mathbf{H} \in \mathbb{R}^{n \times n}$ with $\tilde{O}(m)$ non-zero entries, such that $H_{i,j} \leq 0$ for $i \neq j$, and a classical vector $\mathbf{b} \in \mathbb{R}^n$, we can find in time $\tilde{O}(m)$ a vector $\mathbf{x} \in \mathbb{R}^n$ such that $\|\mathbf{x}\|_\infty = O(\log n)$ and*

$$\frac{1}{2} \mathbf{x}^T \mathbf{H} \mathbf{x} + \langle \mathbf{b}, \mathbf{x} \rangle \leq \frac{1}{2} \cdot \min_{\|\mathbf{z}\|_\infty \leq 1} (\frac{1}{2} \mathbf{z}^T \mathbf{H} \mathbf{z} + \langle \mathbf{b}, \mathbf{z} \rangle).$$

A k -oracle \mathcal{A} gives rise to an iterative method for minimizing a second-order robust function g : starting from $\mathbf{x}_0 \in \mathbb{R}^n$, we define a sequence $\mathbf{x}^{(0)}, \mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \dots$ by

$$\mathbf{x}^{(i+1)} = \mathbf{x}^{(i)} + \frac{1}{k} \Delta_i, \quad \Delta_i = \mathcal{A} \left(\frac{e^2}{k^2} \mathbf{H}_i, \frac{1}{k} \mathbf{b}_i \right)$$

where \mathbf{H}_i is an approximate Hessian at $\mathbf{x}^{(i)}$, and \mathbf{b}_i is an approximate gradient at $\mathbf{x}^{(i)}$. The following theorem, which is an adaptation of [12, Thm. 3.4], upper bounds the progress made in each iteration. We defer its proof to Appendix B.

► **Theorem 3.5.** *Let $g: \mathbb{R}^n \rightarrow \mathbb{R}$ be a second-order robust function with respect to ℓ_∞ , let $\mathbf{x} \in \mathbb{R}^n$ be a starting point, and suppose \mathbf{x}^* is a minimizer of g . Assume that we are given*

1. *a vector $\mathbf{b} \in \mathbb{R}^n$ such that $\|\mathbf{b} - \nabla g(\mathbf{x})\|_1 \leq \delta$,*
2. *two SDD matrices \mathbf{H}_m and \mathbf{H}_a with non-positive off-diagonal entries, such that there exists $\delta_a \geq 0$ and symmetric \mathbf{H}'_m and \mathbf{H}'_a satisfying $\nabla^2 g(\mathbf{x}) = \mathbf{H}'_m + \mathbf{H}'_a$ and $\frac{2}{3} \mathbf{H}_m \preceq \mathbf{H}'_m \preceq \frac{4}{3} \mathbf{H}_m$, $\|\mathbf{H}_a - \mathbf{H}'_a\|_1 \leq \delta_a$.*

Let $k = O(\log n)$ be such that there exists a k -oracle \mathcal{A} for the class of SDD-matrices with non-positive off-diagonal entries (cf. Theorem 3.4). Then for $\mathbf{H} = \mathbf{H}_m + \mathbf{H}_a$ and $\Delta = \mathcal{A} \left(\frac{4e^2}{3k^2} \mathbf{H}, \frac{1}{k} \mathbf{b} \right)$, the vector $\mathbf{x}' = \mathbf{x} + \frac{1}{k} \Delta$ satisfies

$$g(\mathbf{x}') - g(\mathbf{x}^*) \leq \left(1 - \frac{1}{4e^4 \max(kR_\infty, 1)} \right) (g(\mathbf{x}) - g(\mathbf{x}^*)) + \frac{e^2 \delta_a}{k^2} + \frac{3}{2} \delta,$$

where R_∞ is the ℓ_∞ -radius of the sublevel set $\{\mathbf{x}' : g(\mathbf{x}') \leq g(\mathbf{x})\}$ about \mathbf{x} .

3.2 A second-order robust potential for matrix scaling and its properties

Given a sparse matrix $\mathbf{A} \in \mathbb{R}_{\geq 0}^{n \times n}$, a desired error $\varepsilon > 0$, and some number $B > 0$, we consider the regularized potential function $\tilde{f}(x, y)$ given by

$$\tilde{f}(\mathbf{x}, \mathbf{y}) = f(\mathbf{x}, \mathbf{y}) + \frac{\varepsilon^2}{ne^B} \left(\sum_i (e^{x_i} + e^{-x_i}) + \sum_j (e^{y_j} + e^{-y_j}) \right),$$

where f is the commonly-used potential function from Equation (1.2). In [12], the same regularization term is used, but with a different weight (since they aim for ℓ_2 -scaling and we aim for ℓ_1 -scaling). The following is an adaptation of [12, Lem. 4.10], see Appendix B.2.

► **Lemma 3.6.** *Assume \mathbf{A} is asymptotically scalable, with $\|\mathbf{A}\|_1 \leq 1$, and $\mu > 0$ its smallest non-zero entry. Let $B > 0$ and $\varepsilon > 0$ be given. Then the regularized potential \tilde{f} satisfies the following properties:*

1. \tilde{f} is second-order robust with respect to ℓ_∞ , and its Hessian is SDD;
2. we have $f(\mathbf{z}) \leq \tilde{f}(\mathbf{z})$ for any $\mathbf{z} = (\mathbf{x}, \mathbf{y})$,
3. for all \mathbf{z} such that $\tilde{f}(\mathbf{z}) \leq \tilde{f}(\mathbf{0})$, we have $\|\mathbf{z}\|_\infty \leq B + \ln(4n + (n \ln(1/\mu)/\varepsilon^2))$, and
4. for any \mathbf{z}_ε such that $f(\mathbf{z}_\varepsilon) \leq f^* + \varepsilon^2$ and $\|\mathbf{z}_\varepsilon\|_\infty \leq B$, one has $\tilde{f}(\mathbf{z}_\varepsilon) \leq f^* + 5\varepsilon^2$. In particular, if such a \mathbf{z}_ε exists, then $|f^* - \tilde{f}^*| \leq 5\varepsilon^2$.

In order to use Theorem 3.5 to minimize f , we need to show how to approximate both the gradient and Hessian of \tilde{f} . We first consider the Hessian of \tilde{f} , which can be written as the sum of the Hessian of f and the Hessian of the regularizer $\tilde{f} - f$. We have

$$\begin{aligned} \nabla^2 f(\mathbf{x}, \mathbf{y}) &= \begin{bmatrix} \text{diag}(\mathbf{r}(\mathbf{A}(\mathbf{x}, \mathbf{y}))) & \mathbf{A}(\mathbf{x}, \mathbf{y}) \\ \mathbf{A}(\mathbf{x}, \mathbf{y})^T & \text{diag}(\mathbf{c}(\mathbf{A}(\mathbf{x}, \mathbf{y}))) \end{bmatrix}, \\ \nabla^2(\tilde{f} - f)(\mathbf{x}, \mathbf{y}) &= \frac{\varepsilon^2}{ne^B} \begin{bmatrix} \text{diag}(e^{\mathbf{x}} + e^{-\mathbf{x}}) & \mathbf{0} \\ \mathbf{0} & \text{diag}(e^{\mathbf{y}} + e^{-\mathbf{y}}) \end{bmatrix}. \end{aligned} \quad (3.1)$$

Note that computing $\nabla^2 \tilde{f}(\mathbf{x}, \mathbf{y})$ up to high precision can be done using $\tilde{O}(m)$ classical queries to \mathbf{A} , \mathbf{x} , and \mathbf{y} . Below we show how to obtain a sparse approximation of $\nabla^2 \tilde{f}(\mathbf{x}, \mathbf{y})$ using only $\tilde{O}(\sqrt{mn})$ quantum queries. We will do so in the sense of condition (2) of Theorem 3.5 where we take \mathbf{H}'_m to be a (high-precision) additive approximation of $\nabla^2 f(\mathbf{x}, \mathbf{y})$, and $\mathbf{H}'_a = \nabla^2 \tilde{f}(\mathbf{x}, \mathbf{y}) - \mathbf{H}'_m$.

We first obtain a multiplicative spectral approximation of (a high-precision additive approximation of) $\nabla^2 f(\mathbf{x}, \mathbf{y})$. In order to do so we use its structure: it is similar to a Laplacian matrix. This allows us to use the recent quantum Laplacian sparsifier of Apers and de Wolf [6]. For a full proof, carefully keeping track of the bit-complexity, we refer to the full version [17].

► **Lemma 3.7.** *Given quantum query access to \mathbf{x}, \mathbf{y} and sparse quantum query access to \mathbf{A} , such that $\|\mathbf{A}(\mathbf{x}, \mathbf{y})\|_1 \leq C$, we can compute an SDD matrix \mathbf{H}_m with $\tilde{O}(n)$ non-zero entries, each off-diagonal entry non-negative, such that there exist symmetric \mathbf{H}'_m and $\mathbf{H}'_{a,f}$ satisfying $\mathbf{H}'_m + \mathbf{H}'_{a,f} = \nabla^2 f(\mathbf{x}, \mathbf{y})$, and $0.9\mathbf{H}_m \preceq \mathbf{H}'_m \preceq 1.1\mathbf{H}_m$, $\|\mathbf{H}'_{a,f}\|_1 \leq \delta_a$, in time $\tilde{O}(\sqrt{mn} \text{polylog}(C/\delta_a))$.*

Similarly, we can efficiently compute an additive approximation of the Hessian of the regularization term $\tilde{f} - f$ as long as \mathbf{x} and \mathbf{y} have ℓ_∞ -norm not much larger than B , using the expression given in Equation (3.1).

► **Lemma 3.8.** *Given quantum query access to \mathbf{x}, \mathbf{y} with $\|\mathbf{x}\|_\infty, \|\mathbf{y}\|_\infty \leq B + \ln(4n + (n \ln(1/\mu)/\varepsilon^2))$, we can compute a non-negative diagonal matrix $\mathbf{H}_{a,\tilde{f}}$ that satisfies $\|\mathbf{H}_{a,\tilde{f}} - \nabla^2(\tilde{f} - f)(\mathbf{x}, \mathbf{y})\|_1 \leq \delta_a$, in time $\tilde{O}(n \log(1/\delta_a \mu) \text{polylog}(\varepsilon))$.*

► **Theorem 3.9.** *Given quantum query access to \mathbf{x}, \mathbf{y} with $\|\mathbf{x}\|_\infty, \|\mathbf{y}\|_\infty \leq B + \ln(4n + (n \ln(1/\mu)/\varepsilon^2))$, and sparse quantum query access to \mathbf{A} , if $\|\mathbf{A}(\mathbf{x}, \mathbf{y})\|_1 \leq C$, then we can compute (classical descriptions of) an SDD matrix \mathbf{H}_m with $\tilde{O}(n)$ non-zero entries, with all of the off-diagonal entries non-negative, and a non-negative diagonal matrix \mathbf{H}_a such that there exist symmetric $\mathbf{H}'_m, \mathbf{H}'_a$ with $\mathbf{H}'_m + \mathbf{H}'_a = \nabla^2 \tilde{f}(\mathbf{x}, \mathbf{y})$ and*

$$0.9\mathbf{H}_m \preceq \mathbf{H}'_m \preceq 1.1\mathbf{H}_m, \quad \|\mathbf{H}_a - \mathbf{H}'_a\|_1 \leq \delta_a$$

in quantum time $\tilde{O}(\sqrt{mn} \text{polylog}(C/\mu\delta_a))$.

Proof. Let \mathbf{H}_m be the matrix obtained from Lemma 3.7, and let \mathbf{H}_a be the matrix $\mathbf{H}_{a,\tilde{f}}$ obtained from Lemma 3.8. Then \mathbf{H} satisfies the desired properties, with \mathbf{H}'_m as in Lemma 3.7, and $\mathbf{H}'_a = \mathbf{H}'_{a,f} + \nabla^2(\tilde{f} - f)(\mathbf{x}, \mathbf{y})$ with $\mathbf{H}'_{a,f}$ as in Lemma 3.7. ◀

In order to obtain a good approximation of the gradient of \tilde{f} , which is given by

$$\nabla \tilde{f}(\mathbf{x}, \mathbf{y}) = \left[\mathbf{r}(\mathbf{A}(\mathbf{x}, \mathbf{y})) - \mathbf{r} \right] + \frac{\varepsilon^2}{ne^B} \begin{bmatrix} e^{\mathbf{x}} - e^{-\mathbf{x}} \\ e^{\mathbf{y}} - e^{-\mathbf{y}} \end{bmatrix},$$

we can use similar techniques as the prior work on quantum algorithms for matrix scaling [5]. For computing the i -th row marginal, these are based on a careful implementation of amplitude estimation on the unitary that prepares states that are approximately of the form

$$\sum_j |0\rangle \sqrt{A_{ij} e^{x_i + y_j}} |j\rangle + |1\rangle \sqrt{1 - A_{ij} e^{x_i + y_j}} |j\rangle,$$

assuming that the i -th row of $\mathbf{A}(\mathbf{x}, \mathbf{y})$ is properly normalized. The output is an estimate of the i -th row marginal with multiplicative error $1 \pm \delta$, which translates into additive error $\delta \cdot r_i(\mathbf{A}(\mathbf{x}, \mathbf{y}))$; we refer to [5, Thm. 4.5 (arXiv)] for a more precise statement. The part of the gradient coming from the regularization term is dealt with similarly as in Lemma 3.8.

► **Lemma 3.10.** *Given quantum query access to \mathbf{x}, \mathbf{y} and sparse quantum query access to \mathbf{A} , if $\|\mathbf{A}(\mathbf{x}, \mathbf{y})\|_1 \leq C$, we can find a classical description of a vector $\mathbf{b} \in \mathbb{R}^n$ such that $\|\mathbf{b} - \nabla \tilde{f}(\mathbf{x}, \mathbf{y})\|_1 \leq \delta \cdot C$ in quantum time $\tilde{O}(\sqrt{mn}/\delta \cdot \text{polylog}(C/\mu))$.*

The following lemma and corollary help us ensure that throughout the algorithm, $\|\mathbf{A}(\mathbf{x}, \mathbf{y})\|_1$ is bounded above by a constant; if $\|\mathbf{A}(\mathbf{x}, \mathbf{y})\|_1$ is too large, we can change the overall scaling of the matrix and decrease the regularized potential (so in particular, we stay in the sublevel set of the regularized potential).

► **Lemma 3.11.** *Let \mathbf{x}, \mathbf{y} be such that $\tilde{f}(\mathbf{x}, \mathbf{y}) \leq \tilde{f}(\mathbf{0}, \mathbf{0})$, and assume $\|\mathbf{A}(\mathbf{x}, \mathbf{y})\|_1 \geq C'$ where $C' > 1$. Let $\mathbf{x}' = \mathbf{x} - \ln(\gamma)\mathbf{1}$ where $1 \leq \gamma \leq C'$. Then $\tilde{f}(\mathbf{x}', \mathbf{y}) - \tilde{f}(\mathbf{x}, \mathbf{y}) \leq (\frac{1}{\gamma} - 1)C' + \ln(\gamma) + (\gamma - 1) \left(\ln(1/\mu) + \frac{4\varepsilon^2}{e^B} \right)$.*

Proof. We have

$$\begin{aligned} & \tilde{f}(\mathbf{x}', \mathbf{y}) - \tilde{f}(\mathbf{x}, \mathbf{y}) \\ &= \left(\frac{1}{\gamma} - 1 \right) \|\mathbf{A}(\mathbf{x}, \mathbf{y})\|_1 + \ln(\gamma) + \frac{\varepsilon^2}{ne^B} \left(\frac{1}{\gamma} - 1 \right) \left(\sum_i e^{x_i} \right) + \frac{\varepsilon^2}{ne^B} (\gamma - 1) \left(\sum_i e^{-x_i} \right) \end{aligned}$$

$$\begin{aligned}
&\leq \left(\frac{1}{\gamma} - 1\right) \|\mathbf{A}(\mathbf{x}, \mathbf{y})\|_1 + \ln(\gamma) + 0 + \frac{\varepsilon^2}{ne^B}(\gamma - 1) \left(\sum_i e^{-x_i}\right) \\
&\leq \left(\frac{1}{\gamma} - 1\right) C' + \ln(\gamma) + (\gamma - 1) \left(\ln(1/\mu) + \frac{4\varepsilon^2}{e^B}\right)
\end{aligned}$$

where for the last inequality we use $\|\mathbf{A}(\mathbf{x}, \mathbf{y})\|_1 \geq C'$ for the first term and Equation (B.5) for the last term. \blacktriangleleft

An appropriate choice of C' and γ makes the bound in the above lemma non-positive.

► **Corollary 3.12.** *Let $\varepsilon \leq 1$ and $\mu \leq 1$, set $\gamma = 2$ and $C' = 2(\ln(2/\mu) + 4\varepsilon^2/e^B)$. Then, if $\|\mathbf{A}(\mathbf{x}, \mathbf{y})\|_1 \geq C'$ and $\tilde{f}(\mathbf{x}, \mathbf{y}) \leq \tilde{f}(\mathbf{0}, \mathbf{0})$, we have $\tilde{f}(\mathbf{x}', \mathbf{y}) \leq \tilde{f}(\mathbf{x}, \mathbf{y})$.*

3.3 Quantum box-constrained scaling

Combining the above leads to a quantum algorithm for matrix scaling that is based on classical box-constrained Newton methods. See Algorithm 1 for its formal definition. In Theorem 3.13 we analyze its output.

■ **Algorithm 1** Quantum box-constrained Newton method for matrix scaling.

Input: Oracle access to $\mathbf{A} \in [\mu, 1]^{n \times n}$ with $\|\mathbf{A}\|_1 \leq 1$ and $\mu > 0$, error $\varepsilon > 0$, targets $\mathbf{r}, \mathbf{c} \in \mathbb{R}_{>0}^n$ with $\|\mathbf{r}\|_1 = 1 = \|\mathbf{c}\|_1$, diameter bound $B \geq 1$, classical k -oracle \mathcal{A} for SDD matrices with non-negative off-diagonal entries

Output: Vectors $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$ with $\|(\mathbf{x}, \mathbf{y})\|_\infty \leq B + \ln(4n + (n \ln(1/\mu)/\varepsilon^2))$

- 1 set $T = \lceil 4e^4 \max(kB + \ln(4n + (n \ln(1/\mu)/\varepsilon^2)), 1) \cdot \ln\left(\frac{\ln(1/\mu) + 2\varepsilon^2/e^B}{\varepsilon^2/2}\right) \rceil$;
- 2 set $C' = 2\lceil \ln(2/\mu) + 8\varepsilon^2/e^B \rceil$;
- 3 set $\varepsilon' = \lfloor \varepsilon^2/8e^4 \max(k(B + \ln(4n + (n \ln(1/\mu)/\varepsilon^2))), 1) \rfloor$;
- 4 store $\mathbf{x}^{(0)}, \mathbf{y}^{(0)} = \mathbf{0} \in \mathbb{R}^n$ in QCRAM;
- 5 **for** $i = 0, \dots, T - 1$ **do**
- 6 compute $\mathbf{H}_m, \mathbf{H}_a$ s.t. $\mathbf{H}_m + \mathbf{H}_a \approx \nabla^2 \tilde{f}(\mathbf{x}^{(i)}, \mathbf{y}^{(i)})$ as in Theorem 3.9 with $\delta_a = \varepsilon' k^2 / 2e^2$;
- 7 compute $\mathbf{b} \approx \nabla \tilde{f}(\mathbf{x}^{(i)}, \mathbf{y}^{(i)})$ as in Lemma 3.10 at $\mathbf{x}^{(i)}, \mathbf{y}^{(i)}$ with $\delta = \varepsilon'/3$;
- 8 compute $\Delta = \mathcal{A}(\frac{4e^2}{3k^2} \cdot (\mathbf{H}_m + \mathbf{H}_a), \frac{\mathbf{b}}{k})$;
- 9 compute $(\mathbf{x}^{(i+1)}, \mathbf{y}^{(i+1)}) = (\mathbf{x}^{(i)}, \mathbf{y}^{(i)}) + \frac{1}{k}\Delta$ and store in QCRAM;
- 10 set **flag** = **true**;
- 11 **while** **flag** **do**
- 12 Compute $C'/2$ -additive approximation γ of $\|\mathbf{A}(\mathbf{x}^{(i+1)}, \mathbf{y}^{(i+1)})\|_1$;
- 13 **if** $\gamma \leq 3C'/2$ **then**
- 14 set **flag** = **false**;
- 15 **else**
- 16 update $\mathbf{x}^{(i+1)} \leftarrow \mathbf{x}^{(i+1)} - \ln(2)\mathbf{1}$ in QCRAM;
- 17 **end if**
- 18 **end while**
- 19 **end for**
- 20 **return** $(\mathbf{x}, \mathbf{y}) = (\mathbf{x}^{(T)}, \mathbf{y}^{(T)})$;

► **Theorem 3.13.** *Let $\mathbf{A} \in [0, 1]^{n \times n}$ with m non-zero entries, $\mathbf{r}, \mathbf{c} \in \mathbb{R}_{\geq 0}^n$ such that $\|\mathbf{r}\|_1 = 1 = \|\mathbf{c}\|$, and assume \mathbf{A} is asymptotically (\mathbf{r}, \mathbf{c}) -scalable. Let $\varepsilon > 0$, let $B \geq 1$, and assume there exist $(\mathbf{x}_\varepsilon, \mathbf{y}_\varepsilon)$ such that $\|(\mathbf{x}_\varepsilon, \mathbf{y}_\varepsilon)\|_\infty \leq B$ and $f(\mathbf{x}_\varepsilon, \mathbf{y}_\varepsilon) - f^* \leq \varepsilon^2$. Furthermore, let \mathcal{A} be the $O(\log(n))$ -oracle of Theorem 3.4. Then Algorithm 1 with these parameters outputs, with probability $\geq 2/3$, vectors \mathbf{x}, \mathbf{y} such that $f(\mathbf{x}, \mathbf{y}) - f^* \leq 6\varepsilon^2$ and runs in quantum time $\tilde{O}(B^2 \sqrt{mn}/\varepsilon^2)$.*

Proof. In every iteration, the matrices $\mathbf{H}_m, \mathbf{H}_a$ and the vector \mathbf{b} are such that they satisfy the requirements of Theorem 3.5, hence

$$\tilde{f}(\mathbf{x}^{(i+1)}, \mathbf{y}^{(i+1)}) - \tilde{f}^* \leq \left(1 - \frac{1}{4e^4 \max(kR_\infty, 1)}\right) (\tilde{f}(\mathbf{x}^{(i)}, \mathbf{y}^{(i)}) - \tilde{f}^*) + \frac{e^2 \delta_a}{k^2} + \frac{3\delta}{2}$$

where $R_\infty \leq B + \ln(4n + (n \ln(1/\mu)/\varepsilon^2))$ is the ℓ_∞ -radius of the sublevel set $\{(\mathbf{x}, \mathbf{y}) : \tilde{f}(\mathbf{x}, \mathbf{y}) \leq \tilde{f}(\mathbf{0}, \mathbf{0})\}$ about $(\mathbf{0}, \mathbf{0})$, whose upper bound follows from Lemma 3.6. From here on, we write $M = 4e^4 \max(kR_\infty, 1)$. The choice of δ_a and δ in the algorithm is such that $e^2 \delta_a / k^2 + 3\delta/2 \leq \frac{\varepsilon^2}{2M}$, hence we can also bound the progress by

$$\tilde{f}(\mathbf{x}^{(i+1)}, \mathbf{y}^{(i+1)}) - \tilde{f}^* \leq \left(1 - \frac{1}{M}\right) (\tilde{f}(\mathbf{x}^{(i)}, \mathbf{y}^{(i)}) - \tilde{f}^*) + \frac{\varepsilon^2}{2M}.$$

Corollary 3.12 shows that if $\|\mathbf{A}(\mathbf{x}^{(i+1)}, \mathbf{y}^{(i+1)})\|_1$ is larger than C' , then we can shift \mathbf{x} by $-\ln(2)\mathbf{1}$, this halves $\|\mathbf{A}(\mathbf{x}^{(i+1)}, \mathbf{y}^{(i+1)})\|_1$ and does not increase the regularized potential. Repeating this roughly $\log_2(\|\mathbf{A}(\mathbf{x}^{(i+1)}, \mathbf{y}^{(i+1)})\|_1 / C')$ many times² reduces $\|\mathbf{A}(\mathbf{x}^{(i+1)}, \mathbf{y}^{(i+1)})\|_1$ to at most $C = 2C'$. Determining when to stop this process requires a procedure to distinguish between the cases $\|\mathbf{A}(\mathbf{x}^{(i+1)}, \mathbf{y}^{(i+1)})\|_1 \leq C'$ and $\|\mathbf{A}(\mathbf{x}^{(i+1)}, \mathbf{y}^{(i+1)})\|_1 \geq 2C'$ (if in between C' and $2C'$ either continuing or stopping is fine). Such a procedure can be implemented by computing a $C'/2$ -additive approximation of $\|\mathbf{A}(\mathbf{x}^{(i+1)}, \mathbf{y}^{(i+1)})\|_1$, which can be done using $\tilde{O}(\sqrt{mn} \text{polylog}(C'/\mu))$ quantum queries, see (the proof of) [5, Lemma 4.6 (arXiv)]. Therefore, throughout the algorithm we may assume that $\|\mathbf{A}(\mathbf{x}^{(i+1)}, \mathbf{y}^{(i+1)})\|_1 \leq 2C' = C$.

It remains to show that $\tilde{f}(\mathbf{x}^{(T)}, \mathbf{y}^{(T)}) - \tilde{f}^* \leq \varepsilon^2$ for our choice of T . Note that we have

$$\begin{aligned} \tilde{f}(\mathbf{x}^{(T)}, \mathbf{y}^{(T)}) - \tilde{f}^* &\leq \left(1 - \frac{1}{M}\right)^T (\tilde{f}(\mathbf{0}, \mathbf{0}) - \tilde{f}^*) + \sum_{i=0}^{T-1} \left(1 - \frac{1}{M}\right)^{T-i-1} \cdot \frac{\varepsilon^2}{2M} \\ &\leq \left(1 - \frac{1}{M}\right)^T (\tilde{f}(\mathbf{0}, \mathbf{0}) - \tilde{f}^*) + \left(1 - \left(1 - \frac{1}{M}\right)^T\right) \cdot \frac{\varepsilon^2}{2} \\ &\leq \left(1 - \frac{1}{M}\right)^T \left(f(\mathbf{0}, \mathbf{0}) - f^* + \frac{2\varepsilon^2}{e^B}\right) + \frac{\varepsilon^2}{2} \\ &\leq \left(1 - \frac{1}{M}\right)^T \left(\ln(1/\mu) + \frac{2\varepsilon^2}{e^B}\right) + \frac{\varepsilon^2}{2} \leq \varepsilon^2 \end{aligned}$$

where in the third inequality we use Lemma 3.6, and in the last inequality we use

$$\begin{aligned} T &= \left\lceil 4e^4 \max(kB + \ln(4n + (n \ln(1/\mu)/\varepsilon^2)), 1) \cdot \ln\left(\frac{\ln(1/\mu) + 2\varepsilon^2/e^B}{\varepsilon^2/2}\right) \right\rceil \\ &\geq \left\lceil M \cdot \ln\left(\frac{\ln(1/\mu) + 2\varepsilon^2/e^B}{\varepsilon^2/2}\right) \right\rceil \geq \frac{1}{\ln(1 - \frac{1}{M})} \cdot \ln\left(\frac{\varepsilon^2/2}{\ln(1/\mu) + \frac{2\varepsilon^2}{e^B}}\right). \end{aligned}$$

² Which is an almost constant number of times: in a single update of the box-constrained method, we take steps of size at most 1 in ℓ_∞ -norm, so individual entries can only grow by a factor e^2 in a single iteration, and the holds same for $\|\mathbf{A}(\mathbf{x}, \mathbf{y})\|_1$.

This implies that $f(\mathbf{x}^{(T)}, \mathbf{y}^{(T)}) - f^* \leq \tilde{f}(\mathbf{x}^{(T)}, \mathbf{y}^{(T)}) - \tilde{f}^* + 5\varepsilon^2 \leq 6\varepsilon^2$, where we crucially use the last point of Lemma 3.6 and the assumption that there exist $(\mathbf{x}_\varepsilon, \mathbf{y}_\varepsilon)$ with $\|(\mathbf{x}_\varepsilon, \mathbf{y}_\varepsilon)\|_\infty \leq B$ which ε^2 -minimize f .

Finally we bound the time complexity of Algorithm 1. For each of the quoted results, we use the choice $C = 2C' = \tilde{O}(\ln(n) + \varepsilon^2)$. In each of the T iterations we compute:

1. approximations $\mathbf{H}_m, \mathbf{H}_a$ of $\nabla^2 \tilde{f}(\mathbf{x}^{(i)}, \mathbf{y}^{(i)})$ in time $\tilde{O}(\sqrt{mn} \text{polylog}(1/\varepsilon))$ (using that $C, 1/\mu$ are at most $\text{poly}(n)$),
2. an $\varepsilon'/3$ - ℓ_1 -approximation of $\nabla \tilde{f}(\mathbf{x}^{(i)}, \mathbf{y}^{(i)})$ in time $\tilde{O}(\sqrt{mn}/\varepsilon') = \tilde{O}(B\sqrt{mn}/\varepsilon^2)$,
3. an update Δ in time $\tilde{O}(n)$ using one call to the $k = O(\log(n))$ -oracle on SDD-matrices with $\tilde{O}(n)$ non-zero entries from Theorem 3.4,
4. at most $O(1)$ many times (using the fact that in Algorithm 1 the 1-norm changes by at most a constant factor since $\|\frac{1}{k}\Delta\|_\infty \leq 1$) an $O(\ln(1/\mu) + \varepsilon^2)$ -additive approximation of $\|\mathbf{A}(\mathbf{x}^{(i)}, \mathbf{y}^{(i)})\|_1$ in time $\tilde{O}(\sqrt{mn})$.

Note that the second contribution dominates the others, resulting in an overall time complexity of $\tilde{O}(B^2\sqrt{mn}/\varepsilon^2)$. \blacktriangleleft

The above proof relies on Theorem 3.5 to show that the (regularized) potential decreases in each iteration. This decrease depends on the precision used for the marginal estimation in that iteration and one can show that the choice of precision in Algorithm 1 is asymptotically optimal, see the full version [17].

Algorithm 1 takes as (part of the) input a bound B on the ℓ_∞ -norm of an ε^2 -minimizer of f . For the purpose of matrix scaling, one can avoid knowing such a bound in advance, by running the algorithm for successive powers of 2 (i.e., $B = 1, B = 2, B = 4, \dots$) and testing whether the output yields an ε -scaling or not. Verifying whether given \mathbf{x}, \mathbf{y} yield an ε -scaling of \mathbf{A} can be done in time $\tilde{O}(\sqrt{mn}/\varepsilon^2)$. Note that this gives an algorithm for ε -scaling whose complexity depends on a diameter bound for ε^2 -minimizers of f , rather than a diameter bound for ε -scaling vectors. Furthermore, such an approach does not work for the task of finding an ε^2 -minimizer of f , as we do not know how to test this property efficiently.

► **Corollary 3.14.** *For asymptotically-scalable matrices $\mathbf{A} \in \mathbb{R}_{\geq 0}^{n \times n}$ with m non-zero entries, one can find $O(\varepsilon)$ - ℓ_1 -scaling vectors (\mathbf{x}, \mathbf{y}) of \mathbf{A} to target marginals $\mathbf{r}, \mathbf{c} \in \mathbb{R}_{> 0}^n$ with $\|\mathbf{r}\|_1 = 1 = \|\mathbf{c}\|_1$ in time $\tilde{O}(R_\infty^2\sqrt{mn}/\varepsilon^2)$, where R_∞ is such that there exists an ε^2 -approximate minimizer $(\mathbf{x}_\varepsilon, \mathbf{y}_\varepsilon)$ of f with $R_\infty = \|(\mathbf{x}_\varepsilon, \mathbf{y}_\varepsilon)\|_\infty + \ln(4n + (n \ln(1/\mu)/\varepsilon^2))$.*

For the general case mentioned above, we do not have good (i.e., polylogarithmic) bounds on the parameter R_∞ . We do have such bounds when \mathbf{A} is entrywise positive: it is well-known that such an \mathbf{A} can be exactly scaled to uniform marginals with scaling vectors (\mathbf{x}, \mathbf{y}) such that $\|(\mathbf{x}, \mathbf{y})\|_\infty = O(\log(\|\mathbf{A}\|_1/\mu))$ (cf. [22, Lem. 1], [12, Lem. 4.11]). In particular, this implies that there exists a minimizer $(\mathbf{x}^*, \mathbf{y}^*)$ of f with $\|(\mathbf{x}^*, \mathbf{y}^*)\|_\infty = O(\log(\|\mathbf{A}\|_1/\mu)) = \tilde{O}(1)$ and therefore we have the following corollary.

► **Corollary 3.15.** *For entrywise-positive matrices \mathbf{A} , one can find an ε - ℓ_1 -scaling of \mathbf{A} to uniform marginals in time $\tilde{O}(n^{1.5}/\varepsilon^2)$.*

References

- 1 Zeyuan Allen-Zhu, Yuanzhi Li, Rafael Oliveira, and Avi Wigderson. Much faster algorithms for matrix scaling. In *Proceedings of IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS'17)*, pages 890–901, 2017. [arXiv:1704.02315](#).
- 2 Jason Altschuler, Jonathan Niles-Weed, and Philippe Rigollet. Near-linear time approximation algorithms for optimal transport via Sinkhorn iteration. In *Advances in Neural Information Processing Systems*, volume 30, pages 1964–1974, 2017. [arXiv:1705.09634](#).

- 3 Andris Ambainis. Quantum lower bounds by quantum arguments. *Journal of Computer and System Sciences*, 64(4):750–767, 2002. Earlier version in STOC’00. [arXiv:quant-ph/0002066](#).
- 4 Edward Anderson, Zhaojun Bai, Christian Bischof, L Susan Blackford, James Demmel, Jack Dongarra, Jeremy Du Croz, Anne Greenbaum, Sven Hammarling, Alan McKenney, et al. *LAPACK Users’ guide*. SIAM, 1999.
- 5 Joran van Apeldoorn, Sander Gribling, Yinan Li, Harold Nieuwboer, Michael Walter, and Ronald de Wolf. Quantum Algorithms for Matrix Scaling and Matrix Balancing. In *48th International Colloquium on Automata, Languages, and Programming (ICALP 2021)*, volume 198, pages 110:1–110:17, 2021. [arXiv:2011.12823v1](#), [doi:10.4230/LIPIcs.ICALP.2021.110](#).
- 6 Simon Apers and Ronald de Wolf. Quantum speedup for graph sparsification, cut approximation and laplacian solving. *Proceedings of IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS’20)*, pages 637–648, 2020. [arXiv:1911.07306](#).
- 7 Robert Beals, Harry Buhrman, Richard Cleve, Michele Mosca, and Ronald de Wolf. Quantum lower bounds by polynomials. *Journal of the ACM (JACM)*, 48(4):778–797, 2001. [arXiv:quant-ph/9802049](#).
- 8 Yvonne M. M. Bishop, Stephen E. Fienberg, and Paul W. Holland. *Discrete Multivariate Analysis: Theory and Practice*. MIT Press, 1975.
- 9 David T. Brown. A note on approximations to discrete probability distributions. *Information and control*, 2(4):386–392, 1959.
- 10 Peter Bürgisser, Cole Franks, Ankit Garg, Rafael Oliveira, Michael Walter, and Avi Wigderson. Towards a theory of non-commutative optimization: geodesic 1st and 2nd order methods for moment maps and polytopes. In *Proceedings of 60th IEEE Annual Symposium on Foundations of Computer Science (FOCS’19)*, pages 845–861, 2019. [arXiv:1910.12375](#).
- 11 Deeparnab Chakrabarty and Sanjeev Khanna. Better and simpler error analysis of the Sinkhorn–Knopp algorithm for matrix scaling. *Mathematical Programming*, pages 1–13, 2020.
- 12 Michael B. Cohen, Aleksander Madry, Dimitris Tsipras, and Adrian Vladu. Matrix scaling and balancing via box constrained Newton’s method and interior point methods. In *Proceedings of IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS’17)*, pages 902–913, 2017. [doi:10.1109/FOCS.2017.88](#).
- 13 Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. In *Advances in Neural Information Processing Systems*, volume 26, pages 2292–2300, 2013. [arXiv:1306.0895](#).
- 14 W. Edwards Deming and Frederick F. Stephan. On a least squares adjustment of a sampled frequency table when the expected marginal totals are known. *The Annals of Mathematical Statistics*, 11(4):427–444, 1940.
- 15 Pavel E. Dvurechensky, Alexander V. Gasnikov, and Alexey Kroshnin. Computational Optimal Transport: Complexity by Accelerated Gradient Descent Is Better Than by Sinkhorn’s Algorithm. In *ICML*, 2018. [arXiv:1802.04367](#).
- 16 Ankit Garg, Leonid Gurvits, Rafael Oliveira, and Avi Wigderson. Operator scaling: theory and applications. *Foundations of Computational Mathematics*, pages 1–68, 2019. Earlier version in FOCS’16.
- 17 Sander Gribling and Harold Nieuwboer. Improved quantum lower and upper bounds for matrix scaling, 2021. [arXiv:2109.15282v1](#).
- 18 Peter Høyer, Troy Lee, and Robert Špalek. Negative weights make adversaries stronger. In *Proceedings of the 39th Annual ACM SIGACT Symposium on Theory of Computing (STOC’07)*, pages 526–535, 2007. [arXiv:quant-ph/0611054](#), [doi:10.1145/1250790.1250867](#).
- 19 Martin Idel. A review of matrix scaling and Sinkhorn’s normal form for matrices and positive maps, 2016. [arXiv:1609.06349](#).
- 20 B. Kalantari, I. Lari, F. Ricca, and B. Simeone. On the complexity of general matrix scaling and entropy minimization via the RAS algorithm. *Mathematical Programming*, 112:371–401, 2008.

- 21 Bahman Kalantari and Leonid Khachiyan. On the rate of convergence of deterministic and randomized RAS matrix scaling algorithms. *Operations Research Letters*, 14(5):237–244, 1993. doi:10.1016/0167-6377(93)90087-W.
- 22 Bahman Kalantari and Leonid Khachiyan. On the complexity of nonnegative-matrix scaling. *Linear Algebra and its Applications*, 240:87–103, 1996. doi:10.1016/0024-3795(94)00188-X.
- 23 J. Kruithof. Telefoonverkeersrekening. *De Ingenieur*, 52:E15–E25, 1937.
- 24 Rasmus Kyng, Yin Tat Lee, Richard Peng, Sushant Sachdeva, and Daniel A. Spielman. Sparsified Cholesky and multigrid solvers for connection Laplacians. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing (STOC’16)*, pages 842–850, 2016. arXiv:1512.01892, doi:10.1145/2897518.2897640.
- 25 Troy Lee, Rajat Mittal, Ben Reichardt, Robert Špalek, and Mario Szegedy. Quantum query complexity of state conversion. *Proceedings of IEEE 52nd Annual Symposium on Foundations of Computer Science (FOCS’11)*, pages 344–353, 2011. arXiv:1011.3020, doi:10.1109/FOCS.2011.75.
- 26 Troy Lee and Jérémie Roland. A strong direct product theorem for quantum query complexity. *Computational Complexity*, 22(2):429–462, 2013. Earlier version in CCC’12. arXiv:1104.4468.
- 27 Yin Tat Lee, Richard Peng, and Daniel A. Spielman. Sparsified Cholesky solvers for SDD linear systems, 2015. arXiv:1506.08204.
- 28 Nathan Linial, Alex Samorodnitsky, and Avi Wigderson. A deterministic strongly polynomial algorithm for matrix scaling and approximate permanents. *Combinatorica*, 20(4):545–568, 2000.
- 29 Ashwin Nayak and Felix Wu. The quantum query complexity of approximating the median and related statistics. In *Proceedings of the 31st Annual ACM SIGACT Symposium on Theory of Computing (STOC’99)*, pages 384–393, 1999. arXiv:quant-ph/9804066, doi:10.1145/301250.301349.
- 30 Uriel G. Rothblum and Hans Schneider. Scalings of matrices which have prespecified row sums and column sums via optimization. *Linear Algebra and its Applications*, 114:737–764, 1989.
- 31 Richard Sinkhorn. A relationship between arbitrary positive matrices and doubly stochastic matrices. *The Annals of Mathematical Statistics*, 35(2):876–879, 1964.

A

 Missing proofs for Section 2

A.1 Strong convexity properties of the potential

Lemma 2.8 shows that a vector $(\mathbf{x}, \mathbf{y}) \in V$ for which $\|\nabla f(\mathbf{x}, \mathbf{y})\|_2$ is small, is close to the minimizer of f . Here we prove this lemma (see Corollary A.6) using strong convexity properties of f . In Lemma A.1 we show that the Hessian of f restricted to V has smallest eigenvalue at least $n \cdot \mu(\mathbf{x}, \mathbf{y})$ where $\mu(\mathbf{x}, \mathbf{y})$ is the smallest entry appearing in $(A_{ij}e^{x_i+y_j})_{i,j}$. In Lemma A.3 we show that $\mu(\mathbf{x}^*, \mathbf{y}^*) = \Theta(1/n^2)$. This implies that $\mu(\mathbf{x}, \mathbf{y}) = \Theta(1/n^2)$ for all (\mathbf{x}, \mathbf{y}) that are a constant distance away from $(\mathbf{x}^*, \mathbf{y}^*)$ in the ℓ_∞ -norm, in other words, f is $\Theta(1/n)$ -strongly convex around its minimizer. Lemma A.5 summarizes these lemmas: it gives a quantitative bound on the distance to a minimizer, in terms of the gradient.

► **Lemma A.1.** *Let \mathbf{A} be an entrywise non-negative $n \times n$ matrix and let $f: V \subset \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}$ be the potential for this matrix as given in (1.2), where V is the orthogonal complement of $(\mathbf{1}_n, -\mathbf{1}_n)$. Then $\nabla^2 f(\mathbf{x}, \mathbf{y}) \succeq \mu(\mathbf{x}, \mathbf{y}) \cdot n \cdot \mathbf{P}_V$ where \mathbf{P}_V is the projection onto V and $\mu(\mathbf{x}, \mathbf{y})$ is the smallest entry appearing in $\mathbf{A}(\mathbf{x}, \mathbf{y})$. In particular, f is strictly convex on V .*

Proof. The Hessian of the potential $f(\mathbf{x}, \mathbf{y}) = \sum_{i,j=1}^n A_{ij}e^{x_i+y_j} - \langle \mathbf{r}, \mathbf{x} \rangle - \langle \mathbf{c}, \mathbf{y} \rangle$ is given by

$$\nabla^2 f(\mathbf{x}, \mathbf{y}) = \begin{bmatrix} \text{diag}(\mathbf{r}(\mathbf{A}(\mathbf{x}, \mathbf{y}))) & \mathbf{A}(\mathbf{x}, \mathbf{y}) \\ \mathbf{A}(\mathbf{x}, \mathbf{y})^T & \text{diag}(\mathbf{c}(\mathbf{A}(\mathbf{x}, \mathbf{y}))) \end{bmatrix}.$$

We give a lower bound on the non-zero eigenvalues of the Hessian as follows. Conjugating the Hessian with the $2n \times 2n$ matrix $\text{diag}(\mathbf{I}, -\mathbf{I})$ preserves the spectrum, i.e., changing the signs of the off-diagonal $\mathbf{A}(\mathbf{x}, \mathbf{y})$ blocks yields a matrix which one can recognize as the weighted Laplacian of a complete bipartite graph. We denote by $\mu(\mathbf{x}, \mathbf{y})$ the smallest entry of $\mathbf{A}(\mathbf{x}, \mathbf{y})$ and we use \mathbf{J} for the $n \times n$ all-ones matrix. Then

$$\begin{bmatrix} \text{diag}(\mathbf{r}(\mathbf{A}(\mathbf{x}, \mathbf{y}))) & -\mathbf{A}(\mathbf{x}, \mathbf{y}) \\ -\mathbf{A}(\mathbf{x}, \mathbf{y})^T & \text{diag}(\mathbf{c}(\mathbf{A}(\mathbf{x}, \mathbf{y}))) \end{bmatrix} \succeq \begin{bmatrix} n\mu(\mathbf{x}, \mathbf{y})\mathbf{I} & -\mu(\mathbf{x}, \mathbf{y})\mathbf{J} \\ -\mu(\mathbf{x}, \mathbf{y})\mathbf{J} & n\mu(\mathbf{x}, \mathbf{y})\mathbf{I} \end{bmatrix} = \mu(\mathbf{x}, \mathbf{y}) \begin{bmatrix} n\mathbf{I} & -\mathbf{J} \\ -\mathbf{J} & n\mathbf{I} \end{bmatrix},$$

where the PSD inequality follows because the difference of the terms is the weighted Laplacian of the bipartite graph with weighted bipartite adjacency matrix $\mathbf{A}(\mathbf{x}, \mathbf{y}) - \mu(\mathbf{x}, \mathbf{y})\mathbf{J}$, which has non-negative entries. Now observe that the last term $\begin{bmatrix} n\mathbf{I} & -\mathbf{J} \\ -\mathbf{J} & n\mathbf{I} \end{bmatrix}$ is the (unweighted) Laplacian of the complete bipartite graph $K_{n,n}$, whose spectrum is $2n, n, 0$ with multiplicities $1, 2n-2$ and 1 respectively. The zero eigenvalue corresponds to the all-ones vector of length $2n$ and it is easy to see that indeed $(\mathbf{1}, -\mathbf{1})$ also lies in the kernel of $\nabla^2 f(\mathbf{x}, \mathbf{y})$. This shows that the non-zero eigenvalues of $\nabla^2 f(\mathbf{x}, \mathbf{y})$ are at least $n \cdot \mu(\mathbf{x}, \mathbf{y})$, and that it has a one-dimensional eigenspace corresponding to 0 , spanned by the vector $(\mathbf{1}, -\mathbf{1})$. Hence, $\nabla^2 f(\mathbf{x}, \mathbf{y}) \succeq \mu(\mathbf{x}, \mathbf{y}) \cdot n \cdot \mathbf{P}_V$. \blacktriangleleft

We now bound the smallest entry of the rescaled matrix. For this we use the following lemma (cf. [20, Lem. 6.2], [5, Cor. C.3 (arXiv)]) which bounds the variation norm of the scaling vectors $(\mathbf{x}^*, \mathbf{y}^*)$ of an exact scaling.

► **Lemma A.2.** *Let $\mathbf{A} \in [\mu, \nu]^{n \times n}$ and let $(\mathbf{x}^*, \mathbf{y}^*) \in \mathbb{R}^n \times \mathbb{R}^n$ be such that $\mathbf{A}(\mathbf{x}^*, \mathbf{y}^*)$ is exactly (\mathbf{r}, \mathbf{c}) -scaled. Then*

$$x_{\max}^* - x_{\min}^* \leq \ln \frac{\nu}{\mu} + \ln \frac{r_{\max}}{r_{\min}} \quad \text{and} \quad y_{\max}^* - y_{\min}^* \leq \ln \frac{\nu}{\mu} + \ln \frac{c_{\max}}{c_{\min}}.$$

► **Lemma A.3.** *Let $\mathbf{A} \in [\mu, \nu]^{n \times n}$ be an entrywise-positive matrix with $\|\mathbf{A}\|_1 = 1$ and let $f: V \subset \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}$ be the potential for this matrix as given in (1.2), where V is the orthogonal complement of $(\mathbf{1}_n, -\mathbf{1}_n)$. Let $(\mathbf{x}^*, \mathbf{y}^*) \in V$ be the unique minimizer of f in V . Then $\mu(\mathbf{x}^*, \mathbf{y}^*) \geq \frac{1}{n^2} \left(\frac{\mu}{\nu}\right)^3$. Moreover, for every $(\mathbf{x}, \mathbf{y}) \in V$ we have $\mu(\mathbf{x}, \mathbf{y}) \geq \mu(\mathbf{x}^*, \mathbf{y}^*)e^{-2\|(\mathbf{x}, \mathbf{y}) - (\mathbf{x}^*, \mathbf{y}^*)\|_\infty}$.*

Proof. By Lemma A.1 f is strictly convex on V . We also know that \mathbf{A} is exactly scalable. Hence f has a unique minimizer $(\mathbf{x}^*, \mathbf{y}^*)$. By Lemma A.2 we know that the variation norm of x^* and y^* are bounded by $\ln(\nu/\mu)$. Hence, for every $i, i', j, j' \in [n]$ we have

$$\left| \ln \left(\frac{e^{x_i^* + y_j^*}}{e^{x_{i'}^* + y_{j'}^*}} \right) \right| \leq |x_i^* - x_{i'}^*| + |y_j^* - y_{j'}^*| = 2\ln(\nu/\mu).$$

Therefore, the ratio between entries of $\mathbf{A}(\mathbf{x}^*, \mathbf{y}^*)$ is bounded:

$$\left| \frac{\mathbf{A}(\mathbf{x}^*, \mathbf{y}^*)_{ij}}{\mathbf{A}(\mathbf{x}^*, \mathbf{y}^*)_{i'j'}} \right| \leq \left| \frac{A_{ij}}{A_{i'j'}} \right| \left| \left(\frac{e^{x_i^* + y_j^*}}{e^{x_{i'}^* + y_{j'}^*}} \right) \right| \leq \frac{\nu}{\mu} e^{2\ln(\nu/\mu)} = \left(\frac{\nu}{\mu} \right)^3.$$

Since the sum of the entries of $\mathbf{A}(\mathbf{x}^*, \mathbf{y}^*)$ equals 1, this implies that the smallest entry of $\mathbf{A}(\mathbf{x}^*, \mathbf{y}^*)$ is at least $\mu(\mathbf{x}^*, \mathbf{y}^*) \geq \frac{1}{n^2} \left(\frac{\mu}{\nu}\right)^3$. Finally, for $(\mathbf{x}, \mathbf{y}) \in V$ and all $i, j \in [n]$ we have $A_{ij}e^{x_i + y_j} \geq A_{ij}e^{x_i^* + y_j^* - 2\|(\mathbf{x}, \mathbf{y}) - (\mathbf{x}^*, \mathbf{y}^*)\|_\infty}$, so taking the minimum over all i, j gives $\mu(\mathbf{x}, \mathbf{y}) \geq \mu(\mathbf{x}^*, \mathbf{y}^*)e^{-2\|(\mathbf{x}, \mathbf{y}) - (\mathbf{x}^*, \mathbf{y}^*)\|_\infty}$. \blacktriangleleft

Finally, to obtain a diameter bound for the set of points with a small gradient we will use the following (well-known) lemma.

► **Lemma A.4.** *Assume $g: \mathbb{R}^d \rightarrow \mathbb{R}$ is a C^2 convex function such that $\nabla g(\mathbf{0}) = \mathbf{0}$, and assume that for all $\mathbf{x} \in \mathbb{R}^d$ with $\|\mathbf{x}\|_\infty \leq r$, we have $\nabla^2 g(\mathbf{x}) \succeq \lambda I$. Then*

$$\|\nabla g(\mathbf{x})\|_2 \geq \lambda \|\mathbf{x}\|_2 \min(1, r/\|\mathbf{x}\|_\infty) \geq \lambda \min(\|\mathbf{x}\|_\infty, r).$$

In particular, to guarantee that $\|\mathbf{x}\|_\infty \leq C$ for $C \geq 0$, it suffices to show that $\|\nabla g(\mathbf{x})\|_2 < \lambda \min(C, r)$ (strict inequality is necessary as it forces $\min(\|\mathbf{x}\|_\infty, r) = \|\mathbf{x}\|_\infty$).

Proof. Fix $\mathbf{x} \in \mathbb{R}^n$ and consider $h: \mathbb{R} \rightarrow \mathbb{R}$ defined by $h(t) = g(t\mathbf{x})$. Then h is convex, $\partial_{t=0}h(t) = 0$ and $\partial_{t=s}^2 h(t) \geq 0$ for all $s \in \mathbb{R}$. Now assume for $s \in \mathbb{R}$ that $|s|\|\mathbf{x}\|_\infty \leq r$. Then

$$\partial_{t=s}^2 h(t) = \partial_{t=s}(Dg(t\mathbf{x})[\mathbf{x}]) = D^2 g(s\mathbf{x})[\mathbf{x}, \mathbf{x}] = \mathbf{x}^T \nabla^2 g(s\mathbf{x}) \mathbf{x} \geq \lambda \|\mathbf{x}\|_2^2.$$

For $s \geq 0$ this yields a lower bound on $\langle \nabla g(s\mathbf{x}), \mathbf{x} \rangle$ of the form

$$\begin{aligned} \langle \nabla g(s\mathbf{x}), \mathbf{x} \rangle &= \partial_{t=s} h(t) = \int_0^s \partial_{t=\tau}^2 h(t) d\tau \geq \int_0^{\min(s, r/\|\mathbf{x}\|_\infty)} \partial_{t=\tau}^2 h(t) d\tau \\ &\geq \int_0^{\min(s, r/\|\mathbf{x}\|_\infty)} \lambda \|\mathbf{x}\|_2^2 d\tau = \lambda \|\mathbf{x}\|_2^2 \min(s, r/\|\mathbf{x}\|_\infty), \end{aligned}$$

where the first inequality follows from the convexity of h . Setting $s = 1$ and using the Cauchy–Schwarz inequality gives $\|\nabla g(\mathbf{x})\|_2 \|\mathbf{x}\|_2 \geq \lambda \|\mathbf{x}\|_2^2 \min(1, r/\|\mathbf{x}\|_\infty)$ so

$$\|\nabla g(\mathbf{x})\|_2 \geq \lambda \|\mathbf{x}\|_2 \min(1, r/\|\mathbf{x}\|_\infty) \geq \lambda \|\mathbf{x}\|_\infty \min(1, r/\|\mathbf{x}\|_\infty) = \lambda \min(\|\mathbf{x}\|_\infty, r). \quad \blacktriangleleft$$

► **Lemma A.5.** *Let $\mathbf{A} \in [\mu, \nu]^{n \times n}$ be an entrywise non-negative matrix with $\|\mathbf{A}\|_1 = 1$ and let $f: V \subset \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}$ be the potential for this matrix as given in (1.2), where V is the orthogonal complement of $(\mathbf{1}_n, -\mathbf{1}_n)$. Let $(\mathbf{x}^*, \mathbf{y}^*)$ be the unique minimizer of f in V and let $0 < \delta < 1$. Let $(\mathbf{x}, \mathbf{y}) \in V$ be such that $\|\nabla f(\mathbf{x}, \mathbf{y})\|_2 < \delta \cdot \frac{1}{n} \left(\frac{\mu}{\nu}\right)^3 e^{-2}$. Then $\|(\mathbf{x}, \mathbf{y}) - (\mathbf{x}^*, \mathbf{y}^*)\|_\infty \leq \delta$.*

Proof. Lemma A.1 shows that $\nabla^2 f(\mathbf{x}, \mathbf{y}) \succeq n \cdot \mu(\mathbf{x}, \mathbf{y}) \cdot \mathbf{P}_V$, where \mathbf{P}_V is the orthogonal projector on V . Lemma A.3 shows that $\mu(\mathbf{x}, \mathbf{y}) \geq \mu(\mathbf{x}^*, \mathbf{y}^*) e^{-2\|(\mathbf{x}, \mathbf{y}) - (\mathbf{x}^*, \mathbf{y}^*)\|_\infty} \geq \frac{1}{n^2} \left(\frac{\mu}{\nu}\right)^3 e^{-2\|(\mathbf{x}, \mathbf{y}) - (\mathbf{x}^*, \mathbf{y}^*)\|_\infty}$. Hence, for (\mathbf{x}, \mathbf{y}) with $\|(\mathbf{x}, \mathbf{y}) - (\mathbf{x}^*, \mathbf{y}^*)\|_\infty \leq 1$, we have $\nabla^2 f(\mathbf{x}, \mathbf{y}) \succeq \frac{1}{n} \left(\frac{\mu}{\nu}\right)^3 e^{-2} \cdot \mathbf{P}_V$. It then follows from Lemma A.4 that if $\|\nabla f(\mathbf{x}, \mathbf{y})\|_2 < \delta \cdot \frac{1}{n} \left(\frac{\mu}{\nu}\right)^3 e^{-2}$, then $\|(\mathbf{x}, \mathbf{y}) - (\mathbf{x}^*, \mathbf{y}^*)\|_\infty \leq \delta$. \blacktriangleleft

Observe that for \mathbf{A}^σ the ratio between its largest and smallest entry is $\frac{b+1}{b-1} \leq 3$. This gives the following corollary, proving Lemma 2.8.

► **Corollary A.6.** *Let \mathbf{A}^σ be as in Section 2.1 and let f be the associated potential. Let $(\mathbf{x}^*, \mathbf{y}^*)$ be the unique exact scaling of \mathbf{A}^σ in V . If $(\mathbf{x}, \mathbf{y}) \in V$ is such that $\|\nabla f(\mathbf{x}, \mathbf{y})\|_2 < \frac{\delta}{27ne^2}$, then $\|(\mathbf{x}, \mathbf{y}) - (\mathbf{x}^*, \mathbf{y}^*)\|_\infty \leq \delta$.*

B Missing proofs for Section 3

B.1 Minimizing a second-order robust function

Before giving the proof of Theorem 3.5, we introduce the following notation. For a symmetric matrix \mathbf{H} and $\mathbf{b}, \mathbf{z} \in \mathbb{R}^n$, we denote

$$Q(\mathbf{H}, \mathbf{b}, \mathbf{z}) = \langle \mathbf{b}, \mathbf{z} \rangle + \frac{1}{2} \mathbf{z}^T \mathbf{H} \mathbf{z}.$$

We will use the following easily-verified properties of Q repeatedly.

► **Lemma B.1.** *For symmetric matrices \mathbf{H}, \mathbf{H}' and vectors $\mathbf{b}, \mathbf{b}', \mathbf{z}$, we have the following estimates:*

1. *If $\mathbf{H} \preceq \mathbf{H}'$, then $Q(\mathbf{H}, \mathbf{b}, \mathbf{z}) \leq Q(\mathbf{H}', \mathbf{b}, \mathbf{z})$.*
2. *If $\|\mathbf{H} - \mathbf{H}'\|_1 \leq \delta_a$, then $|Q(\mathbf{H}, \mathbf{b}, \mathbf{z}) - Q(\mathbf{H}', \mathbf{b}, \mathbf{z})| \leq \frac{1}{2}\delta_a \|\mathbf{z}\|_\infty^2$.*
3. *We have $|Q(\mathbf{H}, \mathbf{b}, \mathbf{z}) - Q(\mathbf{H}, \mathbf{b}', \mathbf{z})| = |\langle \mathbf{b} - \mathbf{b}', \mathbf{z} \rangle| \leq \|\mathbf{b} - \mathbf{b}'\|_1 \|\mathbf{z}\|_\infty$.*

Proof of Theorem 3.5. We follow the proof of [12, Thm. 3.4], and use their implementation of a k -oracle \mathcal{A} for $k = O(\log n)$, as detailed in Theorem 3.4. That is, \mathcal{A} takes as input an SDD matrix \mathbf{H} with $\tilde{O}(m)$ non-zero entries (off-diagonal entries ≤ 0) and a vector \mathbf{b} , and outputs a vector \mathbf{z} such that $\|\mathbf{z}\|_\infty \leq k$ and

$$Q(\mathbf{H}, \mathbf{b}, \mathbf{z}) \leq \frac{1}{2} \inf_{\|\mathbf{z}'\|_\infty \leq 1} Q(\mathbf{H}, \mathbf{b}, \mathbf{z}').$$

Then for $\mathbf{x}' = \mathbf{x} + \frac{1}{k}\Delta$, $\Delta = \mathcal{A}\left(\frac{4e^2}{3k^2}\mathbf{H}, \frac{1}{k}\mathbf{b}\right)$ we have

$$\begin{aligned} Q\left(\frac{4e^2}{3}\mathbf{H}, \mathbf{b}, \frac{1}{k}\Delta\right) &= Q\left(\frac{4e^2}{3k^2}\mathbf{H}, \frac{1}{k}\mathbf{b}, \Delta\right) \leq \frac{1}{2} \inf_{\|\mathbf{z}\|_\infty \leq 1} Q\left(\frac{4e^2}{3k^2}\mathbf{H}, \frac{1}{k}\mathbf{b}, \mathbf{z}\right) \\ &= \frac{1}{2} \inf_{\|\mathbf{z}\|_\infty \leq 1} Q\left(\frac{4e^2}{3}\mathbf{H}, \mathbf{b}, \mathbf{z}/k\right) = \frac{1}{2} \inf_{\|\mathbf{z}\|_\infty \leq \frac{1}{k}} Q\left(\frac{4e^2}{3}\mathbf{H}, \mathbf{b}, \mathbf{z}\right). \end{aligned}$$

Note that the second-order robustness of g implies that for $\tilde{\mathbf{x}} \in \mathbb{R}^n$ with $\|\mathbf{x} - \tilde{\mathbf{x}}\|_\infty \leq 1$, we have quadratic lower and upper bounds

$$Q\left(\frac{1}{e^2}\nabla^2 g(\mathbf{x}), \nabla g(\mathbf{x}), \tilde{\mathbf{x}} - \mathbf{x}\right) \leq g(\tilde{\mathbf{x}}) - g(\mathbf{x}) \leq Q\left(e^2\nabla^2 g(\mathbf{x}), \nabla g(\mathbf{x}), \tilde{\mathbf{x}} - \mathbf{x}\right). \quad (\text{B.1})$$

The remainder of the proof is structured as follows. We first compare quadratics involving $\nabla^2 g(\mathbf{x})$ and $\nabla g(\mathbf{x})$ to quadratics involving the approximations \mathbf{H} and \mathbf{b} in Equations (B.2) and (B.3). Using these estimates we then obtain a local progress bound over an ℓ_∞ -ball of radius $1/k$, see Equation (B.4). Finally, we convert this local bound into a more global estimate.

The properties of the approximate Hessian and gradient guarantee that

$$\begin{aligned} &Q\left(e^2\nabla^2 g(\mathbf{x}), \nabla g(\mathbf{x}), \tilde{\mathbf{x}} - \mathbf{x}\right) \\ &\leq Q\left(e^2\nabla^2 g(\mathbf{x}), \mathbf{b}, \tilde{\mathbf{x}} - \mathbf{x}\right) + \delta \\ &= Q\left(e^2\mathbf{H}'_m, \mathbf{b}, \tilde{\mathbf{x}} - \mathbf{x}\right) + Q\left(e^2\mathbf{H}'_a, \mathbf{b}, \tilde{\mathbf{x}} - \mathbf{x}\right) - \langle \mathbf{b}, \tilde{\mathbf{x}} - \mathbf{x} \rangle + \delta \\ &\leq Q\left(\frac{4e^2}{3}\mathbf{H}_m, \mathbf{b}, \tilde{\mathbf{x}} - \mathbf{x}\right) + Q\left(e^2\mathbf{H}_a, \mathbf{b}, \tilde{\mathbf{x}} - \mathbf{x}\right) + \frac{e^2}{2}\delta_a \|\tilde{\mathbf{x}} - \mathbf{x}\|_\infty^2 - \langle \mathbf{b}, \tilde{\mathbf{x}} - \mathbf{x} \rangle + \delta \\ &\leq Q\left(\frac{4e^2}{3}\mathbf{H}_m, \mathbf{b}, \tilde{\mathbf{x}} - \mathbf{x}\right) + Q\left(\frac{4e^2}{3}\mathbf{H}_a, \mathbf{b}, \tilde{\mathbf{x}} - \mathbf{x}\right) + \frac{e^2}{2}\delta_a \|\tilde{\mathbf{x}} - \mathbf{x}\|_\infty^2 - \langle \mathbf{b}, \tilde{\mathbf{x}} - \mathbf{x} \rangle + \delta \\ &= Q\left(\frac{4e^2}{3}\mathbf{H}, \mathbf{b}, \tilde{\mathbf{x}} - \mathbf{x}\right) + \frac{e^2}{2}\delta_a \|\tilde{\mathbf{x}} - \mathbf{x}\|_\infty^2 + \delta. \end{aligned} \quad (\text{B.2})$$

35:22 Improved Quantum Lower and Upper Bounds for Matrix Scaling

Furthermore, we also have the upper bound

$$\begin{aligned}
& Q\left(\frac{4e^2}{3}\mathbf{H}, \mathbf{b}, \tilde{\mathbf{x}} - \mathbf{x}\right) \\
&= Q\left(\frac{4e^2}{3}\mathbf{H}_m, \mathbf{b}, \tilde{\mathbf{x}} - \mathbf{x}\right) + Q\left(\frac{4e^2}{3}\mathbf{H}_a, \mathbf{b}, \tilde{\mathbf{x}} - \mathbf{x}\right) - \langle \mathbf{b}, \tilde{\mathbf{x}} - \mathbf{x} \rangle \\
&\leq Q(2e^2\mathbf{H}'_m, \mathbf{b}, \tilde{\mathbf{x}} - \mathbf{x}) + Q(2e^2\mathbf{H}_a, \mathbf{b}, \tilde{\mathbf{x}} - \mathbf{x}) - \langle \mathbf{b}, \tilde{\mathbf{x}} - \mathbf{x} \rangle \\
&\leq Q(2e^2\mathbf{H}'_m, \mathbf{b}, \tilde{\mathbf{x}} - \mathbf{x}) + Q(2e^2\mathbf{H}'_a, \mathbf{b}, \tilde{\mathbf{x}} - \mathbf{x}) + e^2\delta_a\|\tilde{\mathbf{x}} - \mathbf{x}\|_\infty^2 - \langle \mathbf{b}, \tilde{\mathbf{x}} - \mathbf{x} \rangle \\
&\leq Q(2e^2\mathbf{H}'_m, \mathbf{b}, \tilde{\mathbf{x}} - \mathbf{x}) + Q(2e^2\mathbf{H}'_a, \mathbf{b}, \tilde{\mathbf{x}} - \mathbf{x}) + e^2\delta_a\|\tilde{\mathbf{x}} - \mathbf{x}\|_\infty^2 - \langle \mathbf{b}, \tilde{\mathbf{x}} - \mathbf{x} \rangle \\
&= Q(2e^2\nabla^2 g(\mathbf{x}), \mathbf{b}, \tilde{\mathbf{x}} - \mathbf{x}) + e^2\delta_a\|\tilde{\mathbf{x}} - \mathbf{x}\|_\infty^2 \\
&\leq Q(2e^2\nabla^2 g(\mathbf{x}), \nabla g(\mathbf{x}), \tilde{\mathbf{x}} - \mathbf{x}) + e^2\delta_a\|\tilde{\mathbf{x}} - \mathbf{x}\|_\infty^2 + \delta.
\end{aligned} \tag{B.3}$$

Let \mathbf{v}_L and \mathbf{v}_U be the minimizers of quadratics over the ℓ_∞ -ball of radius $1/k$:

$$\mathbf{v}_L = \underset{\|\mathbf{v}\|_\infty \leq 1/k}{\operatorname{argmin}} Q\left(\frac{1}{e^2}\nabla^2 g(\mathbf{x}), \nabla g(\mathbf{x}), \mathbf{v}\right), \quad \mathbf{v}_U = \underset{\|\mathbf{v}\|_\infty \leq 1/k}{\operatorname{argmin}} Q(2e^2\nabla^2 g(\mathbf{x}), \nabla g(\mathbf{x}), \mathbf{v}).$$

Then by the guarantees of the k -oracle, we have

$$\begin{aligned}
Q\left(\frac{4e^2}{3}\mathbf{H}, \mathbf{b}, \frac{1}{k}\Delta\right) &\leq \frac{1}{2} \inf_{\|\mathbf{v}\|_\infty \leq 1/k} Q\left(\frac{4e^2}{3}\mathbf{H}, \mathbf{b}, \mathbf{v}\right) \\
&\leq \frac{1}{2} \inf_{\|\mathbf{v}\|_\infty \leq 1/k} (Q(2e^2\nabla^2 g(\mathbf{x}), \nabla g(\mathbf{x}), \mathbf{v}) + e^2\delta_a\|\mathbf{v}\|_\infty^2 + \delta) \\
&\leq \frac{1}{2} Q(2e^2\nabla^2 g(\mathbf{x}), \nabla g(\mathbf{x}), \mathbf{v}_U) + \frac{e^2\delta_a}{2k^2} + \frac{1}{2}\delta,
\end{aligned}$$

where the second inequality uses Equation (B.3), and the norm bounds $\|\mathbf{v}\|_\infty \leq 1/k \leq 1$ (to apply the inequality). Using the quadratic upper bound from Equation (B.1) on $g(\mathbf{x} + \frac{1}{k}\Delta) - g(\mathbf{x})$ and Equation (B.2), this yields

$$\begin{aligned}
g(\mathbf{x} + \frac{1}{k}\Delta) - g(\mathbf{x}) &\leq Q(e^2\nabla^2 g(\mathbf{x}), \nabla g(\mathbf{x}), \frac{1}{k}\Delta) \leq Q\left(\frac{4e^2}{3}\mathbf{H}, \mathbf{b}, \frac{1}{k}\Delta\right) + \frac{e^2}{2}\delta_a + \delta \\
&\leq \frac{1}{2} Q(2e^2\nabla^2 g(\mathbf{x}), \nabla g(\mathbf{x}), \mathbf{v}_U) + \frac{e^2\delta_a}{k^2} + \frac{3}{2}\delta,
\end{aligned}$$

We can then further upper bound this using

$$Q(2e^2\nabla^2 g(\mathbf{x}), \nabla g(\mathbf{x}), \mathbf{v}_U) \leq Q\left(2e^2\nabla^2 g(\mathbf{x}), \nabla g(\mathbf{x}), \frac{\mathbf{v}_L}{2e^4}\right) = \frac{1}{2e^4} Q\left(\frac{1}{e^2}\nabla^2 g(\mathbf{x}), \nabla g(\mathbf{x}), \mathbf{v}_L\right)$$

where the inequality uses that $\mathbf{v}_U = \underset{\|v\|_\infty \leq 1/k}{\operatorname{argmin}} Q(2e^2\nabla^2 g(\mathbf{x}), \nabla g(\mathbf{x}), v)$ and $\|\mathbf{v}_L\|_\infty \leq 1/k$. Collecting estimates, we obtain

$$g(\mathbf{x} + \frac{1}{k}\Delta) - g(\mathbf{x}) \leq \frac{1}{4e^4} Q\left(\frac{1}{e^2}\nabla^2 g(\mathbf{x}), \nabla g(\mathbf{x}), \mathbf{v}_L\right) + \frac{e^2\delta_a}{k^2} + \frac{3}{2}\delta. \tag{B.4}$$

We now convert this to a more global estimate. Let \mathbf{x}^* be a global minimizer of g . Set $\mathbf{y} = \mathbf{x} + \frac{1}{\max(kR_\infty, 1)}(\mathbf{x}^* - \mathbf{x})$, so that $\|\mathbf{y} - \mathbf{x}\|_\infty \leq \frac{1}{k}$. For the lower bound

$$g_L(\tilde{\mathbf{x}}) = g(\mathbf{x}) + Q\left(\frac{1}{e^2}\nabla^2 g(\mathbf{x}), \nabla g(\mathbf{x}), \tilde{\mathbf{x}} - \mathbf{x}\right)$$

on $g(\tilde{\mathbf{x}})$ we see that $g_L(\mathbf{x} + \mathbf{v}_L) \leq g_L(\mathbf{y}) \leq g(\mathbf{y})$ since $\mathbf{x} + \mathbf{v}_L$ minimizes $g_L \leq g$ over the ℓ_∞ -ball of radius $1/k$ around \mathbf{x} . By convexity of g we get

$$g(\mathbf{y}) = g\left(\mathbf{x} + \frac{1}{\max(kR_\infty, 1)}(\mathbf{x}^* - \mathbf{x})\right) \leq \left(1 - \frac{1}{\max(kR_\infty, 1)}\right)g(\mathbf{x}) + \frac{1}{\max(kR_\infty, 1)}g(\mathbf{x}^*)$$

so

$$g(\mathbf{x}) - g_L(\mathbf{x} + \mathbf{v}_L) \geq g(\mathbf{x}) - g(\mathbf{y}) \geq \frac{1}{\max(kR_\infty, 1)}(g(\mathbf{x}) - g(\mathbf{x}^*)).$$

Using this estimate in Equation (B.4), this gives

$$g(\mathbf{x}) - g\left(\mathbf{x} + \frac{1}{k}\Delta\right) \geq \frac{1}{4e^4 \max(kR_\infty, 1)}(g(\mathbf{x}) - g(\mathbf{x}^*)) - \left(\frac{e^2 \delta_a}{k^2} + \frac{3}{2}\delta\right),$$

which after rearranging and rewriting $\mathbf{x}' = \mathbf{x} + \frac{1}{k}\Delta$ reads

$$g(\mathbf{x}') - g(\mathbf{x}^*) \leq \left(1 - \frac{1}{4e^4 \max(kR_\infty, 1)}\right)(g(\mathbf{x}) - g(\mathbf{x}^*)) + \frac{e^2 \delta_a}{k^2} + \frac{3}{2}\delta. \quad \blacktriangleleft$$

B.2 Approximating the Hessian of the regularized potential

Proof of Lemma 3.6. The first point is easy to verify, as is the second point (the regularization term is always positive). For the third point, suppose we have a \mathbf{z} such that $\tilde{f}(\mathbf{z}) \leq \tilde{f}(0)$. Then

$$\frac{\varepsilon^2}{ne^B} \left(\sum_i (e^{x_i} + e^{-x_i}) + \sum_j (e^{y_j} + e^{-y_j}) \right) \leq f(\mathbf{0}) - f(\mathbf{z}) + \frac{\varepsilon^2}{ne^B} \cdot 4n \leq \ln(1/\mu) + \frac{4\varepsilon^2}{e^B}. \quad (\text{B.5})$$

where the last inequality follows from the potential bound $f(\mathbf{0}) - f^* \leq \ln(1/\mu)$ (which depends on $\|\mathbf{A}\|_1 \leq 1$; in general the upper bound is $\|\mathbf{A}\|_1 - 1 + \ln(1/\mu)$). Since each of the regularization terms is positive, we may restrict ourselves to a single term and see that $e^{x_i} + e^{-x_i} \leq \frac{e^B n \ln(1/\mu)}{\varepsilon^2} + 4n$, from which we may deduce

$$|x_i| \leq \ln \left(\frac{e^B n \ln(1/\mu)}{\varepsilon^2} + 4n \right) = B + \ln \left(\frac{n \ln(1/\mu)}{\varepsilon^2} + \frac{4n}{e^B} \right) \leq B + \ln \left(\frac{n \ln(1/\mu)}{\varepsilon^2} + 4n \right),$$

where the last inequality uses $e^B \geq 1$ (recall $B > 0$). The same upper bound holds for $|y_j|$.

For the last point, note that if $\mathbf{z}_\varepsilon = (\mathbf{x}, \mathbf{y})$, then $e^{x_i} + e^{-x_i} \leq 2e^B$ and similarly for y , so

$$\tilde{f}(\mathbf{z}_\varepsilon) \leq f(\mathbf{z}_\varepsilon) + \frac{\varepsilon^2}{ne^B} \cdot 4ne^B = f(\mathbf{z}_\varepsilon) + 4\varepsilon^2 \leq f^* + 5\varepsilon^2.$$

If such a \mathbf{z}_ε exists, then $f^* \leq \tilde{f}^* \leq \tilde{f}(\mathbf{z}_\varepsilon) \leq f^* + 5\varepsilon^2$. ◀