

Spectral Sparsification by Deterministic Discrepancy Walk*

Lap Chi Lau[†]

Robert Wang[‡]

Hong Zhou[§]

Abstract

Spectral sparsification and discrepancy minimization are two well-studied areas that are closely related. Building on recent connections between these two areas, we generalize the “deterministic discrepancy walk” framework by Pesenti and Vladu [SODA 23] for vector discrepancy to matrix discrepancy, and use it to give a simpler proof of the matrix partial coloring theorem of Reis and Rothvoss [SODA 20]. Moreover, we show that this matrix discrepancy framework provides a unified approach for various spectral sparsification problems, from stronger notions including unit-circle approximation and singular-value approximation to weaker notions including graphical spectral sketching and effective resistance sparsification. In all of these applications, our framework produces improved results with a simpler and deterministic analysis.

1 Introduction

The aim of this work is to simplify and unify some previous work on spectral sparsification and discrepancy minimization. For a better perspective of our results, we will first briefly review the main results and techniques in these two areas, and then see recent connections between them and also some recent motivations.

Spectral Sparsification Karger [Kar99] introduced the notion of a cut sparsifier of a graph, which is a sparse graph over its vertex set that approximately preserves the weight of all cuts. Benczúr and Karger [BK15] designed a non-uniform random sampling algorithm and proved that any graph on n vertices has a cut sparsifier with $O(n \log n)$ edges. This is a highly influential result that has many applications in designing fast graph algorithms.

Spielman and Teng [ST11] introduced a stronger notion called spectral sparsification. Given a graph G , a sparse graph H on the same vertex set is called an ϵ -spectral sparsifier of G if $(1 - \epsilon)L_G \preceq L_H \preceq (1 + \epsilon)L_G$ where L_G and L_H are the Laplacian matrices of G and H respectively¹. They proved that any graph on n vertices has an ϵ -spectral sparsifier with $O((n \text{ polylog } n)/\epsilon^2)$ edges, and used it to design an important algorithm for solving Laplacian equations [ST14].

The stronger notion of spectral sparsification admits the following linear algebraic formulation of the problem. Given vectors $v_1, \dots, v_m \in \mathbb{R}^n$ such that $\sum_{i=1}^m v_i v_i^\top = I_n$, find a sparse reweighing $s \in \mathbb{R}_{\geq 0}^m$ such that $\sum_{i=1}^m s(i) \cdot v_i v_i^\top \approx I_n$. Using this formulation, Spielman and Srivastava [SS11] proved that any graph on n vertices has an ϵ -spectral sparsifier with $O((n \log n)/\epsilon^2)$ edges, generalizing the result of Benczúr and Karger, by analyzing a natural non-uniform sampling algorithm using matrix concentration inequalities. Subsequently, Batson, Spielman and Srivastava [BSS12] proved the striking result that any graph on n vertices has an ϵ -spectral sparsifier with only $O(n/\epsilon^2)$ edges². Their algorithm is quite different from random sampling, based on an ingenious potential function and a deterministic incremental approach.

Furthermore, Marcus, Spielman and Srivastava [MSS15] extended the potential function argument significantly to prove Weaver’s conjecture [Wea04] about a discrepancy problem³: Given vectors v_1, \dots, v_m such that $\sum_{i=1}^m v_i v_i^\top = I_n$ and $\|v_i\|_2^2 \leq \epsilon$ for $1 \leq i \leq m$, there is a partitioning $S_1 \cup S_2 = [m]$ such that $(\frac{1}{2} - \sqrt{\epsilon})^2 \cdot I_n \preceq \sum_{i \in S_1} v_i v_i^\top \preceq (\frac{1}{2} + \sqrt{\epsilon})^2 \cdot I_n$ for $j \in \{1, 2\}$. The results and techniques for spectral sparsification and Weaver’s conjecture have far-reaching consequences with many subsequent works.

*The full version of the paper can also be accessed at <https://arxiv.org/abs/2408.06146>.

[†]Cheriton School of Computer Science, University of Waterloo, Canada. Email: lapchi@uwaterloo.ca.

[‡]Cheriton School of Computer Science, University of Waterloo, Canada. Email: robert.wang2@uwaterloo.ca

[§]Corresponding author, School of Mathematics and Statistics, Fuzhou University, China. Email: hong.zhou@fzu.edu.cn.

¹Note that a spectral sparsifier is a cut sparsifier but not necessarily vice versa.

²Until now, there is no alternative approach to prove the existence of a linear-sized cut sparsifier, without using the spectral sparsification formulation.

³Weaver proved that the discrepancy problem is equivalent to a major open problem in mathematics called the Kadison-Singer problem. We will describe its interpretation as a discrepancy minimization problem.

Discrepancy Minimization Discrepancy theory is a broad area with various settings [Cha00, Mat99]. We consider the combinatorial setting where there are m subsets S_1, \dots, S_m of the ground set $[n]$ with $m \geq n$, and the goal is to find a coloring $s \in \{\pm 1\}^m$ to minimize the maximum discrepancy defined as $\max_{i \in [m]} |\sum_{j \in S_i} s(j)|$. A famous result by Spencer [Spe85] proved that there exists a coloring with discrepancy $O(\sqrt{n \log(m/n)})$, beating the bound obtained by a random coloring. His proof is non-constructive, based on the pigeonhole principle, and it was a long-standing open question to design an efficient algorithm to find such a coloring.

Bansal [Ban10] made the breakthrough in obtaining the first polynomial time algorithm to find a coloring matching Spencer's result when $m = n$. Lovett and Meka [LM15] simplified Bansal's algorithm and fully recovered Spencer's result. Both algorithms in [Ban10, LM15] are based on some interesting random walks similar to Brownian motion with elegant analyses. Rothvoss [Rot17] provided a beautiful alternative algorithmic approach to recover Spencer's result using convex geometry, reducing the problem to lower bounding the volume of a convex body⁴ (a polytope in Spencer's setting). These results laid the foundation for many subsequent developments in algorithmic discrepancy theory in recent years.

The combinatorial discrepancy problem can also be formulated as a vector discrepancy problem: Given vectors $v_1, \dots, v_n \in \mathbb{R}^m$, find a coloring $s \in \{\pm 1\}^n$ to minimize $\|\sum_i s(i) \cdot v_i\|_\infty$. A natural generalization is the following matrix discrepancy problem: Given matrices $A_1, \dots, A_n \in \mathbb{R}^{m \times m}$, find a coloring $s \in \{\pm 1\}^n$ to minimize $\|\sum_i s(i) \cdot A_i\|_{\text{op}}$. The matrix Spencer conjecture⁵ [Zou12, Mek14] has generated much interest: Given symmetric matrices $A_1, \dots, A_n \in \mathbb{R}^{n \times n}$, if $\|A_i\|_{\text{op}} \leq 1$ for $1 \leq i \leq n$, then there is a coloring $s \in \{\pm 1\}^n$ with $\|\sum_i s(i) \cdot A_i\|_\infty = O(\sqrt{n})$. Recently, Bansal, Jiang, and Meka [BJM23] made significant progress towards resolving this conjecture and revealed interesting connections to random matrix theory and free probability [BBvH23].

Connections While the setting of matrix discrepancy minimization seems naturally similar to that of spectral sparsification⁶, there are no known concrete connections between the two areas until relatively recently.

In one direction, Reis and Rothvoss [RR20] were the first to use techniques in algorithmic discrepancy theory to construct linear-sized spectral sparsifiers. The main technical result in [RR20] is the following matrix partial coloring theorem⁷: Given symmetric matrices $A_1, \dots, A_m \in \mathbb{R}^{n \times n}$ that satisfies $\sum_{i=1}^m |A_i| \preceq I_n$, there is a partial fractional coloring $x \in [-1, +1]^m$ such that there are $\Omega(m)$ coordinates in $\{\pm 1\}$ and $\|\sum_{i=1}^m x(i) \cdot A_i\|_{\text{op}} \leq O(\sqrt{n/m})$. This matrix partial coloring theorem can be applied recursively to construct a sparse reweighting for spectral sparsification [RR20], which we will explain in Section 4. To prove the matrix partial coloring theorem, Reis and Rothvoss used the convex geometric approach and bounded the volume of the operator norm ball $\{x \in \mathbb{R}^m : \|\sum_{i=1}^m x(i) \cdot A_i\|_{\text{op}} \leq O(\sqrt{n/m})\}$. To lower bound the volume of the norm ball, they used the potential function in [BSS12] as a barrier function to show that a guided Gaussian random walk will stay in the norm ball with high probability. Recently, this approach was extended by Jambulapati, Reis and Tian [JRT24] to construct linear-sized spectral sparsifier that are degree preserving. Furthermore, Sachdeva, Thudi, and Zhao [STZ24] used the result in [BJM23] for the matrix Spencer problem to construct better spectral sparsifiers for Eulerian directed graphs.

In the other direction, Bansal, Laddha and Vempala [BLV22] and Pesenti and Vladu [PV23] used the potential functions in spectral sparsification for discrepancy minimization. Both of these algorithms can be understood as “deterministic walks”, tracking the discrepancy of a continuously evolving partial fractional coloring guided by a potential function. The potential function used in [BLV22] is similar to that in [BSS12], while the potential function used in [PV23] is based on a regularized optimization formulation developed by Allen-Zhu, Liao and Orecchia [AZLO15] for spectral sparsification. They both showed that their approach provides a unifying algorithmic framework to recover many best known results in discrepancy minimization.

It is emerging from these recent connections that the potential functions in [BSS12, AZLO15] and the Brownian/deterministic discrepancy walks in [Ban10, LM15, RR20, BLV22, PV23] are the two underlying components for both spectral sparsification and discrepancy minimization.

⁴If the volume of a convex body in \mathbb{R}^n is “large enough”, then one can obtain a point in the convex body with $\Omega(n)$ coordinates being $\{\pm 1\}$ by simply projecting a random Gaussian vector to the convex body.

⁵It is not difficult to see that it already generalizes Spencer's result when each matrix is a diagonal matrix.

⁶Note that Weaver's conjecture can be restated as finding a coloring $\{\pm 1\}^m$ such that $\|\sum_{i=1}^m s(i) \cdot v_i v_i^\top\|_{\text{op}} \leq O(\epsilon)$

⁷One advantage of this theorem is that it works for any symmetric matrices, not just for rank one matrices of the form vv^\top , and this will be useful for our applications.

Motivations Besides being a natural goal in developing a simple and unifying algorithmic framework for both spectral sparsification and discrepancy minimization, there are recent applications that require better techniques for constructing stronger notions of spectral sparsifiers. Motivated by problems related to solving directed Laplacian equations [CKP⁺17] and log-space derandomization of directed connectivity [AKM⁺20, APP⁺23], several new notions of spectral sparsification for directed graphs were proposed, including standard approximation [CKP⁺17], unit-circle approximation [AKM⁺20] and singular-value approximation [APP⁺23]. We will present the formal definitions and mention some properties of these notions in Section 2.3.

One common requirement of these notions is that the sparsifiers need to be degree preserving. It is not known how to adapt the standard techniques such as independent random sampling in [SS11] and the potential function based incremental algorithm in [BSS12, AZLO15] to satisfy this requirement. The constructions in [CKP⁺17, APP⁺23] and a recent improvement by Sachdeva, Thudi and Zhao are based on the short cycle decomposition technique developed in [CGP⁺23], which is a dependent rounding method that samples alternating edges in even cycles to ensure degree preservation.

There are motivations in replacing the cycle decomposition technique by discrepancy walks. An open problem is to determine whether there are linear-sized sparsifiers for these stronger notions of spectral sparsification [APP⁺23]. The cycle decomposition technique inherits some polylogarithmic losses and it is not even known how to use it to construct linear-sized spectral sparsifiers in the standard setting [BSS12], while the discrepancy approach was used successfully to construct linear-sized degree-preserving spectral sparsifiers [JRT24]. Another motivation is to develop a unifying approach for spectral sparsification and discrepancy minimization, to consolidate our understanding for more challenging problems in these fields.

1.1 Our Results We generalize the deterministic discrepancy walk approach by Pesenti and Vladu [PV23] for the vector discrepancy setting to the matrix discrepancy setting, and obtain a considerably simpler proof of the matrix partial coloring theorem of Reis and Rothvoss [RR20]. Moreover, we show that this provides a unifying framework to derive simpler and sparser constructions for various spectral sparsification problems. One additional advantage is that this gives the first deterministic algorithms for all these applications, which also leads to more elementary proofs and simpler calculations bypassing the matrix concentration inequalities and probabilistic analyses in previous work.

Matrix Partial Coloring and Matrix Sparsification The main technical result is a deterministic version of the matrix partial coloring theorem by Reis and Rothvoss [RR20], augmented with linear constraints as was done in [JRT24]. The linear constraints provide extra flexibility, such as preserving degrees exactly, that are crucial in all applications in this paper.

THEOREM 1.1. (DETERMINISTIC MATRIX PARTIAL COLORING) *Let $A_1, \dots, A_m \in \mathbb{R}^{n \times n}$ be symmetric matrices such that $\sum_{i=1}^m |A_i| \preceq I_d$. Let $\mathcal{C} \subseteq \mathbb{R}^m$ be a set of good partial fractional colorings defined as*

$$\mathcal{C} := \left\{ x \in \mathbb{R}^m : \left\| \sum_{i=1}^m x(i) \cdot A_i \right\|_{\text{op}} \leq 16 \sqrt{\frac{n}{m}} \right\}.$$

In addition, let $\mathcal{H} \subseteq \mathbb{R}^m$ be a linear subspace of dimension $c \cdot m$ for some constant $c \geq \frac{4}{5}$. There is a deterministic polynomial time algorithm that returns a partial fractional coloring $x \in [-1, 1]^m$ such that

$$x \in \mathcal{C} \cap \mathcal{H} \quad \text{and} \quad |\{i \in [m] : x_i = \pm 1\}| = \Omega(m).$$

We simplify the proof of Reis and Rothvoss' result both conceptually and technically. In [RR20], a random walk guided by the potential function in [BSS12] was used to lower bound the Gaussian volume of the operator norm ball \mathcal{C} ⁸, which implies the existence of a good partial coloring through the convex geometric argument in [Rot17]. In contrast, we directly use a deterministic discrepancy walk to construct a good partial coloring, bypassing all convex geometric concepts. Also, we used the regularized optimization framework developed in [AZLO15, PV23] instead of the closely related potential function in [BSS12]. This allows us to bound the

⁸This is actually an over-simplification of what was done in [RR20], which is only made precise in [JRT24] through nontrivial convex geometric arguments. The original statements in [RR20] are more involved.

matrix discrepancy through a standard step-by-step optimization process (like Newton steps), without manually shifting the barrier and dealing with some complicated higher order terms as in [RR20]. This framework is more systematic and leads to conceptually simpler analysis with considerably easier calculations. We will elaborate more about this point in Section 1.2.

By a similar reduction as in [RR20], we obtain a deterministic matrix sparsification result with an additional subspace requirement.

THEOREM 1.2. (DETERMINISTIC MATRIX SPARSIFICATION) *Given positive semidefinite matrices $A_1, A_2, \dots, A_m \in \mathbb{R}^{n \times n}$ such that $\sum_{i=1}^m A_i \preceq I_n$ and a linear subspace $\mathcal{H} \subseteq \mathbb{R}^m$ of dimension $m - O(n)$, there is a polynomial time deterministic algorithm to construct a sparse reweighting $s \in \mathbb{R}_{\geq 0}^m$ with $|\text{supp}(s)| = O(n/\epsilon^2)$ such that $s - \mathbb{1}_m \in \mathcal{H}$ and*

$$\left\| \sum_{i=1}^m s(i) \cdot A_i - \sum_{i=1}^m A_i \right\|_{\text{op}} \leq \epsilon.$$

We note that the earlier deterministic sparsification algorithms in [BSS12] and [AZLO15] are based on a greedy strategy that updates a single input in each iteration. It is not clear how to satisfy the additional subspace requirement in those frameworks, while our algorithm uses a different high-rank perturbation based strategy that makes it possible.

Stronger Notions of Spectral Sparsification We demonstrate how to use Theorem 1.2 to obtain improved constructions for stronger spectral sparsifiers. One easy corollary of Theorem 1.2 is that we can deterministically find degree-preserving ϵ -spectral sparsifiers of size $O(n/\epsilon^2)$, which derandomizes a result in [JRT24] (see Section 4.1).

The unit-circle (UC) approximation introduced in [AKM⁺20] satisfies the property that it preserves the spectrum of all the powers of the random walk matrix of a directed graph (see Section 2.3 for the formal statement), and it was used to study low-space Laplacian solvers for undirected and Eulerian directed graphs. For undirected graphs, this definition simplifies to the following (see Lemma 4.1): H is an ϵ -UC approximation of G if (i) $D_H = D_G$, (ii) $D_H - A_H \approx_\epsilon D_G - A_G$, and (iii) $D_H + A_H \approx_\epsilon D_G + A_G$, where D denotes the diagonal degree matrix, A denotes the adjacency matrix, and \approx_ϵ denotes the usual ϵ -spectral approximation of matrices. We show that a simple reduction to Theorem 1.2 gives linear-sized UC-sparsifiers in deterministic polynomial time, where such linear-sized sparsifiers were not known to exist previously⁹ and can be seen as a strengthening of the linear-sized spectral sparsifiers in [BSS12].

THEOREM 1.3. (LINEAR-SIZED UC-SPARSIFIERS OF UNDIRECTED GRAPHS) *There is a polynomial time deterministic algorithm to compute an ϵ -UC-approximation with $O(n/\epsilon^2)$ edges for any undirected graph on n vertices.*

The singular-value (SV) approximation introduced in [APP⁺23] is a strictly stronger notion than UC-approximation, and satisfies the stronger property that it is preserved under arbitrary products (see Section 2.3 for the formal statement), which has applications in approximating higher powers of directed random walks and solving directed Laplacian systems. Using the cycle decomposition technique in [CGP⁺23] and expander decomposition, it was proved in [APP⁺23] that any directed graph has an ϵ -SV-sparsifier (and hence an ϵ -UC-sparsifier) with $O((n \log^{12} n)/\epsilon^2)$ edges. We show that a simple reduction to Theorem 1.2 gives linear-sized ϵ -SV-sparsifier when the bipartite lift is an expander graph. This simplifies and improves the corresponding $O(n \log^2 n)$ bound obtained using the cycle decomposition technique in [APP⁺23], and implies the following improved bound for general graphs through the same standard expander decomposition technique as in [APP⁺23].

THEOREM 1.4. (IMPROVED SV-SPARSIFIERS OF DIRECTED GRAPHS) *There is a polynomial time deterministic algorithm to compute an ϵ -SV-sparsifier with $O((n \log^5 n)/\epsilon^2)$ edges for any directed graph on n vertices.*

⁹We remark that no explicit constructions of UC-sparsifiers for general directed graphs were given in [AKM⁺20], not even for general undirected graphs, as it is enough to only sparsify squares of graphs for their applications. In an earlier work [CCL⁺15] that studies similar properties for undirected graphs, there were also no explicit constructions for general undirected graphs, but just for squares of graphs.

Graphical Spectral Sketches and Effective Resistance Sparsifiers We also find applications of [Theorem 1.2](#) to obtain improved constructions for weaker notions of spectral sparsification. An ϵ -spectral sketch of a graph G is a randomized data structure that, for any vector x , preserves the quadratic form $x^\top L_G x$ within a $1 \pm \epsilon$ multiplicative factor with high probability [[ACK+16](#), [JS18](#)]. Chu, Gao, Peng, Sachdeva, Sawlani, and Wang [[CGP+23](#)] considered the notion of a graphical spectral sketch, a spectral sketch that is a reweighted subgraph of G . They proved the interesting result that any graph has a graphical ϵ -spectral sketch with only $O((n \text{ polylog } n)/\epsilon)$ edges, with a better dependency on ϵ than an ϵ -spectral sparsifier which requires $\Omega(n/\epsilon^2)$ edges [[BSS12](#)].

In this work, we consider a deterministic version of graphical ϵ -spectral sketch, which is a sparsifier that ϵ -approximates an input graph's Laplacian quadratic forms with respect to a specified set of constraint vectors \mathcal{K} (see [Section 5](#) for a formal definition). Using the discrepancy approach, we obtain a deterministic polynomial time algorithm to construct sparser graphical ϵ -spectral sketches.

THEOREM 1.5. (IMPROVED GRAPHICAL SPECTRAL SKETCHING) *There is a polynomial time deterministic algorithm to construct a graphical ϵ -spectral sketch with $O((n \log^3 n \cdot \max\{1, \sqrt{\log(|\mathcal{K}|/n)}\})/\epsilon)$ edges for any unweighted undirected graph on n vertices, any set of vectors \mathcal{K} in \mathbb{R}^n , and any $\epsilon > 0$.*

Graphical spectral sketches were also used to construct effective resistance sparsifiers in [[CGP+23](#)], which preserve all pairwise effective resistances between vertices in a graph. Note that this is also a notion weaker than spectral sparsification, as the effective resistance between two vertices i and j in G can be written as $b_{i,j}^\top L_G^\dagger b_{i,j}$ where $b_{i,j} := \mathbb{1}_i - \mathbb{1}_j$ and L_G^\dagger is the pseudo-inverse of the Laplacian matrix L_G . Following the reduction in [[CGP+23](#)], we can also extend [Theorem 1.5](#) to produce sparser effective resistance sparsifiers.

THEOREM 1.6. (IMPROVED EFFECTIVE RESISTANCE SPARSIFICATION) *There is a polynomial time deterministic algorithm to construct a reweighted subgraph \tilde{G} with $O((n \log^{3.5} n)/\epsilon)$ edges such that*

$$(1 - \epsilon) \cdot b_{i,j}^\top L_G^\dagger b_{i,j} \leq b_{i,j}^\top L_{\tilde{G}}^\dagger b_{i,j} \leq (1 + \epsilon) \cdot b_{i,j}^\top L_G^\dagger b_{i,j}$$

for all $i, j \in V(G)$ for any unweighted undirected graph G .

Both [Theorem 1.5](#) and [Theorem 1.6](#) improve the previous results by many $\log n$ factors: the graphical spectral sketches and the effective resistance sparsifiers in [[CGP+23](#)] have $\Omega((n \log^{16} n)/\epsilon)$ edges even assuming optimal short cycle decomposition and optimal expander decomposition.

1.2 Technical Overview The proof of [Theorem 1.1](#) is quite straightforward conceptually, once the potential function is fixed. Building on the regularized optimization framework for vector discrepancy by Pesenti and Vladu [[PV23](#)] and for spectral sparsification by Allen-Zhu, Liao, and Orecchia [[AZLO15](#)], we use the potential function

$$\Phi(x) := \max_{M \in \Delta_n} \langle A(x), M \rangle + \frac{2}{\eta} \text{tr}(M^{\frac{1}{2}}),$$

where $A(x) := \sum_{i=1}^m x(i) \cdot A_i$ is a shorthand of the current solution, $\Delta_n := \{M \succcurlyeq 0 \mid \text{tr}(M) = 1\}$ is the set of density matrices, and η is a parameter. Note that the first term is simply $\lambda_{\max}(A(x))$, the maximum eigenvalue of our current solution. The term $2 \text{tr}(M^{\frac{1}{2}})$ is known as the $\ell_{1/2}$ -regularizer, and the potential function $\Phi(x)$ can be seen as a regularized smooth proxy of $\lambda_{\max}(A(x))$. We remark that the optimizer M can be computed and the closed form characterization of $\Phi(x)$ is very similar to the potential function used in [[BSS12](#)]; see [Lemma 3.1](#) and [\(A.1\)](#) in the appendix.

To find the partial fractional coloring, we start from the initial point $x_0 = 0$. In each iteration $t \geq 1$, the natural step is to find a small perturbation $y \in \mathbb{R}^m$ and set $x_t = x_{t-1} + y$, so that $\|x_t\|_2 > \|x_{t-1}\|_2$ and $\Phi(x_t)$ is not much larger than $\Phi(x_{t-1})$. We prove in [Lemma 3.1](#) that

$$\Phi(x + y) - \Phi(x) \leq \text{tr}(MA(y)) + c\eta \text{tr}(M^{\frac{1}{2}}A(y)M^{\frac{1}{2}}A(y)M^{\frac{1}{2}}),$$

for some $|c| \leq 2$ when $\|y\|$ is small enough where M is the optimizer. This bound generalizes¹⁰ a similar bound ([Theorem 3.3](#)) in [[AZLO15](#)] to allow high-rank updates $A(y)$ that are neither positive semidefinite nor negative

¹⁰To be precise, we only deal with $\ell_{1/2}$ -regularizer in this paper, while [Theorem 3.3](#) in [[AZLO15](#)] is about the general ℓ_q -regularizer. However, the same ideas in this paper can be applied to handle general ℓ_q -regularizer, and thus fully generalizes [Theorem 3.3](#) in [[AZLO15](#)]. We omit the details as $\ell_{1/2}$ -regularizer is already enough for the purpose of this paper.

semidefinite, which is a crucial feature that we need to analyze our algorithm. We provide a simpler proof of this bound using elementary convexity argument and Taylor expansion (in which we use a second-order approximation from [RR20]), bypassing more advanced concepts (e.g., mirror descent and Bregman divergence) used in [AZLO15]; see Appendix A.

To choose the updated perturbation y , we use standard arguments in discrepancy minimization to restrict to the subspace that satisfies:

1. y has support only on “active” coordinates, those coordinates in x_{t-1} that are not in $\{\pm 1\}$ yet,
2. $y \in \mathcal{H}$ so that $x_t \in \mathcal{H}$ assuming $x_{t-1} \in \mathcal{H}$,
3. $y \perp x_{t-1}$ so that $\|x_t\|_2^2 = \|x_{t-1}\|_2^2 + \|y\|_2^2$ is increasing, in order to upper bound the number of iterations,
4. the linear term $\text{tr}(MA(y))$ is zero, and
5. the second order term $\text{tr}(M^{\frac{1}{2}}A(y)M^{\frac{1}{2}}A(y)M^{\frac{1}{2}})$ is small.

The nontrivial step is to bound the second order term, for which we use the idea in [PV23] to rewrite it as the quadratic form $y^T N y$ of a matrix, where $N_{ij} = \text{tr}(M^{\frac{1}{2}}A_i M^{\frac{1}{2}}A_j M^{\frac{1}{2}})$. With this formulation, we just need to argue that N has a large eigenspace with small eigenvalues. To do so, we show that a large principal submatrix of N has small trace (in which we borrow another lemma from [RR20]), from which we conclude the theorem by Cauchy interlacing.

For the applications in spectral sparsification, the linear subspace \mathcal{H} in the discrepancy framework provides much flexibility in incorporating additional constraints, such as the degree constraints for degree-preserving sparsifiers. We believe that this perturbation step in the discrepancy framework is a more versatile and powerful dependent rounding method than the cycle decomposition technique. We demonstrate this (together with the recent works [JRT24, STZ24]) by providing simpler and sparser constructions using the discrepancy approach, for all applications previously using the cycle decomposition technique, including degree-preserving sparsifiers, Eulerian sparsifiers, graphical spectral sketches, effective resistance sparsifiers, and singular-value sparsifiers. The reductions to matrix partial coloring for UC-approximation in Theorem 1.3 and SV-approximation in Theorem 1.4 are very simple and direct, while the reduction for graphical spectral sketches in Theorem 1.5 requires some problem-specific insights.

We emphasize that almost all technical ideas used in this work were known in the literature, but we put them together in a more streamlined manner, considerably simplifying many calculations as well as bypassing many advanced concepts (such as matrix concentration inequalities, probabilistic analyses, convex geometry, mirror descent and Bregman divergence, etc). Our proofs are basically self-contained without using any blackbox (except two basic lemmas from [RR20]), and are arguably easier than [BSS12, AZLO15] even in the basic setting (without keeping track of two potential functions). Also, it is the same approach that works for both discrepancy minimization (as shown in [PV23]) and spectral sparsification. We believe that the simplifications and the unification consolidate our understanding and will stimulate further progress.

2 Background

Given two functions f, g , we use $f \lesssim g$ to denote the existence of a positive constant $c > 0$, such that $f \leq c \cdot g$ always holds. For positive integers n , we use $[n]$ to denote the set $\{1, 2, \dots, n\}$.

2.1 Linear Algebra Throughout this paper, we use subscripts to index different objects, v_1, v_2, \dots for vectors and A_1, A_2, \dots for matrices. To avoid confusion, we write the i -th entry of a vector v as $v(i)$, and write the (i, j) -th entry of a matrix A as $A(i, j)$. Given matrices A_1, \dots, A_m and a vector $y \in \mathbb{R}^m$, we will use the shorthand $A(y) := \sum_{i=1}^m y(i) \cdot A_i$.

We write $\mathbf{1}$ as the all-one vector and write $\mathbf{1}_S$ as the indicator vector of a set S . For a vector $v \in \mathbb{R}^n$ and a set $S \subseteq [n]$, we write $v|_S$ as the $|S|$ -dimensional vector restricting v to the coordinates in S . Given a vector $v \in \mathbb{R}^n$, we write $\text{diag}(v) \in \mathbb{R}^{n \times n}$ as the diagonal matrix with entries from v on the diagonal.

Let $M \in \mathbb{R}^{n \times n}$ be a real symmetric matrix and $M = \sum_{i=1}^n \lambda_i v_i v_i^\top$ be its eigen-decomposition with real eigenvalues and orthonormal basis $\{v_i\}_{i=1}^n$. The largest eigenvalue of M is written as $\lambda_{\max}(M) := \max_{i \in [n]} \lambda_i$ and the operator norm of M is written as $\|M\|_{\text{op}} := \max_{i \in [n]} |\lambda_i|$. We denote $M^\dagger := \sum_{\lambda_i \neq 0} \lambda_i^{-1} v_i v_i^\top$ as the pseudoinverse of M and denote $|M| := \sum_{i=1}^n |\lambda_i| v_i v_i^\top$ as the matrix obtained by taking the absolute value over

all eigenvalues of M . We say that M is positive semidefinite, denoted by $M \succcurlyeq 0$, if all eigenvalues of M are non-negative. We denote $\Delta_n := \{M \succcurlyeq 0 \mid \text{tr}(M) = 1\} \subseteq \mathbb{R}^{n \times n}$ as the set of all n -dimensional density matrices, where $\text{tr}(M)$ is defined as the sum of the diagonal entries of M . It is well-known that $\text{tr}(M) = \sum_{i=1}^n \lambda_i$, and so the following inequality follows from a simple application of Cauchy-Schwartz.

LEMMA 2.1. $\text{tr}(M^{\frac{1}{2}}) \leq \sqrt{n}$ for any $M \in \Delta_n$.

The following three basic facts will also be used. One is that the maximum eigenvalue of a matrix A can be formulated as $\lambda_{\max}(A) = \max_{M \in \Delta_n} \langle A, M \rangle$. Another is that if $A \succcurlyeq 0$ then $\text{tr}(AB) \leq \|B\|_{\text{op}} \cdot \text{tr}(A)$. The last one is the well-known Cauchy interlacing theorem.

THEOREM 2.1. (CAUCHY INTERLACING THEOREM) *Let $A \in \mathbb{R}^{n \times n}$ be a symmetric matrix with eigenvalues $\alpha_1 \leq \dots \leq \alpha_n$ and $B \in \mathbb{R}^{m \times m}$ be a principal submatrix of A with eigenvalues $\beta_1 \leq \dots \leq \beta_m$. For any $1 \leq i \leq m$,*

$$\alpha_i \leq \beta_i \leq \alpha_{n-m+i}.$$

We will also need the following two lemmas from [RR20]. The first one can be established by applying the Cauchy-Schwarz inequality twice.

LEMMA 2.2. ([RR20, LEMMA 10]) *Let A, B, C be symmetric matrices with $A, B \succcurlyeq 0$. Then*

$$\text{tr}(ACBC) \leq \text{tr}(A|C|) \cdot \text{tr}(B|C|).$$

The second one is a second-order approximation of the trace inverse function under perturbations.

LEMMA 2.3. ([RR20, LEMMA 11]) *Let A, B be symmetric matrices with $A \succ 0$ and $\|\eta A^{-1}B\|_{\text{op}} \leq \frac{1}{2}$ for some $\eta > 0$. Then there is a value $c \in [-2, 2]$ so that*

$$\text{tr}((A - \eta B)^{-1}) = \text{tr}(A^{-1}) + \eta \text{tr}(A^{-1}BA^{-1}) + c\eta^2 \text{tr}(A^{-1}BA^{-1}BA^{-1}).$$

2.2 Spectral Graph Theory Let $G = (V, E, w)$ be an undirected graph with weight $w_e \in \mathbb{R}_+$ on each edge $e \in E$. For two vertices $u, v \in V$, we write $u \sim v$ to denote that u and v are neighbors. Given a vertex $v \in V$, we use $d_G(v) := \sum_{u: u \sim v} w_{uv}$ to denote its weighted degree, and we will refer to its number of neighbors as its combinatorial degree.

Let $A_G \in \mathbb{R}^{n \times n}$ be the adjacency matrix of G , with $A_G(u, v) = w_{uv}$ if $uv \in E$ and 0 otherwise. Let D_G be the diagonal degree matrix of G with $D_G(v, v) = d_G(v)$. The Laplacian matrix of G is defined as $L_G := D_G - A_G$. It is well-known that the Laplacian matrix is positive semidefinite, and can be written as the sum of rank-one matrices $L_G = \sum_{e \in E} w_e b_e b_e^T$, where $b_e = \mathbb{1}_u - \mathbb{1}_v$ is the (signed) incidence vector of an edge. For a connected graph G , the random walk matrix is defined as $W_G := D_G^{-1}A_G$, which is a row-stochastic matrix with row sums equal to one. The normalized Laplacian matrix is defined as $\mathcal{L}_G := D_G^{-1/2}L_G D_G^{-1/2}$, where all eigenvalues of \mathcal{L}_G are between 0 and 2. Note that \mathcal{L}_G and W_G are similar matrices with the same eigenvalues. We may drop the subscript G from our matrices when the context is clear.

One class of graphs that will be of particular importance for our sparsification results is the expander graphs. Let $\lambda_i(\mathcal{L})$ be the i -th smallest eigenvalue of \mathcal{L} . For any scalar λ , we say that G is a λ -expander if $\lambda_2(\mathcal{L}_G) \geq \lambda$. We will often show that sparsifiers are easier to construct for expander graphs, and then apply the following well-known expander decomposition result (see e.g. [KVV04]) to construct sparsifiers for general graphs.

THEOREM 2.2. (EXPANDER DECOMPOSITION) *Given an unweighted undirected graph G with n vertices, there is a polynomial time deterministic algorithm to decompose its edges into subgraphs G_1, G_2, \dots, G_k for some k , such that each G_i is an $\Omega(1/\log^2 n)$ -expander and each vertex is contained in at most $O(\log n)$ subgraphs.*

We will use the following simple facts about spectral properties of bipartite graphs.

FACT 2.1. *For any bipartite graph G on $2n$ vertices, $\lambda_{2n-i+1}(\mathcal{L}_G) = 2 - \lambda_i(\mathcal{L}_G)$ for $i \in [n]$.*

FACT 2.2. *A graph G is bipartite if and only if $\lambda_{\max}(\mathcal{L}_G) = 2$.*

We will also consider sparsifying random walks on directed graphs. Given a directed graph G with adjacency matrix A , let D_{in} and D_{out} be the diagonal in-degree and out-degree matrices of G respectively. The random walk matrix of G is defined as $W = D_{\text{out}}^{-1}A$. We say that G is Eulerian if $D_{\text{in}} = D_{\text{out}}$.

2.3 Spectral Sparsification for Directed Graphs We follow the presentation in [APP⁺23] to define the new notions of spectral sparsification for directed graphs [CKP⁺17, AKM⁺20, APP⁺23] in a unifying way. The following is a definition of matrix approximation for general non-symmetric matrices, in which the aim is to preserve $x^\top Ay$ for all vectors x and y (but not just quadratic form $x^\top Ax$ as in the symmetric case) with respect to some appropriate error matrices.

DEFINITION 2.1. (MATRIX APPROXIMATION) *Let $A, \tilde{A} \in \mathbb{C}^{n \times n}$ and $E, F \in \mathbb{C}^{n \times n}$ be positive semidefinite matrices. We say that \tilde{A} is an ϵ -matrix approximation of A with respect to error matrices E and F if any of the following equivalent conditions hold:*

1. $|x^*(A - \tilde{A})y| \leq \frac{\epsilon}{2}(x^*Ex + y^*Fy)$ for all $x, y \in \mathbb{C}^n$,
2. $\|E^{1/2}(A - \tilde{A})F^{1/2}\|_{\text{op}} \leq \epsilon$ and $\ker(E) \subseteq \ker((A - \tilde{A})^\top)$ and $\ker(F) \subseteq \ker(A - \tilde{A})$.

Here $\ker(M)$ denotes the (right) kernel of a matrix M .

Standard Approximation Cohen, Kelner, Peebles, Peng, Rao, Sidford and Vladu [CKP⁺17] introduced the following definition of spectral sparsification for directed graphs, and used it to design fast algorithms for solving directed Laplacian equations.

DEFINITION 2.2. (STANDARD APPROXIMATION OF MATRICES AND OF GRAPHS) *Let $A, \tilde{A} \in \mathbb{C}^{n \times n}$. We say that \tilde{A} is an ϵ -standard approximation of A if \tilde{A} is an ϵ -matrix-approximation of A with respect to the error matrices $E = F = D - \frac{1}{2}(A + A^*)$, where $D := \frac{1}{2}(D_A + D_{A^*})$ and D_A, D_{A^*} are the diagonal matrix of the row sums of A, A^* respectively.*

Let G and \tilde{G} be directed graphs with adjacency matrices A and \tilde{A} . We say that \tilde{G} is an ϵ -standard approximation of G if \tilde{A} is an ϵ -standard approximation of A . Note that $D = \frac{1}{2}(D_{\text{out}} + D_{\text{in}})$ in this case and thus the error matrix is exactly the Laplacian matrix of the underlying undirected graph of G .

Note that as the kernel of any Laplacian matrix contains the all-ones vector, the kernel condition in Definition 2.1 implies that standard approximation is always degree-preserving, as $(A - \tilde{A})\mathbf{1} = 0$ implies the outdegrees are the same and $(A - \tilde{A})^\top \mathbf{1} = 0$ implies the indegrees are the same. Thus, even for undirected graphs, it is a stronger definition than the spectral approximation of Spielman and Teng. Recently, Sachdeva, Thudi, and Zhao [STZ24] give a polynomial time algorithm to construct an ϵ -standard sparsifier with $\tilde{O}((n \log^2 n)/\epsilon^2)$ edges (ignoring poly log log n factors) for Eulerian directed graphs, using the discrepancy result in [BJM23] for the Matrix Spencer problem.

Unit-Circle Approximation Ahmadijad, Kelner, Murtagh, Peebles, Sidford, and Vadhan [AKM⁺20] introduced a stronger notion called unit-circle (UC) approximation, and used it to design deterministic low-space algorithm for estimating random walk probabilities of Eulerian directed graphs.

DEFINITION 2.3. (UNIT-CIRCLE APPROXIMATION) *Let G and \tilde{G} be directed graphs with adjacency matrices A and \tilde{A} respectively. We say that \tilde{G} is an ϵ -UC approximation of G if for all $z \in \mathbb{C}$ with $|z| = 1$, the matrix $z\tilde{A}$ is an ϵ -standard approximation of the matrix zA .*

It is immediate from the definition that UC approximation implies standard approximation by taking $z = 1$. Suppose \tilde{G} is an ϵ -UC approximation of G , and W and \tilde{W} are the random walk matrices of G and \tilde{G} respectively. An interesting property of UC approximation is that \tilde{W}^k is an $O(\epsilon)$ -UC approximation of W^k for all integers $k \geq 1$. In addition to preserving powers, UC approximation preserves a stronger property called the cycle lift of random walk matrices, which was a key property in analyzing the low-space Laplacian solver; see Lemma 1.4 in [AKM⁺20] for these two properties.

For undirected graphs, the adjacency matrix has only real eigenvalues, and it turns out, this means it suffices to check the condition in Definition 2.3 for only $z \in \{\pm 1\}$. Thus UC approximation admits a simple definition yet still carries strong properties that are not captured by standard spectral approximations.

LEMMA 2.4. (UC APPROXIMATION FOR UNDIRECTED GRAPHS [AKM⁺20, LEMMA 3.7]) *Let G and \tilde{G} be undirected graphs with adjacency matrices A and \tilde{A} respectively. Then \tilde{A} is an ϵ -UC-approximation of A if and only if \tilde{A} is an ϵ -standard approximation of A and $-\tilde{A}$ is an ϵ -standard-approximations of $-A$.*

For directed graphs, however, UC approximation is not easy to certify, and it was not known how to construct UC sparsifiers directly for arbitrary directed graphs. In [AKM⁺20], UC sparsifiers are only constructed for squares of Eulerian graphs, which is sufficient for the purpose of designing Eulerian Laplacian solvers using small space.

Singular Value Approximation More recently, Ahmadinejad, Peebles, Pyne, Sidford, and Vadhan [APP⁺23] proposed an even stronger notion called singular-value (SV) approximation, and used it to design nearly linear time algorithms for estimating stationary probabilities of general directed graphs.

DEFINITION 2.4. (SINGULAR VALUE APPROXIMATION) *Let G and \tilde{G} be directed graphs with adjacency matrices A and \tilde{A} respectively. We say that \tilde{G} is an ϵ -SV approximation of G if \tilde{A} is an ϵ -matrix approximation of A with respect to the error matrices $E = D_{\text{out}} - A D_{\text{in}}^\dagger A^\top$ and $F = D_{\text{in}} - A^\top D_{\text{out}}^\dagger A$, where D_{out} and D_{in} are the diagonal outdegree and indegree matrices of G . We may also say that \tilde{A} is an ϵ -SV approximation of A when it is clear from the context.*

It was proved in [APP⁺23] that SV approximation implies UC approximation and is strictly stronger. One useful property of SV approximation is being preserved under arbitrary matrix product operations: Given two sequences of directed random walk matrices W_1, \dots, W_k and $\tilde{W}_1, \dots, \tilde{W}_k$, if each \tilde{W}_i is an ϵ -SV approximation of W_i , then the matrix $\prod_{i=1}^k \tilde{W}_i$ is an $(\epsilon + O(\epsilon^2))$ -SV approximation of $\prod_{i=1}^k W_i$. Moreover, SV approximation is preserved under the following operations that allow for efficient constructions.

LEMMA 2.5. *SV approximation is preserved under the following operations:*

1. (Decomposition:) *If \tilde{A}_1 and \tilde{A}_2 are ϵ -SV-approximations of matrices A_1 and A_2 respectively, then $\tilde{A}_1 + \tilde{A}_2$ is an ϵ -SV approximation of $A_1 + A_2$.*
2. (Bipartite Lift:) *\tilde{A} is an ϵ -SV approximation of A if and only if the symmetric matrix*

$$\begin{bmatrix} 0 & \tilde{A} \\ \tilde{A}^\top & 0 \end{bmatrix} \text{ is an } \epsilon\text{-SV approximation of } \begin{bmatrix} 0 & A \\ A^\top & 0 \end{bmatrix}.$$

The first property says that SV approximation is linear, which is crucial for decomposition-based construction algorithms. The second property says that to construct an SV sparsifier of a directed graph $G = (V, \vec{E})$, it is equivalent to construct an SV sparsifier of its bipartite lift, which is an undirected bipartite graph. Thus, unlike for standard or UC approximation, SV sparsification of directed graphs can be reduced to SV sparsification of undirected bipartite graphs. This will allow us to apply the partial coloring algorithm in [Theorem 1.1](#) to construct SV sparsifiers for both directed and undirected graphs.

3 Deterministic Matrix Partial Coloring

The goal of this section is to prove [Theorem 1.1](#), with a proof overview given in [Section 1.2](#). We will first present the deterministic discrepancy walk framework used in previous work in [Section 3.1](#). Then, we will describe the potential function and use it to bound the maximum eigenvalue in [Section 3.2](#). Finally, we will describe the restricted subspace in [Section 3.3](#) and bound the second-order term using a spectral argument in [Section 3.4](#) to complete the proof.

3.1 The Deterministic Discrepancy Walk Framework We use the deterministic discrepancy walk framework in [LRR17, BLV22, PV23] for computing a partial coloring with small discrepancy. The idea is to use a potential function that bounds the discrepancy of our current solution to guide a deterministic walk in a suitable high dimensional subspace.

Deterministic Discrepancy Walk for Matrix Partial Coloring

Input: $A_1, \dots, A_m \in \mathbb{R}^{n \times n}$ such that $\sum_{i=1}^m |A_i| \preceq I_n$ and a potential function $\Phi : \mathbb{R}^m \mapsto \mathbb{R}_+$

Output: $x \in [-1, 1]^m$ such that x has small discrepancy $\|A(x)\|_{\text{op}}$ and many coordinates in $\{\pm 1\}$.

1. Initialization: Set $x_0 = 0_m$ as the initial point. Set $H_1 = [m]$ be the initial set of active coordinates. Set $t = 1$. Let $\alpha = 1/\text{poly}(m)$ be the maximal step size.
2. **While** $m_t := |H_t| > \frac{3}{4}m$ **do**
 - (a) Pick y_t to be a unit vector from an appropriate subspace $\mathcal{U}_t \subseteq \mathbb{R}^m$ satisfying $y_t \perp x_{t-1}$, $\text{supp}(y_t) \subseteq H_t$, and $\Phi(x_{t-1} + y_t) - \Phi(x_{t-1})$ is bounded.
 - (b) Let δ_t be the largest step size such that $\delta_t \leq \alpha$ and $x_{t-1} + \delta_t y_t \in [-1, 1]^m$.
 - (c) Update $x_t \leftarrow x_{t-1} + \delta_t y_t$. Update $t \leftarrow t + 1$.
 - (d) Update $H_t := \{i \in [n] \mid |x_{t-1}(i)| < 1\}$ be the set of fractional coordinates.
3. Return $x := x_T$ where T is the last iteration.

To bound the discrepancy of the final solution x_T , we use a potential function Φ such that $\Phi(x)$ is a smooth estimation of the discrepancy of the solution x . We then bound the final discrepancy by $\Phi(x_T) = \Phi(0) + \sum_{t=1}^T (\Phi(x_t) - \Phi(x_{t-1}))$. So, if $\Phi(0)$ is small and $\Phi(x_t) - \Phi(x_{t-1})$ is bounded for all t , then we can upper bound the discrepancy of x_T . The main task of the analysis under this framework then is to show that for a suitable potential function, and any $x_{t-1} \in [-1, 1]^m$, there is a large enough subspace \mathcal{U}_t such that for any unit vector $y_t \in \mathcal{U}_t$, the potential increase $\Phi(x_{t-1} + \delta_t y_t) - \Phi(x_{t-1})$ is small if the step size δ_t is at most α . This implies that if the dimension of the allowed subspace \mathcal{H} is also large enough, then there is always an update direction $y_t \in \mathcal{U}_t \cap \mathcal{H}$, which would imply that the final solution has small discrepancy and is in the allowed subspace \mathcal{H} as well.

3.2 Potential Function and Maximum Eigenvalue Bound Allen-Zhu, Liao, Orecchia [AZLO15] developed a regularized optimization framework to derive the spectral sparsification result in [BSS12] in a more principled way. Recall that the maximum eigenvalue $\lambda_{\max}(A)$ of a matrix A can be formulated as $\max_{M \in \Delta_n} \langle A, M \rangle$ where Δ_n is the set of density matrices. In this framework, the potential function is a regularized version of the maximum eigenvalue

$$\Phi(x) := \max_{M \in \Delta_n} \langle A(x), M \rangle - \frac{1}{\eta} \cdot \phi(M),$$

where the regularizer $\phi(M)$ is a nonpositive strongly convex function and η is a parameter controlling the contribution of the regularization term. They showed that the negative entropy regularizer $\phi(M) := \langle M, \log M \rangle$ can be used to obtain a deterministic algorithm to recover the $O((n \log n)/\epsilon^2)$ spectral sparsification result in [SS11], and the $\ell_{1/2}$ -regularizer $\phi(M) = -2 \text{tr}(M^{\frac{1}{2}})$ can be used to recover the $O(n/\epsilon^2)$ spectral sparsification result in [BSS12].

Potential Function Naturally, we set our potential function to be

$$(3.1) \quad \Phi(x) := \max_{M \in \Delta_n} \langle A(x), M \rangle + \frac{2}{\eta} \cdot \text{tr}(M^{\frac{1}{2}}),$$

which was also used in [AZLO15]. It follows from Lemma 2.1 that $\Phi(x)$ satisfies $\lambda_{\max}(A(x)) \leq \Phi(x) \leq \lambda_{\max}(A(x)) + 2\sqrt{n}/\eta$. For the final partial coloring $x_T = \sum_{t=1}^T \delta_t y_t$ returned by the general framework,

$$\lambda_{\max}(A(x_T)) \leq \Phi(x_T) = \Phi(0) + \sum_{t=1}^T (\Phi(x_{t-1} + \delta_t y_t) - \Phi(x_{t-1})) \leq \frac{2\sqrt{n}}{\eta} + \sum_{t=1}^T (\Phi(x_{t-1} + \delta_t y_t) - \Phi(x_{t-1})).$$

Maximum Eigenvalue and Operator Norm The potential function is used to upper bound $\lambda_{\max}(A(x))$, but the goal in Theorem 1.1 is to upper bound $\|A(x)\|_{\text{op}}$. A standard trick in discrepancy theory is to handle the

operator norm via the reduction

$$\|A(x)\|_{\text{op}} = \lambda_{\max} \left(\sum_{i=1}^m x(i) \cdot \begin{pmatrix} A_i & 0 \\ 0 & -A_i \end{pmatrix} \right).$$

This only increases the dimension of the matrices by a factor of two. We will abuse notation to still use A_i to denote the block diagonal matrix $\begin{pmatrix} A_i & 0 \\ 0 & -A_i \end{pmatrix}$ throughout this section.

We remark that this is an advantage of the framework by Reis and Rothvoss [RR20] that allows the input matrices A_1, \dots, A_m to be arbitrary symmetric matrices (not just for rank one symmetric matrices of the form vv^\top) so that the above reduction works. In previous work [BSS12, AZLO15] for spectral sparsification, two potential functions are used to keep track of the maximum eigenvalue and the minimum eigenvalue separately, and thus the algorithm and the analysis are more involved. This is one place where the discrepancy walk framework provides simplification even for the standard setting in [BSS12, AZLO15], as keeping track of only one potential function makes the algorithm and the analysis conceptually and technically simpler.

Maximum Eigenvalue Bound The next task is to bound the potential increase. The regret bound (Theorem 3.3) in [AZLO15] may serve this purpose. However, their bound has an additional assumption that the updates being either positive semidefinite or negative semidefinite, which does not apply to our matrices $\begin{pmatrix} A_i & 0 \\ 0 & -A_i \end{pmatrix}$. Using some observations from [RR20], we remove this restriction and obtain the following lemma.

LEMMA 3.1. (POTENTIAL INCREASE) *Given symmetric matrices $A_1, \dots, A_m \in \mathbb{R}^{n \times n}$ and $x \in \mathbb{R}^m$, the unique optimizer in (3.1) is $M = (uI_n - \eta A(x))^{-2}$ where $u \in \mathbb{R}$ is the unique value such that $M \in \Delta_n$. For any $y \in \mathbb{R}^m$, the change of potential function can be bounded by*

$$\Phi(x + y) - \Phi(x) \leq \text{tr}(MA(y)) + c\eta \text{tr} \left(M^{\frac{1}{2}} A(y) M^{\frac{1}{2}} A(y) M^{\frac{1}{2}} \right),$$

where $|c| \leq 2$ as long as $\left\| M^{\frac{1}{2}} \cdot \eta A(y) \right\|_{\text{op}} \leq \frac{1}{2}$.

Using elementary convexity argument and Taylor expansion, we present a simple proof of Lemma 3.1 in Appendix A, which avoids more advanced concepts in optimization (e.g., mirror descent and Bregman divergence) that were used in [AZLO15].

With Lemma 3.1, we set $x = x_{t-1}$, $y = \delta_t y_t$, and $M = (u_t I_n - \eta A(x_{t-1}))^{-2}$ for each iteration t to bound

$$(3.2) \quad \lambda_{\max}(A(x_T)) \leq \frac{2\sqrt{n}}{\eta} + \sum_{t=1}^T \left(\text{tr} \left(M_t A(\delta_t y_t) \right) + c_t \eta \text{tr} \left(M_t^{\frac{1}{2}} A(\delta_t y_t) M_t^{\frac{1}{2}} A(\delta_t y_t) M_t^{\frac{1}{2}} \right) \right).$$

Note that we can set the maximal step size of δ_t for all t to be $\alpha := \frac{1}{2\eta}$ to ensure that¹¹

$$(3.3) \quad \left\| M^{\frac{1}{2}} \cdot \eta A(\delta_t y_t) \right\|_{\text{op}} \leq \eta \left\| \sum_{i=1}^m \delta_t \cdot y_t(i) \cdot A_i \right\|_{\text{op}} \leq \eta \cdot \delta_t \cdot \|y_t\|_{\infty} \cdot \left\| \sum_{i=1}^m |A_i| \right\|_{\text{op}} \leq \eta \cdot \delta_t \leq \frac{1}{2},$$

where we used the fact that M is a density matrix, y_t is a unit vector and the assumption in Theorem 1.1 that $\sum_{i=1}^m |A_i| \preceq I_n$. Therefore, we can assume that $|c_t| \leq 2$ for all t when we apply Lemma 3.1.

3.3 Restricted Subspaces The key task is to find an appropriate unit vector y_t so that the potential increase in (3.2) is bounded. The essence of many results in discrepancy minimization is to argue that there are not too many bad directions, so that as long as the degree of freedom is large (that is, $m \gg n$) then there is a large subspace of good update directions.

To bound the potential increase in each iteration in (3.2), the nontrivial part is to bound the second order term. We use the idea in [PV23] to write the second order term as the quadratic form of a matrix N_t and

¹¹To look ahead, we will set $\eta \approx \sqrt{m}$ so that $\alpha \leq 1/\text{poly}(m)$.

restrict y_t to lie in the small eigenspace of N_t . Formally, let H_t be the active coordinates in the t -th iteration and $m_t = |H_t|$. The second order term in (3.2) can be written as

$$\text{tr} \left(M_t^{\frac{1}{2}} A(y_t) M_t^{\frac{1}{2}} A(y_t) M_t^{\frac{1}{2}} \right) = \sum_{i,j \in H_t} y_t(i) \cdot y_t(j) \cdot \text{tr} \left(M_t^{\frac{1}{2}} A_i M_t^{\frac{1}{2}} A_j M_t^{\frac{1}{2}} \right) = (y_t|_{H_t})^\top N_t (y_t|_{H_t}),$$

where $y_t|_{H_t} \in \mathbb{R}^{m_t}$ is obtained by restricting y_t to coordinates in H_t , and N_t is an $(m_t \times m_t)$ -dimensional matrix defined as

$$N_t := \left\{ \text{tr} \left(M_t^{\frac{1}{2}} A_i M_t^{\frac{1}{2}} A_j M_t^{\frac{1}{2}} \right) \right\}_{i,j \in H_t}.$$

Note that N_t is a positive semidefinite matrix¹². Let $N_t = \sum_{i=1}^{m_t} \lambda_i u_i u_i^\top$ be the eigenvalue decomposition of N_t with $0 \leq \lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_{m_t}$. In order to bound the second order term, we restrict $y_t|_{H_t}$ to lie in the eigenspace spanned by $\{u_1, \dots, u_{m_t/3}\}$. With these notations set up, we can formally define the good subspace $\mathcal{U} := U^0 \cap U^1 \cap U^2 \cap U^3$ for the unit update vector y_t where

$$\begin{aligned} U^0 &= \{y \in \mathbb{R}^m \mid y_i = 0 \ \forall i \notin H_t\}, \\ U^1 &= \{y \in \mathbb{R}^m \mid y \perp x_{t-1}\}, \\ U^2 &= \left\{ y \in \mathbb{R}^m \mid \text{tr}(M_t A(y)) = \sum_{i=1}^m y(i) \text{tr}(M_t A_i) = 0 \right\}, \\ U^3 &= \{y \in \mathbb{R}^m \mid y|_{H_t} \in \text{span}\{u_1, u_2, \dots, u_{m_t/3}\}\}. \end{aligned}$$

As discussed in the overview in Section 1.2, the subspace U^0 is to ensure that only active coordinates are updated, U^1 is to ensure that $\|x_t\|_2^2$ is monotone increasing in order to bound the number of iterations of the algorithm, U^2 is to ensure that the linear term in (3.2) is zero, and U^3 is to ensure that the second order term in (3.2) is “small”. As $m_t \geq \frac{3}{4}m$ when the algorithm has not terminated, it follows that

$$(3.4) \quad \dim(\mathcal{U}) \geq m_t - 2 - \frac{2m_t}{3} = \frac{m_t}{3} - 2 \geq \frac{m}{4} - 2.$$

The remaining task is to upper bound the eigenvalue of the low eigenspace in order to upper bound the second order term.

LEMMA 3.2. (LOW EIGENSPACE) *Given $A_1, \dots, A_m \in \mathbb{R}^{n \times n}$ such that $\sum_{i=1}^m |A_i| \preccurlyeq I_n$, any unit vector $y \in U^0 \cap U^3$ satisfies*

$$\text{tr} \left(M_t^{\frac{1}{2}} A(y) M_t^{\frac{1}{2}} A(y) M_t^{\frac{1}{2}} \right) \leq \frac{9\sqrt{n}}{m_t^2}.$$

We will use a spectral argument to prove Lemma 3.2 in the next subsection. In the rest of this subsection, we first assume Lemma 3.2 to finish the proof of Theorem 1.1.

Proof of Theorem 1.1 Given the input matrices A_1, \dots, A_m such that $\sum_{i=1}^m |A_i| \preccurlyeq I_n$ and the linear subspace \mathcal{H} , we apply the deterministic discrepancy walk algorithm for matrix partial coloring with the subspace $\mathcal{U}_t = \mathcal{U} \cap \mathcal{H}$. By (3.4),

$$\dim(\mathcal{U}_t) \geq \dim(\mathcal{U}) - \dim(\mathcal{H}^\perp) \geq \frac{m}{4} - 2 - \frac{m}{5} > 0$$

as long as $m = \Omega(1)$, thus in each iteration there is always a unit vector $y_t \in \mathcal{U}_t$. As shown in (3.3), by taking the maximal step size to be $\alpha = \frac{1}{2\eta}$, we can assume that $|c_t| \leq 2$ in (3.2), and so

$$\lambda_{\max}(A(x_T)) \leq \frac{2\sqrt{n}}{\eta} + \sum_{t=1}^T \left(\delta_t \text{tr}(M_t A(y_t)) + 2\eta \delta_t^2 \text{tr} \left(M_t^{\frac{1}{2}} A(y_t) M_t^{\frac{1}{2}} A(y_t) M_t^{\frac{1}{2}} \right) \right).$$

¹²To see this, write $N_t(i, j) = \text{tr}(M_t^{\frac{1}{2}} A_i M_t^{\frac{1}{2}} A_j M_t^{\frac{1}{2}}) = \langle M_t^{\frac{1}{4}} A_i M_t^{\frac{1}{4}}, M_t^{\frac{1}{4}} A_j M_t^{\frac{1}{4}} \rangle$, and so N_t can be written as $X^\top X$ where the i -th column of X is $\text{vec}(M_t^{\frac{1}{4}} A_i M_t^{\frac{1}{4}})$, the vectorization of the matrix $M_t^{\frac{1}{4}} A_i M_t^{\frac{1}{4}}$.

Since $y_t \in U^2$, the linear term is $\text{tr}(M_t A(y_t)) = 0$. As $y_t \in U^0 \cap U^3$, by [Lemma 3.2](#), the second order term is $\text{tr}(M_t^{\frac{1}{2}} A(y_t) M_t^{\frac{1}{2}} A(y_t) M_t^{\frac{1}{2}}) \leq 9\sqrt{n}/m_t^2$. Therefore,

$$\lambda_{\max}(A(x_T)) \leq \frac{2\sqrt{n}}{\eta} + \frac{18\eta\sqrt{n}}{m_t^2} \sum_{t=1}^T \delta_t^2 \leq \frac{2\sqrt{n}}{\eta} + \frac{32\eta\sqrt{n}}{m^2} \sum_{t=1}^T \delta_t^2,$$

where the last inequality holds as $m_t \geq 3m/4$ before the while loop terminates. Since $y_t \in U^1$ (such that $y_t \perp x_{t-1}$ for all $t \in [T]$) and $\|x_T\| \in [-1, 1]^m$, it follows that $m \geq \|x_T\|_2^2 = \sum_{t=1}^T \|\delta_t y_t\|_2^2 = \sum_{t=1}^T \delta_t^2$. Therefore, by setting $\eta = \frac{1}{4}\sqrt{m}$, we conclude that

$$\lambda_{\max}(A(x_T)) \leq \frac{2\sqrt{n}}{\eta} + \frac{32\eta\sqrt{n}}{m} \leq 16\sqrt{\frac{n}{m}}.$$

Finally, to show the polynomial runtime, we bound the number of iterations of the deterministic walk algorithm. Note that in step 2(c), the update either (i) freezes a new coordinate, or (ii) the squared length of the solution $\|x_t\|^2 = \|x_{t-1}\|^2 + \alpha^2$ increases by α^2 as $y_t \perp x_{t-1}$. Clearly, the number of the first type of iterations is at most m . The number of the second type of iterations is at most m/α^2 as $\|x_T\|^2 \leq m$. Therefore, by our choice of $\alpha = \frac{1}{2\eta} = \frac{2}{\sqrt{m}}$, the total number of iterations is at most $m/\alpha^2 + m = O(m^2)$.

3.4 Spectral Argument We prove [Lemma 3.2](#) in this subsection using a spectral argument. We remark that the calculations are similar to that in [\[RR20, page 13\]](#), but we adapt them to a slightly different setting.

For a unit vector $y \in U^0$, the restricted vector $y|_{H_t}$ is a unit vector in \mathbb{R}^{H_t} . Since $y \in U^3$, the second order term is bounded by the eigenvalue of N_t in the low eigenspace so that

$$\text{tr}\left(M_t^{\frac{1}{2}} A(y) M_t^{\frac{1}{2}} A(y) M_t^{\frac{1}{2}}\right) = (y|_{H_t})^\top N_t (y|_{H_t}) \leq \lambda_{m_t/3}(N_t).$$

To bound $\lambda_{m_t/3}(N_t)$, our idea is to upper bound the trace of a large principal submatrix \tilde{N}_t of N_t and to use Cauchy interlacing theorem to bound $\lambda_{m_t/3}$. Let

$$S = \left\{ i \in H_t \mid \text{tr}(M_t^{\frac{1}{2}} |A_i|) \geq \frac{3 \text{tr}(M_t^{\frac{1}{2}})}{m_t} \right\}$$

be the set of “large” active coordinates. Note that $\sum_{i \in H_t} \text{tr}(M_t^{\frac{1}{2}} |A_i|) \leq \text{tr}(M_t^{\frac{1}{2}}) \cdot \|\sum_{i=1}^m |A_i|\|_{\text{op}} \leq \text{tr}(M_t^{\frac{1}{2}})$ (see [Section 2.1](#) for the first inequality) and each $\text{tr}(M_t^{\frac{1}{2}} |A_i|) \geq 0$, so it follows from Markov’s inequality that $|S| \leq \frac{1}{3}m_t$. Let \tilde{N}_t be the principal submatrix of N restricted to the indices in $H_t - S$. Note that $\dim(\tilde{N}_t) \geq m_t - |S| \geq \frac{2}{3}m_t$. By Cauchy interlacing in [Theorem 2.1](#), it holds that $\lambda_{m_t/3}(N_t) \leq \lambda_{m_t/3}(\tilde{N}_t)$. To bound $\lambda_{m_t/3}(\tilde{N}_t)$, we will simply compute the trace of \tilde{N}_t and use an averaging argument. Applying [Lemma 2.2](#) to each diagonal entry of \tilde{N}_t , we have

$$\text{tr}(\tilde{N}_t) = \sum_{i \in H_t - S} \text{tr}(M_t A_i M_t^{\frac{1}{2}} A_i) \leq \sum_{i \in H_t - S} \text{tr}(M_t |A_i|) \cdot \text{tr}(M_t^{\frac{1}{2}} |A_i|).$$

Now, by the definition of S and the assumption that $\sum_{i=1}^m |A_i| \preceq I_n$ and [Lemma 2.1](#), we obtain that

$$\text{tr}(\tilde{N}_t) \leq \frac{3 \text{tr}(M_t^{1/2})}{m_t} \cdot \sum_{i \in H_t - S} \text{tr}(M_t |A_i|) \leq \frac{3 \text{tr}(M_t^{1/2})}{m_t} \cdot \text{tr}(M_t) \leq \frac{3\sqrt{n}}{m_t}.$$

Since $\dim(\tilde{N}_t) \geq \frac{2}{3}m_t$, the average value of the eigenvalue of \tilde{N}_t is $\text{tr}(\tilde{N}_t)/\dim(\tilde{N}_t) \leq \frac{9}{2}\sqrt{n}/m_t^2$. As \tilde{N}_t is positive semidefinite, by Markov’s inequality, at most half of the eigenvalues of \tilde{N}_t can be greater than $9\sqrt{n}/m_t^2$. Combining the inequalities, we conclude that

$$\text{tr}\left(M_t^{\frac{1}{2}} A(y) M_t^{\frac{1}{2}} A(y) M_t^{\frac{1}{2}}\right) \leq \lambda_{m_t/3}(N_t) \leq \lambda_{m_t/3}(\tilde{N}_t) \leq \lambda_{\dim(\tilde{N}_t)/2}(\tilde{N}_t) \leq \frac{9\sqrt{n}}{m_t^2}.$$

4 From Matrix Partial Coloring to Spectral Sparsification

In this section, we show how to apply the matrix partial coloring [Theorem 1.1](#) to construct spectral sparsifiers. In [Section 4.1](#), we prove [Theorem 1.2](#) using the same reduction as in [\[RR20\]](#), and obtain deterministic degree-preserving spectral sparsification as a corollary. In [Section 4.2](#) and [Section 4.3](#), we use [Theorem 1.2](#) to construct UC sparsifiers as stated in [Theorem 1.3](#) and SV sparsifiers as stated in [Theorem 1.4](#).

4.1 Spectral Sparsification with Linear Subspace Constraints The goal of this subsection is to prove [Theorem 1.2](#), a general matrix sparsification result with a linear subspace constraint. We emphasize that the reduction from matrix sparsification in [Theorem 1.2](#) to matrix partial coloring in [Theorem 1.1](#) is exactly the same as in [\[RR20\]](#), both the algorithm and the analysis. We include them for the completeness of this paper and for the inclusion of the linear subspace constraint.

Spectral Sparsification with Linear Subspace Constraint

Input: positive semidefinite matrices $A_1, A_2, \dots, A_m \in \mathbb{R}^{n \times n}$ such that $\sum_{i=1}^m A_i \preceq I_n$, a linear subspace $\mathcal{H} \subseteq \mathbb{R}^m$ of dimension $m - O(n)$, and target accuracy parameter ϵ .

Output: $s \in \mathbb{R}_{\geq 0}^m$ such that $|\text{supp}(s)| = \Omega(n/\epsilon^2)$ and $s - \vec{1} \in \mathcal{H}$ and $\|\sum_{i=1}^m (s(i) - 1)A_i\|_{\text{op}} \leq O(\epsilon)$.

1. Initialize $s_0(i) = 1$ for $i \in [m]$. Let $t = 1$.
2. **While** $m_t := |\text{supp}(s_{t-1})| > cn/\epsilon^2$ for some fixed constant c **do**
 - (a) Apply the deterministic discrepancy walk algorithm in [Theorem 1.1](#) to find a partial coloring $x_t \in [-1, 1]^m$ such that
 - i. $\|\sum_i x_t(i) \cdot s_{t-1}(i) \cdot A_i\|_{\text{op}} \leq 16\sqrt{\frac{n}{m_t}}$,
 - ii. $|\{i \in \text{supp}(s_{t-1}) \mid x_t(i) = \pm 1\}| = \Omega(m_t)$ and $x_t(i) = 0$ for all $i \notin \text{supp}(s_{t-1})$,
 - iii. $x_t \in \mathcal{H}_t := \{y \in \mathbb{R}^m \mid \text{diag}(s_{t-1}) \cdot y \in \mathcal{H}\} \cap \{y \in \mathbb{R}^m \mid y(i) = 0 \forall i \notin \text{supp}(s_{t-1})\}$.
 - (b) If there are more $x_t(i) = 1$ than $x_t(i) = -1$ then update $x_t \leftarrow -x_t$.
 - (c) Update $s_t(i) \leftarrow s_{t-1}(i) \cdot (1 + x_t(i))$ for all $i \in [m]$.
 - (d) $t \leftarrow t + 1$.
3. **Return:** $s = s_T$ where T is the last iteration.

The idea is to find a partial coloring x with small discrepancy (with a constant fraction of entries of x being -1) and use it to zero-out a constant fraction of the entries in s in each iteration. Informally, if $x \in \{\pm 1\}^m$ is a full coloring with small discrepancy, then we either double-up or zero-out an entry of s in each iteration.

Proof of [Theorem 1.2](#) We start by assuming that the spectral sparsification algorithm can successfully find a partial coloring x_t that meets all the requirements (i)-(iii) in Step 2(a) in each iteration. It is immediate that the final reweighting will have support size at most $O(n/\epsilon^2)$ by the design of the algorithm. Next, the final reweighting s would satisfy the subspace requirement $s - \vec{1} \in \mathcal{H}$ as

$$s_t = s_{t-1} + \text{diag}(s_{t-1}) \cdot x_t \text{ for all } t \quad \text{and} \quad s_0 = \vec{1} \quad \implies \quad s = \vec{1} + \sum_{t=1}^T \text{diag}(s_{t-1}) \cdot x_t$$

and each $\text{diag}(s_{t-1}) \cdot x_t \in \mathcal{H}$ by the requirement (iii) in Step 2(a) in each iteration. For the discrepancy guarantee, first note that Step 2(a)(ii) and Step 2(b) imply that x_t has $\Omega(m_t)$ coordinates with -1 in the support of s_{t-1} , and thus the support size m_t decreases by a constant factor in each iteration by Step 2(c) of the algorithm. Therefore, by a telescoping sum and the triangle inequality, the discrepancy is

$$(4.1) \quad \left\| A(s) - A(\vec{1}_m) \right\|_{\text{op}} \leq \sum_{t=1}^T \left\| \sum_{i=1}^m (s_t(i) - s_{t-1}(i)) \cdot A_i \right\|_{\text{op}} = \sum_{t=1}^T \left\| \sum_{i=1}^m x_t(i) \cdot s_{t-1}(i) \cdot A_i \right\|_{\text{op}} \lesssim \sum_{t=1}^T \sqrt{\frac{n}{m_t}} \lesssim \epsilon,$$

where the second last inequality follows by requirement (i) in Step 2(a) in each iteration, and the last inequality follows as m_t is a geometric sequence dominated by the last term when $m_t > cn/\epsilon^2$. Finally, for the polynomial running time, as the support size m_t decreases by a constant factor in each iteration, the while loop will terminate within $O(\log(\epsilon^2 m/n))$ iterations.

It remains to show that we can indeed find the desired partial coloring x_t in each iteration. We maintain that $\sum_{i=1}^m s_{t-1}(i) \cdot A_i \preceq 2I_n$ at each iteration t , which obviously holds in the first iteration as $s_0 = \vec{1}$. This implies that $\sum_{i=1}^m \frac{1}{2} s_{t-1}(i) \cdot A_i \preceq I_n$ satisfies the requirement of [Theorem 1.1](#). Note that the input subspace \mathcal{H}_t has dimension $m_t - O(n) \geq \frac{4}{5} m_t$ for $m_t = \Omega(n/\epsilon^2)$ when ϵ is smaller than a small enough constant. So we can apply [Theorem 1.1](#) with $\{\frac{1}{2} s_{t-1}(i) A_i\}_{i=1}^m$ as the input matrices and \mathcal{H}_t as the input subspace to find a partial coloring x_t that satisfies all the requirements in the matrix sparsification algorithm in polynomial time. The property $\sum_{i=1}^m s_{t-1}(i) \cdot A_i \preceq 2I_n$ is maintained by the same argument in (4.1) as long as ϵ is smaller than a small enough constant. This completes the proof of [Theorem 1.2](#).

REMARK 4.1. *Note that the input matrices in [Theorem 1.2](#) are assumed to be positive semidefinite (as was done in [\[RR20\]](#)) while the input matrices in [Theorem 1.1](#) are only assumed to be symmetric. We can relax the assumption in [Theorem 1.2](#) to be symmetric as well, by the simple trick of replacing the input symmetric matrices A_i by $\begin{pmatrix} A_i & 0 \\ 0 & |A_i| \end{pmatrix}$, so that the property $\sum_{i=1}^m s_{t-1}(i) \cdot |A_i| \preceq 2I_n$ can be maintained by the discrepancy bound and the above proof would go through.*

Degree-Preserving Spectral Sparsification We show an easy application of [Theorem 1.2](#) to construct degree-preserving spectral sparsifiers in deterministic polynomial time, derandomizing a result in [\[JRT24\]](#). Given an edge-weighted undirected graph $G = (V, E)$, construct a positive semidefinite matrix $A_i = L_G^{\dagger/2} b_{e_i} b_{e_i}^\top L_G^{\dagger/2}$ for each edge $e_i \in E$, so that $\sum_{i=1}^m A_i = I_n$ where $n = |V|$ and $m = |E|$. For the degree constraints, construct the subspace $\mathcal{H} := \{x \mid \sum_{u:u \sim v} x(uv) = 0 \ \forall v \in V\}$ to ensure that $\sum_{u:u \sim v} s(uv) \cdot w(uv) = \sum_{u:u \sim v} w(uv)$ for each vertex $v \in V$. Note that $\dim(\mathcal{H}) \geq m - n$. Therefore, we can apply [Theorem 1.2](#) with A_1, \dots, A_m and \mathcal{H} to obtain a deterministic polynomial time algorithm to construct linear-sized degree-preserving spectral sparsifiers.

4.2 Unit-Circle Sparsification for Undirected Graphs The goal of this subsection is to prove [Theorem 1.3](#) that there are linear-sized UC sparsifiers for undirected graphs. This is a nice application to illustrate our deterministic discrepancy walk framework, as it follows very easily from our matrix sparsification result, while it seems out of reach for the known techniques in [\[SS11, BSS12, AZLO15\]](#). Recall from [Section 2](#) that UC approximation enjoys some interesting property that was not satisfied by standard spectral sparsification, so [Theorem 1.3](#) can be seen as a strengthening of the classical results in [\[BSS12, AZLO15\]](#).

We mentioned in the introduction that \tilde{G} is an ϵ -UC approximation of G if (1) $D_{\tilde{G}} = D_G$, (2) $(1 - \epsilon)L_G \preceq L_{\tilde{G}} \preceq (1 + \epsilon)L_G$, and (3) $(1 - \epsilon)U_G \preceq U_{\tilde{G}} \preceq (1 + \epsilon)U_G$ where $U_G := D_G + A_G$ which is often called the unsigned Laplacian matrix of G . We check this claim formally in the following lemma.

LEMMA 4.1. *Let \tilde{G} be a degree-preserving reweighted subgraph of a connected undirected graph G . Then \tilde{G} is an ϵ -UC approximation of G if*

$$\left\| L_G^{\dagger/2} (L_G - L_{\tilde{G}}) L_G^{\dagger/2} \right\|_{\text{op}} \leq \epsilon \quad \text{and} \quad \left\| U_G^{\dagger/2} (U_G - U_{\tilde{G}}) U_G^{\dagger/2} \right\|_{\text{op}} \leq \epsilon.$$

Proof. Let the adjacency matrices of G and \tilde{G} be A and \tilde{A} respectively. By [Definition 2.3](#), \tilde{G} is an ϵ -UC approximation of G if and only if (i) \tilde{A} is an ϵ -standard approximation of A and (ii) $-\tilde{A}$ is an ϵ -standard approximation of $-A$. For undirected graphs, by [Definition 2.2](#), (i) means that \tilde{A} is an ϵ -matrix approximation of A with respect to the error matrix $D - A = L$, and (ii) means that $-\tilde{A}$ is an ϵ -matrix approximation of $-A$ with respect to the error matrix $D + A = U$.

To check (i), by [Definition 2.1](#), it is equivalent to $\|L_G^{\dagger/2} (A_G - A_{\tilde{G}}) L_G^{\dagger/2}\|_{\text{op}} \leq \epsilon$ and $\ker(L_G) \subseteq \ker(A_G - A_{\tilde{G}})$. Since $D_G = D_{\tilde{G}}$ as \tilde{G} is degree-preserving, the first condition becomes $\|L_G^{\dagger/2} (L_G - L_{\tilde{G}}) L_G^{\dagger/2}\|_{\text{op}} \leq \epsilon$, which is satisfied by the assumption. As G is connected and so $\ker(L_G) = \vec{1}$, the second condition reduces to checking $(A_G - A_{\tilde{G}})\vec{1} = 0$, which is also satisfied as \tilde{G} is degree-preserving.

To check (ii), by [Definition 2.1](#), it is equivalent to $\|U_G^{\dagger/2}(-A_G + A_{\tilde{G}})U_G^{\dagger/2}\|_{\text{op}} \leq \epsilon$ and $\ker(U_G) \subseteq \ker(A_G - A_{\tilde{G}})$. Since $D_G = D_{\tilde{G}}$ and $U_G = D_G + A_G$, the first condition is equivalent to $\|U_G^{\dagger/2}(U_G - U_{\tilde{G}})U_G^{\dagger/2}\|_{\text{op}} \leq \epsilon$, which is satisfied by the assumption. To check the second condition, we characterize the kernel of U_G . The kernel of U_G is non-empty if and only if U_G has a zero eigenvalue. Note that $U_G = 2D_G - L_G$, so it has a zero eigenvalue if and only if $D_G^{-\frac{1}{2}}U_G D_G^{-\frac{1}{2}} = 2I - \mathcal{L}_G$ has a zero eigenvalue if and only if \mathcal{L}_G has an eigenvalue of 2. Hence, by [Fact 2.2](#), U_G has a zero eigenvalue if and only if G is bipartite, and $\ker(U_G) = \text{span}(\vec{\mathbb{1}}_X - \vec{\mathbb{1}}_Y)$ where (X, Y) is the unique bipartition of G as G is connected. Therefore, the second condition reduces to checking $(A_G - A_{\tilde{G}})(\vec{\mathbb{1}}_X - \vec{\mathbb{1}}_Y) = 0$. Since \tilde{G} is a subgraph of G and thus (X, Y) is also a bipartition of \tilde{G} , it is straightforward to verify that $A_G(\vec{\mathbb{1}}_X - \vec{\mathbb{1}}_Y) = D_G(\vec{\mathbb{1}}_X - \vec{\mathbb{1}}_Y)$ and $A_{\tilde{G}}(\vec{\mathbb{1}}_X - \vec{\mathbb{1}}_Y) = D_{\tilde{G}}(\vec{\mathbb{1}}_X - \vec{\mathbb{1}}_Y)$, and thus the second condition is also satisfied because $D_G = D_{\tilde{G}}$. \square

Proof of Theorem 1.3 With [Lemma 4.1](#), it is easy to reduce UC approximation of undirected graphs to the matrix sparsification result in [Theorem 1.2](#). We already know how to do (1) and (2) simultaneously as it is just degree-preserving spectral sparsification. Note that we also know to do (3), as $U_G = \sum_{e \in E} \vec{b}_e \vec{b}_e^\top$ where $\vec{b}_{uv} := \vec{\mathbb{1}}_u + \vec{\mathbb{1}}_v$ is the unsigned edge-incidence vector, and so we can write (3) as a matrix sparsification problem when the input matrices sum to the identity matrix. To do (1), (2), and (3) simultaneously, we just need to do the standard block matrix trick, which is an advantage that the framework in [\[RR20\]](#) offered as it works for arbitrary rank symmetric matrices.

Given an undirected graph $G = (V, E)$, for each edge $e \in E$, construct a positive semidefinite matrix

$$A_e = \begin{bmatrix} L_G^{\dagger/2} \vec{b}_e \vec{b}_e^\top L_G^{\dagger/2} & 0 \\ 0 & U_G^{\dagger/2} \vec{b}_e \vec{b}_e^\top U_G^{\dagger/2} \end{bmatrix},$$

so that $\sum_{e \in E} A_e = I_{2n}$ where $n = |V|$. Also, construct the subspace $\mathcal{H} := \{x \mid \sum_{u:u \sim v} x(uv) = 0 \forall v \in V\}$ to ensure degree preservation as in the previous subsection. Applying [Theorem 1.2](#) with $\{A_e\}_{e \in E}$ and \mathcal{H} will give a deterministic polynomial time algorithm to construct a linear-sized degree-preserving sparsifier \tilde{G} , where the first block of the discrepancy guarantee implies that $\|L_G^{\dagger/2}(L_{\tilde{G}} - L_G)L_G^{\dagger/2}\|_{\text{op}} \leq \epsilon$ and the second block of the discrepancy guarantee implies that $\|U_G^{\dagger/2}(U_{\tilde{G}} - U_G)U_G^{\dagger/2}\|_{\text{op}} \leq \epsilon$. We conclude from [Lemma 4.1](#) that \tilde{G} is a linear-sized ϵ -UC sparsifier of G .

4.3 Singular Value Sparsification for Directed Graphs The goal of this section is to prove [Theorem 1.4](#) about SV sparsifiers for directed graphs. By the second item in [Lemma 2.5](#), the problem of constructing an SV sparsifier of a directed graph can be reduced to the problem of constructing an SV sparsifier of its bipartite lift which is an undirected graph. Henceforth, we will assume that G is an undirected bipartite graph.

By the definition of SV approximation in [Definition 2.4](#), we would like to approximate the adjacency matrix A with respect to the error matrix $D - AD^{-1}A$. Unlike in the case of UC-approximation, we do not know how to reduce the general problem to matrix sparsification where the input matrices sum to the identity matrix. However, we observe that if the bipartite graph is an expander graph, then we can reduce to matrix sparsification where the sum of input matrices has bounded spectral norm. This observation leads to linear-sized SV sparsifier for bipartite $\Omega(1)$ -expander.

THEOREM 4.1. (SV APPROXIMATION FOR BIPARTITE EXPANDERS) *Let G be a connected bipartite graph with $\lambda_2(\mathcal{L}_G) \geq \lambda$. Then there is a polynomial time deterministic algorithm to compute a (ϵ/λ) -SV sparsifier of G with $O(n/\epsilon^2)$ edges.*

We will use the following lemma for SV approximation which is an analog of [Lemma 4.1](#) for UC-sparsification. The proof is to show that $\ker(E) = \text{span}\{\vec{\mathbb{1}}_V, \vec{\mathbb{1}}_X - \vec{\mathbb{1}}_Y\}$ where (X, Y) is the unique bipartition of G , and the rest of the argument follows the same way as in the proof of [Lemma 4.1](#). We omit the straightforward proof.

LEMMA 4.2. *Let \tilde{G} be a degree-preserving reweighted subgraph of a connected bipartite graph G . Let $E = D_G - A_G D_G^{-1} A_G$. Then \tilde{G} is an ϵ -SV approximation of G if*

$$\|E^{\dagger/2}(L_G - L_{\tilde{G}})E^{\dagger/2}\|_{\text{op}} \leq \epsilon.$$

Proof of Theorem 4.1 Given a bipartite graph $G = (V, E)$ with $\lambda_2(\mathcal{L}_G) \geq \lambda$, for each edge $e \in E$, construct a positive semidefinite matrix $A_e = \lambda E^{\dagger/2} b_e b_e^\top E^{\dagger/2}$. We claim that $\sum_{e \in E} A_e \preceq I$. Using the identity $E = (D - A)D^{-1}(D + A) = (D + A)D^{-1}(D - A)$, we see that

$$\sum_e A_e = \lambda E^{\dagger/2} L_G E^{\dagger/2} = \lambda (D + A)^{\dagger/2} D (D + A)^{\dagger/2},$$

where we plug in $E^{\dagger/2} = (D + A)^{\dagger/2} D^{1/2} (D - A)^{\dagger/2}$ on the left and $E^{\dagger/2} = (D - A)^{\dagger/2} D^{1/2} (D + A)^{\dagger/2}$ on the right to obtain the second equality. Some simple manipulations show that

$$\frac{1}{\lambda} \left\| \sum_e A_e \right\| = \left\| (D + A)^{\dagger/2} D (D + A)^{\dagger/2} \right\| = \left\| D^{\frac{1}{2}} (D + A)^{\dagger} D^{\frac{1}{2}} \right\| = \left\| (D^{-\frac{1}{2}} (D + A) D^{-\frac{1}{2}})^{\dagger} \right\| = \left\| (2I - \mathcal{L}_G)^{\dagger} \right\|,$$

where the second equality is by $\|BB^*\| = \|B^*B\|$, the third equality is by $\|B\| = \|(B^\dagger)^\dagger\|$, and the final equality is by $D^{-\frac{1}{2}}AD^{-\frac{1}{2}} = I - \mathcal{L}$. Since the eigenvalues of \mathcal{L}_G of a connected bipartite graph G satisfy $0 = \lambda_1 < \lambda_2 \leq \dots \leq \lambda_{n-1} < \lambda_n = 2$, the claim follows as

$$\frac{1}{\lambda} \left\| \sum_e A_e \right\| = \left\| (2I - \mathcal{L}_G)^{\dagger} \right\| = \frac{1}{2 - \lambda_{n-1}(\mathcal{L}_G)} = \frac{1}{\lambda_2(\mathcal{L}_G)} \leq \frac{1}{\lambda} \implies \sum_e A_e \preceq I,$$

where we used that [Fact 2.1](#) in the last equality and the implication is because $\sum_e A_e$ is a positive semidefinite matrix.

Applying [Theorem 1.2](#) with $\{A_e\}_{e \in E}$ and the degree-constrained subspace $\mathcal{H} := \{x \mid \sum_{u:u \sim v} x(uv) = 0 \ \forall v \in V\}$ will give a deterministic polynomial time algorithm to construct a linear-sized degree preserving sparsifier \tilde{G} , where the discrepancy guarantee implies that $\lambda \|E^{\dagger/2}(L_{\tilde{G}} - L_G)E^{\dagger/2}\| \leq \epsilon$. We conclude from [Lemma 4.2](#) that \tilde{G} is a linear-sized (ϵ/λ) -SV sparsifier of G .

REMARK 4.2. *We note that some similar calculations were done in [APP⁺23, Lemma 4.3] and ours were inspired by theirs, but we would like to point out that their purpose was different and in particular not for the construction of sparsification algorithms.*

Proof of Theorem 1.4 To derive [Theorem 1.4](#), we will use expander decomposition as was done in [APP⁺23]. By [Theorem 2.2](#), every graph G can be decomposed into G_1, \dots, G_k such that each G_i is an $\Omega(1/\log^2 n)$ -expander and each vertex is contained in at most $O(\log n)$ subgraphs. For each G_i , we apply [Theorem 1.4](#) with error parameter $\epsilon' := \Omega(\epsilon/\log^2 n)$ to find a graph \tilde{G}_i which is an $O(\epsilon)$ -SV sparsifier of G_i , with $O((n_i \log^4 n)/\epsilon^2)$ edges where $n_i = |V(G_i)|$. By the linearity property of SV approximation in [Lemma 2.5](#), the graph $\tilde{G} := \tilde{G}_1 \cup \dots \cup \tilde{G}_k$ is an ϵ -SV sparsifier of G . Since each vertex in G is contained in at most $O(\log n)$ subgraphs, it follows that $\sum_{i=1}^k n_i = O(n \log n)$ and thus the total number of edges in \tilde{G} is $O((n \log^5 n)/\epsilon^2)$.

5 Graphical Spectral Sketch and Resistance Sparsifier

In this section, we show how to use the discrepancy framework to construct graphical spectral sketches in [Section 5.1](#) and effective resistance sparsifiers in [Section 5.2](#). We remark that there are some new technical ideas needed for proving [Theorem 1.5](#) for spectral sketches, while [Theorem 1.6](#) for resistance sparsifiers follows easily from a reduction to graphical spectral sketch in [CGP⁺23] and [Theorem 1.5](#). We assume that the input graphs are undirected and unweighted in this section.

5.1 Graphical Spectral Sketch The definition of spectral sketches in [ACK⁺16, JS18, CGP⁺23] is inherently probabilistic. To design deterministic algorithms, we formulate the following deterministic version of ϵ -spectral sketches.

DEFINITION 5.1. (DETERMINISTIC GRAPHIC SPECTRAL SKETCH) *Given a graph $G = (V, E)$ on n vertices and a set of vectors $\mathcal{K} \subseteq \mathbb{R}^n$, we say that a weighted graph \tilde{G} is an ϵ -graphical spectral-sketch with respect to \mathcal{K} if*

$$(1 - \epsilon)z^\top L_G z \leq z^\top L_{\tilde{G}} z \leq (1 + \epsilon)z^\top L_G z \quad \text{for all } z \in \mathcal{K}.$$

Spectral Sketch from Partial Coloring We reduce sparsification to partial coloring as in Section 4.1. In the following deterministic graphical sketch algorithm, the discrepancy requirement of the partial coloring in (5.1) is constructed so that the same argument as in Theorem 1.2 would go through to establish the spectral sketch guarantee.

Deterministic Spectral Sketch Algorithm

Input: an unweighted undirected $G = (V, E)$ on n vertices, a set of vectors $\mathcal{K} \in \mathbb{R}^n$, and a target accuracy parameter ϵ .

Output: a weighted undirected graph $\tilde{G} = (V, \tilde{E})$ with $|\tilde{E}| \leq n \cdot f(n)/\epsilon$ and $z^\top L_{\tilde{G}} z \approx_\epsilon z^\top L_G z \forall z \in \mathcal{K}$.

1. Initialize $s_0(e) = 1$ for $e \in E$. Let $t = 1$.
2. **While** $m_t := |\text{supp}(s_{t-1})| > nf(n)/\epsilon$ for some fixed function $f(n)$ **do**

(a) Find a partial coloring $x_t : E \rightarrow [-1, 1]$ such that

i. For any $z \in \mathcal{K}$,

$$(5.1) \quad \left| \sum_e x_t(e) \cdot s_{t-1}(e) \cdot \langle z, b_e \rangle^2 \right| \leq \frac{nf(n)}{m_t} \cdot z^\top L_G z$$

ii. $|\{i \in \text{supp}(s_{t-1}) \mid x_t(i) = \pm 1\}| = \Omega(m_t)$ and $x_t(i) = 0$ for all $i \notin \text{supp}(s_{t-1})$.

(b) If there are more $x_t(i) = 1$ than $x_t(i) = -1$ then update $x_t \leftarrow -x_t$.

(c) Update $s_t(e) \leftarrow s_t(e) = s_{t-1}(e) \cdot (1 + x_t(e))$ for all $e \in E$.

(d) $t \leftarrow t + 1$.

3. **Return:** \tilde{G} with edges weight $s_T(e)$ for $e \in E$ where T is the last iteration.

LEMMA 5.1. *Assuming that there is a deterministic polynomial time algorithm to find a partial coloring satisfying the requirements in Step 2(a) in each iteration, then the deterministic spectral sketch algorithm is a deterministic polynomial time algorithm that outputs an $O(\epsilon)$ -graphical spectral sketch with respect to \mathcal{K} with $O(nf(n)/\epsilon)$ edges.*

Proof. It is immediate that the graph \tilde{G} has only $O(nf(n)/\epsilon)$ edges by the design of the algorithm. We thus focus on the spectral sketch guarantee. Note that $z^\top L_{\tilde{G}} z = z^\top (\sum_e s_T(e) \cdot b_e b_e^\top) z$ and $z^\top L_G z = z^\top (\sum_e s_0(e) \cdot b_e b_e^\top) z$ for each $z \in \mathcal{K}$. Thus, by a telescoping sum and the triangle inequality,

$$|z^\top L_{\tilde{G}} z - z^\top L_G z| \leq \sum_t \left| z^\top \left(\sum_e x_t(e) \cdot s_{t-1}(e) \cdot b_e b_e^\top \right) z \right| \leq \sum_t \frac{nf(n)}{m_t} \cdot z^\top L_G z,$$

where in the first inequality we used $s_t(e) - s_{t-1}(e) = x_t(e) \cdot s_{t-1}(e)$ by Step 2(c) of the algorithm, and in the second inequality we used the discrepancy bound (5.1) in Step 2(a)(i) of the algorithm. Since m_t is a geometrically decreasing sequence (which follows from Step 2(a)(ii), 2(b), and 2(c) as was explained in the proof of Theorem 1.2), the sum on the right hand side is at most a constant factor of the last term, which is at most $\epsilon \cdot z^\top L_G z$ as $m_t > nf(n)/\epsilon$ in the final iteration. \square

Therefore, the spectral sketch problem is reduced to the partial coloring problem defined in (5.1), for which we need some new technical ideas as described below.

Vector Discrepancy Algorithm Note that the partial coloring in Step 2(a) is a vector discrepancy problem, rather than a matrix discrepancy problem as in Theorem 1.2. The problem of finding a low discrepancy signing $x \in \mathbb{R}^m$ to minimize its projection in a set of input directions is well studied in the literature. To keep our algorithm deterministic, we will make use of the deterministic discrepancy minimization algorithm of Levy, Ramadas and

Rothvoss in [LRR17]. We state their main result in a way that fits into the deterministic discrepancy walk algorithm¹³.

THEOREM 5.1. ([LRR17, THEOREM 1 AND LEMMA 6]) *Suppose we run the deterministic discrepancy walk algorithm with input vectors $a_1, \dots, a_k \in \mathbb{R}^m$ where $k \geq m$. In each iteration t , there exists a polynomial time computable subspace U_t of dimension at least $\frac{4}{5}m$ such that as long as the update direction $y_t \in U_t$ for all $t = 1, 2, \dots, T$, then x_T satisfies $|\langle a_i, x_T \rangle| \leq O(\|a_i\|_2 \sqrt{\log \frac{k}{m}})$ for all $1 \leq i \leq k$.*

The most straightforward way to apply [Theorem 5.1](#) to satisfy the vector discrepancy constraints in (5.1) is to define an $|E|$ -dimensional vector a_z for each $z \in \mathcal{K}$ where $a_z(e) := s_{t-1}(e) \cdot \langle z, b_e \rangle^2$ for each $e \in E$, so that the constraints in (5.1) become $|\langle a_z, x_t \rangle| \leq \frac{nf(n)}{m_t} \cdot z^\top L_G z$ for each $z \in \mathcal{K}$ so as to apply [Theorem 5.1](#). The problem with this approach is that the guarantee in [Theorem 5.1](#) could be much larger than our desired bound¹⁴. We remark that this straightforward approach can be used to obtain the bound $|\langle a, x_t \rangle| \leq \sqrt{\frac{nf(n)}{m_t}} \cdot z^\top L_G z$, and this implies a spectral sketch of size $O(nf(n)/\epsilon^2)$ but that does not improve the spectral sparsification result [BSS12]. The discrepancy bound in the form of (5.1) is crucial for the sparsifier size to have a dependency of $1/\epsilon$ rather than $1/\epsilon^2$.

Main Ideas We explain the main ideas roughly before we present the formal proof below. The starting observation is that by restricting x to the degree-preserving subspace that satisfies $\sum_{u:v \sim u} x(u, v) \cdot s(u, v) = 0$ for all $v \in V$, we can rewrite the left hand side of the discrepancy constraints in (5.1) as

$$\left| \sum_{uv \in E} x(u, v) \cdot s(u, v) \cdot \langle z, b_e \rangle^2 \right| = 2 \left| \sum_{uv \in E} x(u, v) \cdot s(u, v) \cdot z(u) \cdot z(v) \right|.$$

Then the idea is that we only need to sparsify edges from the dense parts of the graph, and so we freeze those edge variables that have large $s(u, v)$ or incident on low degree vertices. By doing so, we get a useful bound that $s(u, v)^2 \lesssim \frac{n^2}{m^2} \cdot d(u) \cdot d(v)$ where $d(v)$ is the degree of vertex v in the input graph G . This will imply that the 2-norm of the (rewritten) discrepancy constraint can be bounded by

$$\sqrt{\sum_{uv \in E} s(u, v)^2 \cdot z(u)^2 \cdot z(v)^2} \lesssim \frac{n}{m} z^\top D z \leq \frac{n}{\lambda m} z^\top L z,$$

where the last inequality¹⁵ follows when the input graph G is a λ -expander. This implies that we can apply [Theorem 5.1](#) to construct spectral sketches for $\Omega(1)$ -expander graphs with only $O(\frac{n}{\epsilon} \sqrt{\log \frac{k}{m}})$ edges. Finally, we use the standard expander decomposition technique to construct spectral sketches for general graphs with $O(n \log^{3.5} n / \epsilon)$ edges.

Restricted Subspaces With the main ideas presented, we describe our restricted subspaces precisely. In the τ -th iteration, we would like to find a partial coloring $x_\tau \in \mathbb{R}^m$ given the current reweighting $s := s_{\tau-1} \in \mathbb{R}^m$, where m is the number of edges in the input graph G .

¹³We briefly sketch a simplistic version of the proof in [LRR17] to illustrate how that fits into the deterministic discrepancy walk algorithm. To bound the discrepancy, they used a potential function based on the multiplicative weights update method, which is defined as $\Phi(x) := \sum_{i=1}^m \Phi_i(x) := \sum_{i=1}^m \exp(\lambda \cdot \langle \frac{a_i}{\|a_i\|}, x \rangle - \lambda^2)$. The subspace U_t defined in [LRR17, Section 2] guarantees that as long as the update direction y_t is chosen from U_t at each iteration, then (1) the total potential $\Phi(x_t)$ does not increase much [LRR17, Lemma 9] and (2) each $\Phi_i(x_t)$ contributes at most a $O(1/m)$ fraction to the total potential [LRR17, Lemma 10]. These conditions ensure that the final partial coloring satisfies the required discrepancy bound as stated. We remark that the constant $\frac{4}{5}$ is arbitrary and can be changed to any larger constant at most one without changing the result. We also remark that the framework in [Section 3](#) can also be used to recover a similar result in [Theorem 5.1](#) using the so called ℓ_q -regularizer in [AZLO15, PV23], but this derivation is omitted to keep the presentation simple.

¹⁴To see this, consider the following example where G is the complete graph, and $\mathcal{K} := \{b_{i,j}\}_{i,j \in V}$ with one constraint for each pair of vertices. Initially $s_0(e) = 1$ for each $e \in E$. Then, setting the input vectors $\{a_{i,j}\}_{i,j \in V}$ as above and applying [Theorem 5.1](#) gives a partial coloring x with discrepancy bound roughly $\|a_{i,j}\|_2 = \sqrt{\sum_e \langle b_{i,j}, b_e \rangle^4} \approx \sqrt{n}$ when G is the complete graph. On the other hand, the desired bound is $\frac{nf(n)}{m_t} \cdot z^\top L_G z \approx f(n)$ since $m \approx n^2$ and $b_{i,j}^\top L_G b_{i,j} = 2n - 2$ in the complete graph. We are interested in the regime where $f(n) = \text{polylog}(n)$ and so the bound returned by [Theorem 5.1](#) using this straightforward approach is too weak.

¹⁵This is not precise for the simplicity of the presentation. We will see the correct version below.

1. The first subspace is the degree-preserving subspace

$$U^1 := \left\{ x \in \mathbb{R}^m \mid \sum_{u:u\sim v} x(u,v) \cdot s(u,v) = 0 \quad \forall v \in V \right\},$$

which ensures that the weighted degrees of s_τ are the same as that in $s_{\tau-1}$ by Step 2(c) of the algorithm.

2. The second subspace is to restrict to the dense parts of the graph. For each vertex v , let $d_s(v) := |\{e \in \text{supp}(s) \mid v \in e\}|$ be the degree of v in the support of s , and $d(v)$ be the degree of v in the input unweighted graph G . Let $E_v := \{uv \in E \mid s(u,v) > 10 \cdot d(v)/d_s(v)\}$ be the set of high weight edges incident on v . Define $E_0 := \cup_{v \in V} E_v$. Also let $V_{\text{low}} := \{v \mid d_s(v) \leq \frac{1}{10}m/n\}$ be set of vertices with low degree in the support of s . And define E_1 to be the set of edges incident to some vertices in V_{low} . The second subspace will be used to freeze the edges in $E_0 \cup E_1$:

$$U^2 = \{x \in \mathbb{R}^m \mid x(e) = 0 \quad \forall e \in E_0 \cup E_1\}.$$

Since we maintain that $\sum_{u:u\sim v} s(u,v) = d(v)$ for all $v \in V$, we see that $|E_v| \leq \frac{1}{10} \cdot d_s(v)$ by an averaging argument, and thus $|E_0| \leq \frac{1}{10} \sum_{v \in V} d_s(v) = \frac{1}{5}m$. And $|E_1| \leq \frac{1}{10}m$ as there are at most n vertices. This implies that $\dim((U^2)^\perp) \leq \frac{3}{10}m$, the complement of the subspace U^2 has dimension at most $\frac{3}{10}m$.

3. The third set of subspaces is for the discrepancy constraints. For each $z \in \mathcal{K}$, we first define a shifted vector \bar{z} as follows. Let $c_z := \sum_v d(v) \cdot z(v) / \sum_v d(v)$ be the center of z with respect to the degrees. Define $\bar{z} = z - c_z \mathbf{1}$. We note that \bar{z} is constructed so that it satisfies the property $\sum_v d(v) \cdot \bar{z}(v) = 0$, which will be used in the spectral argument for expander graphs. Now, for each $z \in \mathcal{K}$, define a vector $a_z \in \mathbb{R}^m$ where

$$a_z(u,v) = \begin{cases} s(u,v) \cdot \bar{z}(u) \cdot \bar{z}(v) & \text{if } uv \in E_s \\ 0 & \text{otherwise} \end{cases}$$

where $E_s := \text{supp}(s) \setminus (E_0 \cup E_1)$. Then the third set of subspaces is defined as

$$U_t^3 := \text{the subspace used in } t\text{-th iteration in Theorem 5.1 when vectors } \{a_z\}_{z \in \mathcal{K}} \text{ are given as input.}$$

Note that there is one subspace for each iteration t of the partial coloring algorithm.

The final set of restricted subspace is defined as $\mathcal{H}_t := U^0 \cap U^1 \cap U_t^3$, for each iteration t of the partial coloring algorithm. By Theorem 5.1, $\dim(U_t^3) \geq \frac{4}{5}m$, and so

$$\dim(\mathcal{H}_t) \geq \dim(U_t^3) - \dim((U^0)^\perp) - \dim((U^1)^\perp) \geq \frac{4m}{5} - n - \frac{3m}{10} = \frac{m}{2} - n,$$

which is large enough as long as $m \geq \Omega(n)$.

Spectral Sketches for Expander Graphs

We are ready to construct spectral sketches for expander graphs.

THEOREM 5.2. *Let G be a λ -expander on n vertices. For any set of vectors \mathcal{K} with $|\mathcal{K}| \geq n$, there is a deterministic polytime algorithm to construct a deterministic ϵ -spectral-sketch with $O\left(\frac{n\sqrt{\log(|\mathcal{K}|/n)}}{\lambda\epsilon}\right)$ edges.*

Proof. We implement the partial coloring subroutine required by Step 2(a) of the deterministic graphical sketch algorithm with $f(n) = \frac{1}{\lambda} \sqrt{\log \frac{|\mathcal{K}|}{m}}$. Suppose we are in the τ -th iteration of the deterministic graphical sketch algorithm. We let $s := s_{\tau-1}$ be the current reweighting, and let $x := x_\tau$ be the partial coloring that we would like to find. For the partial coloring subroutine, we run a deterministic discrepancy walk algorithm (e.g. the one in [LRR17]) with the restricted subspace \mathcal{H}_t in the t -th iteration, and use the output of the algorithm as our partial coloring x for this τ -th iteration. The main task is to check the discrepancy requirement in (5.1) is

satisfied. For each $z \in \mathcal{K}$, since $x \in U^1$, the left hand side of (5.1) can be rewritten as

$$\begin{aligned} & \sum_{e \in E} x(e) \cdot s(e) \cdot \langle b_e, z \rangle^2 \\ &= \sum_{uv \in E} x(u, v) \cdot s(u, v) \cdot (z(u) - z(v) - c_z + c_z)^2 \\ &= \sum_{v \in V} (z(v) - c_z)^2 \sum_{u: u \sim v} x(u, v) \cdot s(u, v) - 2 \sum_{uv \in E} x(u, v) \cdot s(u, v) \cdot (z(u) - c_z) \cdot (z(v) - c_z) \\ &= -2 \sum_{uv \in E} x(u, v) \cdot s(u, v) \cdot \bar{z}(v) \cdot \bar{z}(v). \end{aligned}$$

Since the update direction y_t in the partial coloring subroutine is in U_t for all t , it follows that x satisfies

$$\left| \sum_{e \in E} x(e) \cdot s(e) \cdot \langle z, b_e \rangle^2 \right| = 2 \left| \sum_{uv \in E} x(u, v) \cdot s(u, v) \cdot \bar{z}(v) \cdot \bar{z}(v) \right| = 2 |\langle a_z, x \rangle| \lesssim \|a_z\|_2 \sqrt{\log \frac{|\mathcal{K}|}{m}},$$

where the second equality is by the definition of a_z and U^2 and the inequality is by [Theorem 5.1](#).

It remains to bound $\|a_z\|_2$. Note that for each edge $uv \in E_s$, it is not in E_u or E_v and so $s(u, v) \lesssim \min \left\{ \frac{d(v)}{d_s(v)}, \frac{d(u)}{d_s(u)} \right\}$. Also, for each edge $uv \in E_s$, it is not in E_1 and so $d_s(u), d_s(v) \gtrsim \frac{m}{n}$. These imply that

$$s(u, v)^2 \lesssim \min \left\{ \frac{d(v)}{d_s(v)}, \frac{d(u)}{d_s(u)} \right\}^2 \leq \frac{d(u) \cdot d(v)}{d_s(u) \cdot d_s(v)} \lesssim \frac{n^2}{m^2} \cdot d(u) \cdot d(v).$$

Then it follows from this and the definition of a_z that

$$\|a_w\|_2 = \sqrt{\sum_{(u,v) \in E_s} s(u, v)^2 \cdot \bar{z}(u)^2 \cdot \bar{z}(v)^2} \lesssim \frac{n}{m} \sqrt{\sum_{(u,v) \in E_s} d(u) \cdot d(v) \cdot \bar{z}(u)^2 \cdot \bar{z}(v)^2} \leq \frac{n}{m} \cdot \bar{z}^\top D \bar{z},$$

where the last inequality holds as $(\bar{z}^\top D \bar{z})^2 = \sum_{u,v \in V} d(u) \cdot d(v) \cdot \bar{z}(u)^2 \cdot \bar{z}(v)^2$. Finally, the Courant-Fischer theorem characterizes the second eigenvalue of the normalized Laplacian as

$$\lambda_2(\mathcal{L}_G) = \min_{x: \sum_v d(v) \cdot x(v) = 0} \frac{x^\top L_G x}{x^\top D x}.$$

Since \bar{z} was constructed to satisfy $\sum_v d(v) \cdot \bar{z}(v) = 0$, it follows from Courant-Fischer that $\bar{z}^\top D \bar{z} \geq \lambda_2(\mathcal{L}_G) \cdot \bar{z}^\top L_G \bar{z} \geq \lambda \cdot \bar{z}^\top L_G \bar{z} = \lambda \cdot z^\top L_G z$ as \bar{z} is just a shift of z . Putting together, we conclude that

$$\left| \sum_{e \in E} x(e) \cdot s(e) \cdot \langle z, b_e \rangle^2 \right| \lesssim \|a_z\|_2 \sqrt{\log \frac{|\mathcal{K}|}{m}} \leq \frac{n}{m} \cdot \bar{z}^\top D \bar{z} \cdot \sqrt{\log \frac{|\mathcal{K}|}{m}} \leq \frac{n}{\lambda m} \cdot z^\top L_G z \cdot \sqrt{\log \frac{|\mathcal{K}|}{m}}.$$

Therefore, we have implemented a partial coloring algorithm required by Step 2(a) of the deterministic graphical sketch algorithm with $f(n) = \frac{1}{\lambda} \sqrt{\log \frac{|\mathcal{K}|}{m}}$, and thus the theorem follows from [Lemma 5.1](#). \square

Proof of [Theorem 1.5](#) To construct graphical spectral sketches for arbitrary unweighted undirected graphs, we once again use an expander-decomposition as was done in [\[CGP+23\]](#). By [Theorem 2.2](#), every graph G can be decomposed into G_1, \dots, G_k such that each G_i is an $\Omega(1/\log^2 n)$ -expander and each vertex is contained in at most $O(\log n)$ subgraphs. For each G_i , we apply [Theorem 5.2](#) to compute a deterministic ϵ -graphical spectral sketch \tilde{G}_i of G_i , with $O(n_i \log^2 n \sqrt{\log \frac{|\mathcal{K}|}{n}}/\epsilon)$ edges when $n_i := |V(G_i)|$. Note that $\tilde{G} := \sum_i \tilde{G}_i$ is an ϵ -graphical spectral sketch of G . As each vertex in G is contained in at most $\log n$ subgraphs, it follows that $\sum_{i=1}^k n_i = O(n \log n)$ and thus the total number of edges in \tilde{G} is $O(n \log^3 n \sqrt{\log \frac{|\mathcal{K}|}{n}}/\epsilon)$.

5.2 Effective Resistance Sparsifiers An interesting application of the graphical spectral sketch in [CGP⁺23] was to construct $O(n \text{ polylog } n/\epsilon)$ -sized effective resistance sparsifiers, answering an open question from [ACK⁺16]. They proved the following useful property that relates sketches of Laplacian quadratic forms to sketches of its pseudo-inverse.

LEMMA 5.2. ([CGP⁺23, LEMMA 6.8]) *Let $G = (V, E)$ be an undirected graph. Suppose \tilde{G} is a reweighted subgraph of G satisfying*

1. \tilde{G} is an $O(\sqrt{\epsilon})$ -spectral sparsifier of G ,
2. $(1 - \epsilon) \cdot b_{i,j}^\top L_G^\dagger b_{i,j} \leq b_{i,j}^\top L_G^\dagger L_{\tilde{G}}^\dagger L_G^\dagger b_{i,j} \leq (1 + \epsilon) \cdot b_{i,j}^\top L_G^\dagger b_{i,j}$ for all $i, j \in V$.

Then $(1 - \epsilon) \cdot b_{i,j}^\top L_G^\dagger b_{i,j} \leq b_{i,j}^\top L_{\tilde{G}}^\dagger b_{i,j} \leq (1 + \epsilon) \cdot b_{i,j}^\top L_G^\dagger b_{i,j}$ for all $i, j \in V$.

In other words, they proved that the problem of finding an effective resistance sparsifier can be reduced to finding a sparsifier that is simultaneously a $O(\sqrt{\epsilon})$ -spectral sparsifier (which requires only $O(n/\epsilon)$ edges by [BSS12]) and an $O(\epsilon)$ -graphical spectral sketch with respect to the vectors $\{L^\dagger b_{i,j}\}_{i,j \in V}$. They used the short cycle decomposition technique and expander decomposition to construct effective resistance sparsifiers with $O(n \text{ polylog } n/\epsilon)$ edges, but it is at least $\Omega(n \log^{16} n/\epsilon)$ even assuming optimal short cycle decomposition and optimal expander decomposition.

In this subsection, we show that the deterministic discrepancy framework can be used to construct sparser effective resistance sparsifiers and prove Theorem 1.6. As was done in [CGP⁺23], we first construct sparsifiers on expander graphs and then apply expander decomposition.

Deterministic Effective Resistance Sparsification

Input: a λ -expander $G = (V, E)$, and a target accuracy parameter ϵ .

Output: $\tilde{G} = (V, \tilde{E})$ that is an ϵ -effective resistance sparsifier of G with $|\tilde{E}| \lesssim n\sqrt{\log n}/(\lambda\epsilon)$.

1. Initialize $s_0(e) = 1$ for $e \in E$. Let $t = 1$.
2. **While** $m_t := |\text{supp}(s_{t-1})| > cn\sqrt{\log n}/(\lambda\epsilon)$ for some fixed constant c **do**
 - (a) Apply the matrix partial coloring algorithm with inputs $\{A_e\}_{e \in E}$ where $A_e = L_G^{\dagger/2} b_e b_e^\top L_G^{\dagger/2}$ and vector constraints $\mathcal{K} = \{L_G^\dagger b_{i,j}\}_{i,j \in V}$. Find a partial coloring $x_t : E \rightarrow [-1, 1]$ such that
 - i. the matrix discrepancy is
$$\left\| \sum_e x_t(e) \cdot s_{t-1}(e) \cdot L_G^{\dagger/2} b_e b_e^\top L_G^{\dagger/2} \right\|_{\text{op}} \lesssim \sqrt{\frac{n}{m}},$$
 - ii. the vector discrepancy is
$$\sum_e x_t(e) \cdot s_{t-1}(e) \cdot (b_{i,j}^\top L_G^\dagger b_e)^2 \lesssim \frac{n\sqrt{\log n}}{\lambda m} \cdot b_{i,j}^\top L_G^\dagger b_{i,j} \quad \forall i, j \in V,$$
 - iii. x_t is degree-preserving such that $\sum_{u:u \sim v} x_t(u, v) \cdot s_{t-1}(u, v) = 0$ for all $v \in V$,
 - iv. $|\{i \in \text{supp}(s_{t-1}) \mid x_t(i) = \pm 1\}| = \Omega(m_t)$ and $x_t(i) = 0$ for all $i \notin \text{supp}(s_{t-1})$.
 - (b) If there are more $x_t(i) = 1$ than $x_t(i) = -1$ then update $x_t \leftarrow -x_t$.
 - (c) Update $s_t(e) \leftarrow s_t(e) = s_{t-1}(e) \cdot (1 + x_t(e))$ for all $e \in E$.
 - (d) $t \leftarrow t + 1$
3. **Return:** \tilde{G} with edges weight $s_T(e)$ for $e \in E$ where T is the last iteration.

We prove the existence of a nearly-linear sized effective resistance sparsifiers for $\Omega(1)$ -expander graphs.

THEOREM 5.3. *Given a λ -expander graph G , the deterministic effective resistance sparsification algorithm always returns an ϵ -effective resistance sparsifier with $O(\frac{n\sqrt{\log n}}{\lambda\epsilon})$ edges in deterministic polynomial time.*

Proof. The main task is to show that the partial coloring in Step 2(a) can be implemented in deterministic polynomial time. Suppose we are in the τ -th iteration in the while loop. Let $t \in [T]$ denote the iteration of a particular run of the deterministic walk partial coloring algorithm in Step 2(a). In the proof of [Theorem 1.1](#) for matrix discrepancy, we showed that there exists a subspace U_t of dimension at least $\frac{2}{3}m_t$ such that if y_t is picked from U_t for all iterations t , then the matrix discrepancy requirement in Step 2(a)(i) is satisfied at the end¹⁶. In the proof of [Theorem 5.2](#) for spectral sketches, we showed that under the same framework, there exists a subspace V_t of dimension at least $\frac{1}{2}m_t - n$ such that as long as y_t is picked from V_t for all iterations t , then the vector discrepancy requirement in Step 2(a)(ii) and the degree-preserving constraints in Step 2(a)(iii) are satisfied at the end. Therefore, by choosing $y_t \in U_t \cap V_t$, which still has large enough dimension as long as $m_t > 10n$ say, to incorporate the standard partial coloring constraints in Step 2(a)(iv). Finally, by using the same arguments as in [Theorem 1.2](#) and [Lemma 5.1](#), we see that at the end of the algorithm, our graph \tilde{G} has $O(\frac{n\sqrt{\log n}}{\lambda\epsilon})$ edges and satisfy that it is an $O(\sqrt{\epsilon})$ -spectral sparsifier as well as an $O(\epsilon)$ -spectral sketches with respect to \mathcal{K} , and hence \tilde{G} is an ϵ effective resistance sparsifier of G by [Lemma 5.2](#). \square

Finally, we obtain [Theorem 1.6](#) using [Theorem 5.3](#) and the expander decomposition in [Theorem 2.2](#). The proof is the same as in the proof of [Theorem 1.5](#) and is omitted.

Concluding Remarks

Building on recent works [[RR20](#), [PV23](#)], we developed a unified algorithmic framework for both discrepancy minimization and spectral sparsification, by combining the potential functions from spectral sparsification [[BSS12](#), [AZLO15](#)] and the partial coloring and perturbation updates from discrepancy minimization [[Spe85](#), [Ban10](#), [LM15](#), [LRR17](#)]. We demonstrate this framework by showing simpler and improved constructions for various spectral sparsification problems. The analysis is self-contained and elementary and is considerably simpler than that in [[RR20](#)], and even in the standard setting is more intuitive and arguably simpler than that in [[BSS12](#), [AZLO15](#)].

Together with the results in [[PV23](#)], this framework recovers best known results in many settings in discrepancy minimization and spectral sparsification, but not the most advanced ones such as the matrix Spencer problem [[BJM23](#)] and the Weaver’s discrepancy problem [[MSS15](#)]. It is thus an interesting and important open direction to extend this framework to recover these results. One concrete and intermediate problem in this direction is to recover the result in [[STZ24](#)] for Eulerian sparsifiers, which uses the result in [[BJM23](#)] for matrix Spencer as a black box. Other related problems are to recover some results in approximation algorithms [[AOG15](#), [LZ22b](#), [LZ22a](#)], which built on the results for spectral sparsification [[AZLO15](#), [AZLSW21](#)] and for the Kadison-Singer problem [[MSS15](#), [AOG14](#), [KLS20](#)].

Another open question is to design fast deterministic algorithms for spectral sparsification and discrepancy minimization, which has been open even in the standard settings [[BSS12](#), [Ban10](#)].

Acknowledgements

We thank Valentino Dante Tjowasi for collaborations at the initial stage of this work. Hong Zhou was supported in part by National Key R&D Program of China (Grant No. 2023YFA1010202) and Natural Science Foundation of Fujian Province of China (Grant No. 2023J01414). Lap Chi Lau was supported in part by NSERC Discovery Grant. Robert Wang was supported in part by NSERC Discovery Grant and Canada Graduate Scholarship.

References

- [ACK⁺16] Alexandr Andoni, Jiecao Chen, Robert Krauthgamer, Bo Qin, David P. Woodruff, and Qin Zhang. On sketching quadratic forms. In *ITCS’16—Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science*, pages 311–319, 2016.

¹⁶In the proof of [Theorem 1.1](#), we only showed that $\dim(U_t) \geq \frac{1}{4}m$, but this can be adjusted to have dimension at least $\frac{2}{3}m$ (or any arbitrary constant times m) by making the eigenspace in U^3 larger by a constant factor which would only increase the second-order term by a constant.

- [AKM⁺20] AmirMahdi Ahmadinejad, Jonathan A. Kelner, Jack Murtagh, John Peebles, Aaron Sidford, and Salil P. Vadhan. High-precision estimation of random walks in small space. In *61st IEEE Annual Symposium on Foundations of Computer Science, FOCS 2020*, pages 1295–1306, 2020.
- [AOG14] Nima Anari and Shayan Oveis Gharan. The Kadison-Singer problem for strongly Rayleigh measures and applications to asymmetric TSP. *arXiv preprint arXiv:1412.1143*, 2014.
- [AOG15] Nima Anari and Shayan Oveis Gharan. Effective-resistance-reducing flows, spectrally thin trees, and asymmetric TSP. In *Proceedings of the 56th Annual Symposium on Foundations of Computer Science, FOCS '15*, pages 20–39, Washington, DC, USA, 2015. IEEE.
- [APP⁺23] AmirMahdi Ahmadinejad, John Peebles, Edward Pyne, Aaron Sidford, and Salil P. Vadhan. Singular value approximation and sparsifying random walks on directed graphs. In *64th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2023*, pages 846–854, 2023.
- [AZLO15] Zeyuan Allen-Zhu, Zhenyu Liao, and Lorenzo Orecchia. Spectral sparsification and regret minimization beyond matrix multiplicative updates [extended abstract]. In *STOC'15—Proceedings of the 2015 ACM Symposium on Theory of Computing*, pages 237–245. ACM, New York, 2015.
- [AZLSW21] Zeyuan Allen-Zhu, Yuanzhi Li, Aarti Singh, and Yining Wang. Near-optimal discrete optimization for experimental design: a regret minimization approach. *Math. Program.*, 186(1-2, Ser. A):439–478, 2021.
- [Ban10] Nikhil Bansal. Constructive algorithms for discrepancy minimization. In *Proceedings of the 51th Annual IEEE Symposium on Foundations of Computer Science, FOCS '10*, pages 3–10. IEEE Computer Society, 2010.
- [BBvH23] Afonso S. Bandeira, March T. Boedihardjo, and Ramon van Handel. Matrix concentration inequalities and free probability. *Invent. Math.*, 234(1):419–487, 2023.
- [BJM23] Nikhil Bansal, Haotian Jiang, and Raghu Meka. Resolving matrix Spencer conjecture up to poly-logarithmic rank. In *STOC'23—Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, pages 1814–1819, 2023.
- [BK15] András A. Benczúr and David R. Karger. Randomized approximation schemes for cuts and flows in capacitated graphs. *SIAM J. Comput.*, 44(2):290–319, 2015.
- [BLV22] Nikhil Bansal, Aditi Laddha, and Santosh S. Vempala. A unified approach to discrepancy minimization. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2022*, volume 245 of *LIPICs*, pages 1:1–1:22. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022.
- [BSS12] Joshua Batson, Daniel A. Spielman, and Nikhil Srivastava. Twice-Ramanujan sparsifiers. *SIAM J. Comput.*, 41(6):1704–1721, 2012.
- [CCL⁺15] Dehua Cheng, Yu Cheng, Yan Liu, Richard Peng, and Shang-Hua Teng. Spectral sparsification of random-walk matrix polynomials. *CoRR*, abs/1502.03496, 2015.
- [CGP⁺23] Timothy Chu, Yu Gao, Richard Peng, Sushant Sachdeva, Saurabh Sawlani, and Junxing Wang. Graph sparsification, spectral sketches, and faster resistance computation via short cycle decompositions. *SIAM J. Comput.*, 52(6):FOCS18–85–FOCS18–157, 2023.
- [Cha00] Bernard Chazelle. *The discrepancy method*. Cambridge University Press, Cambridge, 2000. Randomness and complexity.
- [CKP⁺17] Michael B. Cohen, Jonathan Kelner, John Peebles, Richard Peng, Anup B. Rao, Aaron Sidford, and Adrian Vladu. Almost-linear-time algorithms for Markov chains and new spectral primitives for directed graphs. In *STOC'17—Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 410–419. ACM, New York, 2017.
- [JRT24] Arun Jambulapati, Victor Reis, and Kevin Tian. Linear-sized sparsifiers via near-linear time discrepancy theory. In *Proceedings of the 2024 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 5169–5208, 2024.
- [JS18] Arun Jambulapati and Aaron Sidford. Efficient $\tilde{O}(n/\epsilon)$ spectral sketches for the Laplacian and its pseudoinverse. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2487–2503. SIAM, Philadelphia, PA, 2018.
- [Kar99] David R. Karger. Random sampling in cut, flow, and network design problems. *Math. Oper. Res.*, 24(2):383–413, 1999.
- [KLS20] Rasmus Kyng, Kyle Luh, and Zhao Song. Four deviations suffice for rank 1 matrices. *Adv. Math.*, 375:107366, 17, 2020.
- [KVV04] Ravi Kannan, Santosh Vempala, and Adrian Vetta. On clusterings: good, bad and spectral. *J. ACM*, 51(3):497–515, 2004.
- [LM15] Shachar Lovett and Raghu Meka. Constructive discrepancy minimization by walking on the edges. *SIAM J. Comput.*, 44(5):1573–1582, 2015.
- [LRR17] Avi Levy, Harishchandra Ramadas, and Thomas Rothvoss. Deterministic discrepancy minimization via the multiplicative weight update method. In *Integer Programming and Combinatorial Optimization (IPCO)*, pages 380–391, 2017.
- [LZ22a] Lap Chi Lau and Hong Zhou. A local search framework for experimental design. *SIAM J. Comput.*, 51(4):900–951,

2022.

- [LZ22b] Lap Chi Lau and Hong Zhou. A spectral approach to network design. *SIAM J. Comput.*, 51(4):1018–1064, 2022.
- [Mat99] Jiri Matousek. *Geometric discrepancy*, volume 18 of *Algorithms and Combinatorics*. Springer-Verlag, Berlin, 1999. An illustrated guide.
- [Mek14] Raghu Meka. Blog post: Discrepancy and beating the union bound, 2014. URL: <https://windowsontheory.org/2014/02/07/discrepancy-and-beating-the-union-bound/>. Last visited on 2024/06/11.
- [MSS15] Adam W. Marcus, Daniel A. Spielman, and Nikhil Srivastava. Interlacing families II: Mixed characteristic polynomials and the Kadison-Singer problem. *Ann. of Math. (2)*, 182(1):327–350, 2015.
- [PV23] Lucas Pesenti and Adrian Vladu. Discrepancy minimization via regularization. In *Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1734–1758. SIAM, Philadelphia, PA, 2023.
- [Rot17] Thomas Rothvoss. Constructive discrepancy minimization for convex sets. *SIAM J. Comput.*, 46:224–234, 2017.
- [RR20] Victor Reis and Thomas Rothvoss. Linear size sparsifier and the geometry of the operator norm ball. In *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2337–2348. SIAM, Philadelphia, PA, 2020.
- [Spe85] Joel Spencer. Six standard deviations suffice. *Trans. Amer. Math. Soc.*, 289(2):679–706, 1985.
- [SS11] Daniel A. Spielman and Nikhil Srivastava. Graph sparsification by effective resistances. *SIAM J. Comput.*, 40(6):1913–1926, 2011.
- [ST11] Daniel A. Spielman and Shang-Hua Teng. Spectral sparsification of graphs. *SIAM J. Comput.*, 40(4):981–1025, 2011.
- [ST14] Daniel A. Spielman and Shang-Hua Teng. Nearly linear time algorithms for preconditioning and solving symmetric, diagonally dominant linear systems. *SIAM J. Matrix Anal. Appl.*, 35(3):835–885, 2014.
- [STZ24] Sushant Sachdeva, Anvith Thudi, and Yibin Zhao. Better sparsifiers for directed Eulerian graphs. In *51st International Colloquium on Automata, Languages, and Programming (ICALP)*, volume 297 of *LIPICs. Leibniz Int. Proc. Inform.*, pages Art. No. 119, 20, 2024.
- [Wea04] Nik Weaver. The Kadison-Singer problem in discrepancy theory. *Discrete Math.*, 278(1-3):227–239, 2004.
- [Zou12] Anastasios Zouzias. A matrix hyperbolic cosine algorithm and applications. In *Automata, Languages, and Programming - 39th International Colloquium, ICALP 2012, Proceedings, Part I*, volume 7391 of *Lecture Notes in Computer Science*, pages 846–858. Springer, 2012.

A Simple Proof of Lemma 3.1

We first compute the optimizer M . Since Δ_m is a compact set and the regularizer $\phi(M) = -2 \operatorname{tr}(M^{\frac{1}{2}})$ is strongly convex, the optimizer of (3.1) is attained and uniquely defined. Moreover, as the gradient of the $\ell_{1/2}$ -regularizer is $\nabla\phi(M) = M^{-1/2}$, which blows up when M is singular, the optimizer of (3.1) stays in the interior of Δ_m . Therefore, we can apply the KKT condition to (3.1) without the constraint $M \succcurlyeq 0$ (so only the constraint $\operatorname{tr}(M) = \langle M, I \rangle = 1$), and obtain a closed-form characterization of the unique optimizer

$$M = (u_x I_n - \eta A(x))^{-2},$$

where u_x is the unique value such that M is a density matrix (note that u_x is a function of x). We can then compute the closed-form characterization of the potential function. Let $A(x) = \sum_{i=1}^n \lambda_i v_i v_i^\top$ be its eigen-decomposition. Then $M = \sum_{i=1}^n (u_x - \eta \lambda_i)^{-2} v_i v_i^\top$, and so

$$\Phi(x) = \frac{1}{\eta} \operatorname{tr}(M^{\frac{1}{2}}) + \langle A(x), M \rangle + \frac{1}{\eta} \operatorname{tr}(M^{\frac{1}{2}}) = \frac{1}{\eta} \operatorname{tr}(M^{\frac{1}{2}}) + \sum_{i=1}^n \left(\lambda_i (u_x - \eta \lambda_i)^{-2} + \frac{(u_x - \eta \lambda_i)^{-1}}{\eta} \right).$$

Rearranging the terms and using $\sum_{i=1}^n (u_x - \eta \lambda_i)^{-2} = 1$ as M is a density matrix, we obtain that

$$(A.1) \quad \Phi(x) = \frac{1}{\eta} \operatorname{tr}(M^{1/2}) + \frac{u_x}{\eta} \cdot \sum_{i=1}^d (u_x - \eta \lambda_i)^{-2} = \frac{1}{\eta} \operatorname{tr}((u_x I_n - \eta A(x))^{-1}) + \frac{u_x}{\eta}.$$

Note that this is basically the same as the trace inverse potential function in [BSS12].

To bound the potential increase, the following claim was implicitly proved in [AZLO15] using mirror descent and Bregman divergence. Here we provide a simpler proof using elementary convexity arguments.

CLAIM A.1. *If $\|(u_x I_n - \eta A(x))^{-1} \cdot \eta A(y)\|_{\text{op}} < 1$, then it holds that*

$$\Phi(x+y) - \Phi(x) \leq \frac{1}{\eta} \left(\operatorname{tr}((u_x I_n - \eta A(x+y))^{-1}) - \operatorname{tr}((u_x I_n - \eta A(x))^{-1}) \right).$$

Proof. By (A.1), the potential increase is

$$(A.2) \quad \Phi(x+y) - \Phi(x) = \frac{1}{\eta} \left(\underbrace{\text{tr}((u_{x+y}I_n - \eta A(x+y))^{-1})}_{(*)} - \text{tr}((u_x I_n - \eta A(x))^{-1}) \right) + \frac{u_{x+y} - u_x}{\eta}.$$

Given η and x and y , consider the univariate function

$$f(u) = \text{tr}((uI_n - \eta A(x+y))^{-1}).$$

Note that the $(*)$ term is equal to $f(u_{x+y})$. We will use the convexity of f to bound the $(*)$ term.

First, we show that $f(u)$ is convex within the interval between u_x and u_{x+y} . Let $\beta_1 \leq \dots \leq \beta_n$ be the eigenvalues of $\eta A(x+y)$. When $u > \beta_n = \lambda_{\max}(\eta A(x+y))$, we can check that $f(u) = \sum_{i=1}^n (u - \beta_i)^{-1}$ is convex in u (say by the second derivative test). So, we just need to show that both u_{x+y} and u_x are greater than $\lambda_{\max}(\eta A(x+y))$ to establish the convexity within the interval between them. For the former, since $(u_{x+y}I_n - \eta A(x+y))^{-2}$ is positive definite, it follows immediately that $u_{x+y} > \lambda_{\max}(\eta A(x+y))$. For the latter, the assumption $\|(u_x I_n - \eta A(x))^{-1} \cdot \eta A(y)\|_{\text{op}} < 1$ implies that $\eta A(y) \preceq u_x I_n - \eta A(x)$, and thus $u_x > \lambda_{\max}(\eta A(x+y))$. Therefore, $f(u)$ is convex within the interval between u_x and u_{x+y} . This implies that

$$f(u_{x+y}) + f'(u_{x+y})(u_x - u_{x+y}) \leq f(u_x),$$

or equivalently,

$$(A.3) \quad \text{tr}((u_{x+y}I_n - \eta A(x+y))^{-1}) \leq \text{tr}((u_x I_n - \eta A(x+y))^{-1}) - f'(u_{x+y}) \cdot (u_x - u_{x+y}).$$

It remains to compute $f'(u_{x+y})$. The derivative of $f(u)$ when $u > \lambda_{\max}(\eta A(x+y))$ is

$$f'(u) = \text{tr}(\partial_u (uI_n - \eta A(x+y))^{-1}) = -\text{tr}((uI_n - \eta A(x+y))^{-2}).$$

Since $(u_{x+y}I_n - \eta A(x+y))^{-2}$ is a density matrix, when $f'(u)$ is evaluated at u_{x+y} , we have

$$f'(u_{x+y}) = -\text{tr}((u_{x+y}I_n - \eta A(x+y))^{-2}) = -1.$$

Plugging it into (A.3), it follows that

$$\text{tr}((u_{x+y}I_n - \eta A(x+y))^{-1}) \leq \text{tr}((u_x I_n - \eta A(x+y))^{-1}) + (u_x - u_{x+y}).$$

Combining with (A.2), we conclude that

$$\Phi(x+y) - \Phi(x) \leq \frac{1}{\eta} (\text{tr}((u_x I_n - \eta A(x+y))^{-1}) - \text{tr}((u_x I_n - \eta A(x))^{-1})).$$

□

We now apply Lemma 2.3 (from [RR20]) to finish the proof. Set $A := u_x I_n - \eta A(x) = M^{-\frac{1}{2}} \succ 0$ and $B := A(y)$. The assumption $\|M^{\frac{1}{2}} \cdot \eta A(y)\|_{\text{op}} \leq \frac{1}{2}$ in Lemma 3.1 translates to the assumption $\|\eta A^{-1} B\|_{\text{op}} \leq \frac{1}{2}$ in Lemma 2.3. So, we can apply Lemma 2.3 to conclude that there is a value $c \in [-2, 2]$ so that

$$\text{tr}((u_x I_n - \eta A(x+y))^{-1}) - \text{tr}((u_x I_n - \eta A(x))^{-1}) = \eta \text{tr}(MA(y)) + c\eta^2 \text{tr}(M^{\frac{1}{2}} A(y) M^{\frac{1}{2}} A(y) M^{\frac{1}{2}}).$$

Combining with Claim A.1, we proved Lemma 3.1.