

Pipage Rounding, Pessimistic Estimators and Matrix Concentration

Nicholas J. A. Harvey
Department of Computer Science
UBC
Vancouver, Canada
nickhar@cs.ubc.ca

Neil Olver
Department of Mathematics
MIT
Cambridge, USA
olver@math.mit.edu

Abstract

Pipage rounding is a dependent random sampling technique that has several interesting properties and diverse applications. One property that has been particularly useful is negative correlation of the resulting vector. Unfortunately negative correlation has its limitations, and there are some further desirable properties that do not seem to follow from existing techniques. In particular, recent concentration results for sums of independent random matrices are not known to extend to a negatively dependent setting.

We introduce a simple but useful technique called *concavity of pessimistic estimators*. This technique allows us to show concentration of submodular functions and concentration of matrix sums under pipage rounding. The former result answers a question of Chekuri et al. (2009). To prove the latter result, we derive a new variant of Lieb's celebrated concavity theorem in matrix analysis.

We provide numerous applications of these results. One is to *spectrally-thin trees*, a spectral analog of the *thin trees* that played a crucial role in the recent breakthrough on the asymmetric traveling salesman problem. We show a polynomial time algorithm that, given a graph where every edge has effective conductance at least κ , returns an $O(\kappa^{-1} \cdot \log n / \log \log n)$ -spectrally-thin tree. There are further applications to rounding of semidefinite programs, to the column subset selection problem, and to a geometric question of extracting a nearly-orthonormal basis from an isotropic distribution.

1 Introduction

Rounding is a crucial step in the design of many approximation algorithms. Given a fractional vector satisfying some constraints, a rounding method produces an integer vector that satisfies those constraints, either exactly or approximately. *Randomized rounding* [41] [57, Chapter 5], in which the coordinates of the fractional vector are rounded randomly and independently, produces good integer vectors for many applications. *Dependent rounding* methods, in which the resulting integer vector does not have independent coordinates, are important in many scenarios where naive randomized rounding does poorly. Various techniques exist for designing dependent rounding methods (see, e.g., the surveys [48, 4]).

It is common for a rounding scenario to involve two types of constraints: hard constraints, which must be satisfied exactly by the integer solution, and soft constraints, which must be approximately satisfied by the integer solution. Low-congestion multi-path routing [49], max cut with given sizes of parts [1], thin spanning trees [3], and submodular maximization under a matroid constraint [15, 20] are examples of problems whose solutions involve such a rounding scenario. The hard constraint is often membership in an integer polytope that is defined using combinatorial objects (e.g., matchings or matroids). The soft constraints are usually simple linear inequalities.

With randomized rounding, the independent choices lead to concentration of measure phenomena that are useful for handling soft constraints. For example, Chernoff bounds are commonly used to show that linear inequalities are approximately satisfied [41]. The past decade has seen various uses of *matrix* concentration bounds (e.g., [2, 44, 52]) to show that linear matrix inequalities are approximately satisfied by random sampling or rounding. Such uses have occurred in many diverse areas: graph sparsification [46], compressed sensing [54], statistics [21], machine learning [42] and numerical linear algebra [33].

With dependent rounding, concentration phenomena can also occur. *Pipage rounding*, *swap rounding* and *maximum entropy sampling* are dependent rounding techniques that have seen many important uses over the past decade [49, 1, 26, 15, 3, 20]. An important feature in some scenarios is that any Chernoff bound that is valid under independent randomized rounding remains valid under these dependent rounding techniques. This fact is proven by showing that the rounded solution has a *negatively correlated distribution*, then appealing to the fact that Chernoff bounds remain valid under such distributions [39]. Unfortunately, commutativity plays a key role in proving that fact, and these arguments do not seem to extend to matrix concentration bounds, e.g., [2, 37, 44, 52]. Consequently, these matrix inequalities have so far not been combined with dependent rounding.

We prove the first result showing that matrix concentration bounds are usable in a dependent rounding scenario. Our technique is not based on negative correlation, but rather the fortuitous interaction between pipage rounding and various pessimistic estimators. In particular, we show that Tropp’s matrix Chernoff bound [52] has a pessimistic estimator that decreases monotonically under pipage rounding. As a consequence, we can extend the reach of pipage rounding from soft constraints that are linear inequalities to soft constraints that are linear matrix inequalities. Our proof uses non-trivial techniques from matrix analysis and complex analysis; in particular, we prove a new variant of Lieb’s concavity theorem.

1.1 Motivation and Results

One key area where our techniques yield new results is for *thin spanning trees*. These are intriguing objects in graph theory that relate to foundational topics, such as nowhere-zero flows [28], and the asymmetric traveling salesman problem [3]. Given a graph G on n nodes, a spanning tree T of G is α -*thin* if, for every cut, the number of edges of T crossing the cut is at most α times the number of

edges of G crossing the cut. It has been conjectured that any graph with connectivity k has an $f(k)$ -thin spanning tree where $f(k) = O(1/k)$. This would imply a constant factor approximation algorithm for the asymmetric traveling salesman problem [38]. Asadpour et al. [3] give a randomized algorithm to find a spanning tree that is $O(\frac{\log n}{k \log \log n})$ -thin. Later Chekuri et al. [19, 20] gave a simpler algorithm using randomized pipage rounding or swap rounding.

A *spectrally-thin* spanning tree is a stronger notion that is naturally motivated by work on spectral sparsification [46, 5]. A spanning tree T is α -spectrally-thin if $L_T \preceq \alpha L_G$, where L_G refers to the Laplacian of G , and \preceq to the Löwner ordering of Hermitian matrices. In Section 4.3, we show a result on spectrally thin trees that strongly mirrors the result of Asadpour et al.

Theorem 1.1. *There is a deterministic, polynomial-time algorithm that given any graph on n nodes where every edge has effective conductance at least κ , constructs a $O(\frac{\log n}{\kappa \log \log n})$ -spectrally-thin spanning subtree.*

This spectral notion of thinness seems to be an important one, as the recent breakthrough of Marcus et al. [34] implies that $O(1/\kappa)$ -spectrally-thin trees exist. Details of this connection are given in Appendix F. It is unknown if similar techniques can show that $O(1/k)$ -thin trees exist. The best known algorithmic construction of spectrally-thin trees is still Theorem 1.1.

This result is a special case of a result in a more abstract geometric setting. Suppose $V = \{v_1, \dots, v_m\}$ are unit vectors in ℓ_2^n for which $\sum_{i=1}^m v_i v_i^\top$ is a multiple of the identity. Does there exist a subset $V_B = \{v_i : i \in B\}$ that is a *basis* of \mathbb{R}^n and for which the maximum eigenvalue of $\sum_{i \in B} v_i v_i^\top$ is small? The maximum eigenvalue is 1 if and only if V_B is orthonormal, but an arbitrary V need not contain an orthonormal basis. Again, the breakthrough of Marcus et al. [34] yields a non-constructive proof of a basis with maximum eigenvalue $O(1)$; see Appendix F. In Section 4.2, we show how to find in polynomial time a basis $V_B \subseteq V$ for which the maximum eigenvalue of $\sum_{i \in B} v_i v_i^\top$ is $O(\log n / \log \log n)$. Previous constructive techniques [2, 37, 44, 52] only provide a bound of $O(\log n)$.

Our geometric result also relates to the *column subset selection* problem in numerical linear algebra [13, 51, 12, 22] which seeks to “approximate” a matrix A by a small subset of its columns, under various notions of approximation. Define the *stable rank* of A to be the Frobenius norm divided by the spectral norm, all squared; this roughly captures the rank of A , ignoring negligibly small singular values. In numerical linear algebra [13, 12, 22], the number of columns chosen is typically much larger than the stable rank. The operator theory community considers similar questions [10, 11, 47, 51], although the number of columns selected is typically much smaller than the stable rank. In Section 4.4, we show that one can efficiently select a *linearly independent* set of columns of size *equal* to the stable rank, while carefully controlling the maximum singular value.

1.2 Techniques

Our results are based on the pipage rounding technique [1, 49, 26, 15], which has had several interesting uses in the recent literature. Deterministic and randomized forms of pipage rounding exist; our result applies to both of those, as well as to swap rounding. Typical uses of pipage rounding involve some of the following ideas.

- There are processes that iteratively move a point in a matroid base polytope towards an extreme point, while modifying only two coordinates at a time. The exchange properties of matroid bases ensure that this is possible.
- One can define a “potential function” on the matroid base polytope (e.g., the ad hoc functions defined in [1], or the multilinear extension of a submodular function [15]) such that the function is concave or convex in directions that increase one coordinate and decrease another.

- The randomized form of pipage rounding [49, 26, 20] outputs a matroid base whose elements are negatively correlated (more precisely, *negative cylinder dependent*). This ensures that linear functions of that base satisfy the same Chernoff-type concentration bounds that are satisfied under independent rounding.

Our aim is to show that, for various concentration bounds, the final extreme point satisfies the same bounds that would be achieved by independent randomized rounding. For Chernoff bounds this follows from negative correlation, but for other bounds such a result was not previously known.

- Let f be a monotone submodular function defined on the ground set of the matroid. When using randomized pipage rounding, does the value of f at the final extreme point satisfy the same lower tail bound as when using independent rounding? Chekuri et al. [19] conjectured this to be true, and they proved such a result when using swap rounding.
- Let f be a linear function mapping points in the matroid base polytope to symmetric matrices. When using pipage rounding, can the value of f at the final extreme point be guaranteed to satisfy the same eigenvalue bounds as when using independent rounding?

It does not seem easy to answer these questions using negative correlation properties.

We present a new approach that leads to a positive answer to both of these questions. In both cases, we can define a pessimistic estimator [40] that bounds the probability that randomized rounding fails to achieve the desired concentration. We show that these pessimistic estimators are concave when one element’s sampling probability is increased and another’s is decreased by the same amount. Due to that concavity property, the base output by randomized pipage rounding satisfies the same concentration bounds that would be satisfied under independent randomized rounding. For the second question (matrix concentration), the pessimistic estimator can be efficiently evaluated, so deterministic pipage rounding can also be used.

The concavity property of our pessimistic estimator for matrix concentration is a non-trivial fact. We establish that fact by proving a new variant of Lieb’s concavity theorem [32], which is a “masterpiece of matrix analysis” [8] with deep applications in mathematical physics and quantum information theory [16, 23, 35]. Although there is much interest in the mathematical physics community on extensions and variants of Lieb’s theorem, our particular variant does not seem to appear in the literature.

2 Preliminaries

Let $[m] = \{1, \dots, m\}$. For a set $S \subseteq [m]$, the vector $\chi(S) \in \mathbb{R}^m$ is the characteristic vector of S . For a vector $x \in \mathbb{R}^m$ and a set $S \subseteq [m]$, the notation $x(S)$ denotes $\sum_{i \in S} x_i$. The vector e_i denotes the i^{th} standard basis vector of the finite dimensional vector space that is apparent from context. The vector $\vec{1}$ denotes a vector whose components are all ones and whose dimension is apparent from context. We will use \mathbb{R}_+ and \mathbb{Z}_+ to denote the nonnegative and positive reals respectively.

Let \mathbb{S}^n denote the space of symmetric, real matrices of size $n \times n$. Let $\mathbb{S}_+^n, \mathbb{S}_{++}^n \subset \mathbb{S}^n$ respectively denote the cones of positive semidefinite and positive definite matrices. Let $\mathbb{D}^n \subseteq \mathbb{S}^n$ denote the space of $n \times n$ diagonal matrices. Let \preceq denote the Löwner partial order on symmetric matrices, i.e., $A \preceq B$ iff $B - A \in \mathbb{S}_+^n$. Similarly, $A \prec B$ iff $B - A \in \mathbb{S}_{++}^n$. For $A \in \mathbb{S}^n$, let $\lambda_{\max}(A)$ and $\lambda_{\min}(A)$ respectively denote the largest and smallest eigenvalues of A . For $B \in \mathbb{S}^n$, let B^+ denote its Moore-Penrose pseudoinverse. For $B \in \mathbb{S}_+^n$, let $B^{+/2} \in \mathbb{S}_+^n$ denote the positive semidefinite square root of B^+ . The image of B is $\text{im } B$ and the orthogonal projection onto $\text{im } B$ is $I_{\text{im } B}$.

The notation $\|\cdot\|$ denotes the ℓ_2 norm for vectors and the ℓ_2 operator norm for matrices.

If \mathcal{D} is a distribution, $X \sim \mathcal{D}$ means that the random variable X has distribution \mathcal{D} .

3 Concavity of Pessimistic Estimators

In this section we state the known results on pipage rounding and our concavity of pessimistic estimators technique. We then apply this technique in three scenarios, of increasing difficulty: (1) Chernoff bounds, (2) submodular functions, and (3) matrix concentration. The latter two results are new, and in particular are not known to follow using negative correlation. This pessimistic estimator for matrix concentration underlies all applications in Section 4.

3.1 Pipage Rounding

Pipage rounding is a dependent rounding process originating in works of Ageev, Srinivasan and Sviridenko [1, 49]. Calinescu et al. [15] generalized it to a matroid setting. We now state the main results of randomized and deterministic pipage rounding; a proof sketch is given in Appendix A.

Let \mathbf{M} be a matroid on $[m]$ and let $P \subset \mathbb{R}^m$ be its base polytope. For all algorithmic applications in this paper, \mathbf{M} can be presented to the algorithm via an independence oracle. A function $g : P \rightarrow \mathbb{R}$ is said to be *concave under swaps* if

$$\forall p \in P, \forall a, b \in [m], \quad z \mapsto g(p + z(e_a - e_b)) \text{ is concave.} \quad (1)$$

Theorem 3.1 (Randomized Pipage Rounding). *There is a randomized, polynomial-time algorithm that, given $x \in P$, outputs an extreme point \hat{x} of P with $\mathbb{E}[\hat{x}] = x$ and such that, for any g concave under swaps, $\mathbb{E}[g(\hat{x})] \leq g(x)$.*

Theorem 3.2 (Deterministic Pipage Rounding). *There is a deterministic, polynomial-time algorithm that, given $x \in P$ and a value oracle for a function g that is concave under swaps, outputs an extreme point \hat{x} of P with $g(\hat{x}) \leq g(x)$.*

The swap rounding procedure of Chekuri et al. [19, 20] also proves Theorem 3.1 and Theorem 3.2.

For $x \in [0, 1]^m$, let $\mathcal{D}(x)$ be the product distribution on $\{0, 1\}^m$ with marginals given by x , i.e., $\mathbb{P}_{X \sim \mathcal{D}(x)}[X_i = 1] = x_i$. Let $\mathcal{E} \subseteq \{0, 1\}^m$. A *pessimistic estimator* [40, 50] for \mathcal{E} is a function $g : [0, 1]^m \rightarrow \mathbb{R}$ that satisfies

$$\begin{aligned} \mathbb{P}_{X \sim \mathcal{D}(x)}[X \in \mathcal{E}] &\leq g(x) & \forall x \in [0, 1]^m \\ \min \{ g(x - x_i e_i), g(x + (1 - x_i) e_i) \} &\leq g(x) & \forall x \in [0, 1]^m, i \in [m]. \end{aligned} \quad (2)$$

For uses of pessimistic estimators in derandomization, the function g is also required to be efficiently computable. That is not required with their use in randomized pipage rounding as g is not even provided as input to the algorithm.

Claim 3.3 (Concavity of Pessimistic Estimators). *Let $\mathcal{E} \subseteq \{0, 1\}^m$ and let g be a function that satisfies (2) and is concave under swaps.*

Suppose randomized pipage rounding is started at an initial point $x_0 \in P$, and let \hat{x} be the (random) extreme point of P that is output. If $g(x_0) \leq \epsilon$ then $\mathbb{P}[\hat{x} \in \mathcal{E}] \leq \epsilon$.

Suppose deterministic pipage rounding is given oracle access to g and an initial point $x_0 \in P$ with $g(x_0) < 1$. Then the extreme point \hat{x} of P that is output satisfies $\hat{x} \notin \mathcal{E}$.

We omit the proof of Claim 3.3 as it is an easy consequence of Theorem 3.1 and Theorem 3.2.

3.2 Chernoff bound

Let us start with a simple result to illustrate the technique. First we state the Chernoff bound in convenient notation. We discuss only the right tail; an analogous result holds for the left tail. Fix any vector $w \in [0, 1]^m$. For $t \in \mathbb{R}$ and $\theta > 0$, define $g_{t,\theta} : [0, 1]^m \rightarrow \mathbb{R}$ by

$$g_{t,\theta}(x) := e^{-\theta t} \cdot \mathbb{E}_{X \sim \mathcal{D}(x)} \left[e^{\theta w^\top X} \right].$$

Let $\mu = w^\top x$ and $\delta \geq 0$. Then

$$\mathbb{P}_{X \sim \mathcal{D}(x)} \left[w^\top X \geq t \right] \leq \inf_{\theta > 0} g_{t,\theta}(x) \quad \text{and} \quad g_{(1+\delta)\mu, \ln(1+\delta)}(x) \leq \left(\frac{e^\delta}{(1+\delta)^{1+\delta}} \right)^\mu. \quad (3)$$

The following claim is proven in Appendix B.

Claim 3.4. $g_{t,\theta}$ is concave under swaps.

Consequently, Claim 3.3 implies the following result.

Corollary 3.5. If randomized pipage rounding starts at $x_0 \in P$ and outputs the extreme point \hat{x} of P then, $\forall w \in [0, 1]^m$, $\delta \geq 0$,

$$\mathbb{P} \left[w^\top \hat{x} \geq (1+\delta)\mu \right] \leq \left(\frac{e^\delta}{(1+\delta)^{1+\delta}} \right)^\mu \quad (4)$$

where $\mu = w^\top x_0$. Furthermore, if this right-hand side is strictly less than 1, then deterministic pipage rounding outputs an extreme point \hat{x} of P with $w^\top \hat{x} < (1+\delta)\mu$.

The key point is that the right-hand sides of (3) and (4) are the same. Chekuri et al. [19] proved this fact using negative correlation of \hat{x} , generalizing a result of Srinivasan [49].

3.3 Submodular functions

Chekuri et al. [19, Theorem 1.3] prove an analog of the Chernoff bound for concentration of submodular functions under independent rounding. They show that the same bound remains true under swap rounding [19, Theorem 1.4] and ask whether it remains true under pipage rounding.

Formally, let $f : \{0, 1\}^m \rightarrow \mathbb{R}$ be a non-negative, monotone, submodular function with marginals in $[0, 1]$. The *multilinear extension* of f is $F : [0, 1]^m \rightarrow \mathbb{R}$ with $F(x) := \mathbb{E}_{X \sim \mathcal{D}(x)} [f(X)]$. For $t \in \mathbb{R}$ and $\theta < 0$, define $g_{t,\theta} : [0, 1]^m \rightarrow \mathbb{R}$ by

$$g_{t,\theta}(x) := e^{-\theta t} \cdot \mathbb{E}_{X \sim \mathcal{D}(x)} \left[e^{\theta f(X)} \right].$$

The left tail bound of Chekuri et al. is: with $\mu = F(x)$, $\delta \in [0, 1)$,

$$\mathbb{P}_{X \sim \mathcal{D}(x)} [f(X) \leq t] \leq \inf_{\theta < 0} g_{t,\theta}(x) \quad \text{and} \quad g_{(1-\delta)\mu, \ln(1-\delta)}(x) \leq \exp(-\delta^2 \mu / 2).$$

The following claim is proven in Appendix B.

Claim 3.6. $g_{t,\theta}$ is concave under swaps.

Claim 3.3 implies the following result, answering an open question of Chekuri et al. [19, p. 3].

Corollary 3.7. If randomized pipage rounding starts at $x_0 \in P$ and outputs the extreme point \hat{x} of P then, letting $\mu = F(x_0)$, we have $\mathbb{P} [f(\hat{x}) \leq (1-\delta)\mu] \leq \exp(-\delta^2 \mu / 2)$.

Chekuri et al. [20, p. 583] state that this fact does not follow from negative correlation of \hat{x} .

3.4 Matrix Concentration

Tropp [52], improving on Ahlswede-Winter [2] and Oliveira [37], proves a beautiful analog of the Chernoff bound for sums of independent random matrices. We state a simplified form here.

Theorem 3.8. *Let $M_1, \dots, M_m \in \mathbb{S}_+^n$ satisfy $M_i \preceq R \cdot I$. For $t \in \mathbb{R}$ and $\theta > 0$, define $g_{t,\theta} : [0, 1]^m \rightarrow \mathbb{R}$ by*

$$g_{t,\theta}(x) := e^{-\theta t} \cdot \text{tr} \exp \left(\sum_{i=1}^m \log \mathbb{E}_{X \sim \mathcal{D}(x)} \left[e^{\theta X_i M_i} \right] \right).$$

Then, for $\mu \geq \left\| \mathbb{E}_{X \sim \mathcal{D}(x)} \left[\sum_i X_i M_i \right] \right\|$ and $\delta \geq 0$,

$$\mathbb{P}_{X \sim \mathcal{D}(x)} \left[\left\| \sum_i X_i M_i \right\| \geq t \right] \leq \inf_{\theta > 0} g_{t,\theta}(x) \quad \text{and} \quad g_{(1+\delta)\mu, \ln(1+\delta)}(x) \leq n \cdot \left(\frac{e^\delta}{(1+\delta)^{1+\delta}} \right)^{\mu/R}.$$

The following is our main lemma on pessimistic estimators. The proof is in Appendix B.

Lemma 3.9. *$g_{t,\theta}$ is concave under swaps.*

Consequently, Claim 3.3 implies the following result.

Corollary 3.10. *Let P be a matroid base polytope and let $x_0 \in P$. Let $M_1, \dots, M_m \in \mathbb{S}_+^n$ satisfy $M_i \preceq R \cdot I$. Let $\mu \geq \left\| \mathbb{E}_{X \sim \mathcal{D}(x_0)} \left[\sum_i X_i M_i \right] \right\|$. If randomized pipage rounding starts at x_0 and outputs the extreme point $\hat{x} = \chi(S)$ of P then we have*

$$\mathbb{P} \left[\left\| \sum_{i \in S} M_i \right\| \geq (1+\delta)\mu \right] \leq n \cdot \left(\frac{e^\delta}{(1+\delta)^{1+\delta}} \right)^{\mu/R}. \quad (5)$$

Furthermore, if this right-hand side is strictly less than 1, then deterministic pipage rounding outputs an extreme point $\hat{x} = \chi(S)$ of P with $\left\| \sum_{i \in S} M_i \right\| < (1+\delta)\mu$.

The inequalities in Theorem 3.8 involve non-trivial matrix analysis, such as operator concavity of log and Lieb's celebrated concavity theorem [32]. It seems that even those results do not suffice to prove Lemma 3.9. To prove it, we derive a new variant of Lieb's theorem (Theorem 3.12). Lieb [32] proved several related concavity theorems; for us, the most relevant form is:

Theorem 3.11 (Lieb [32]). *Let $L, K \in \mathbb{S}^n$ and $C \in \mathbb{S}_{++}^n$. Then $z \mapsto \text{tr} \exp(L + \log(C + zK))$ is concave in a neighborhood of 0.*

The main technical result of this paper is:

Theorem 3.12. *Let $L \in \mathbb{S}^n$, $C_1, C_2 \in \mathbb{S}_{++}^n$ and $K_1, K_2 \in \mathbb{S}_+^n$. Then the univariate function*

$$z \mapsto \text{tr} \exp \left(L + \log(C_1 + zK_1) + \log(C_2 - zK_2) \right) \quad (6)$$

is concave in a neighborhood of 0.

There are several known approaches to proving Lieb's theorem. The simplest is Tropp's approach [53]; however, his proof is based on joint convexity of quantum entropy, which is itself usually proven using Lieb's theorem. We were unable to prove Theorem 3.12 using Tropp's approach. Lieb's original proof [32], which proves concavity by directly analyzing the second derivative, involves numerous delicate steps of matrix analysis. We were able to adapt this approach to prove a weaker form of Theorem 3.12 that requires some additional commutativity assumptions; details are in

Appendix E. This weaker result suffices to prove Lemma 3.9. Epstein [24] gives an elegant approach to proving Lieb’s theorem using complex analysis, and in particular powerful results concerning *Herglotz functions*. Our proof of Theorem 3.12, which appears in Appendix D, is an adaptation of Epstein’s approach.

Remark. Another well-known matrix concentration inequality is the Ahlswede-Winter [2] inequality, for which pessimistic estimators were studied by Wigderson and Xiao [56]. It is natural to wonder whether we could have used their pessimistic estimators instead. Unfortunately they do not seem applicable for our scenario. The issue is that the Ahlswede-Winter inequality is most effective for analyzing sums of i.i.d. random matrices, due to some inequalities that arise in their analysis. In our scenario, due to the way that pipage rounding works, we require non-i.i.d. product distributions, so it is much more convenient to base our approach on Theorem 3.8.

4 Applications

4.1 Rounding of semidefinite programs

Let \mathbf{M} be a matroid and let $P \subset \mathbb{R}^n$ be its base polytope. Consider the spectrahedron

$$Q := P \cap \left\{ x \in \mathbb{R}^m : \sum_{i=1}^m x_i A_i \preceq B \right\}, \quad (7)$$

where each $A_1, \dots, A_m, B \in \mathbb{S}_+^n$. We think of P as specifying “hard” constraints and the semidefinite constraint as being “soft”.

Theorem 4.1. *Suppose that $A_i \preceq B$ for all i . If randomized pipage rounding starts at $x_0 \in Q$ and outputs the extreme point $\chi(S)$ of P , then $\mathbb{P} \left[\sum_{i \in S} A_i \preceq \alpha B \right] \geq 1 - 1/n$, for some $\alpha = O(\log n / \log \log n)$. Furthermore, if deterministic pipage rounding starts at $x_0 \in Q$, then it outputs an extreme point $\chi(S)$ of P with $\sum_{i \in S} A_i \preceq \alpha B$.*

This theorem is optimal with respect to α , as discussed below. The hypothesis that $A_i \preceq B$ is a “width” condition that commonly arises in optimization and rounding.

Proof. Recall the notation defined in Section 2. Let $M_i = B^{+/2} A_i B^{+/2}$. By standard arguments,

$$\begin{aligned} \sum_{i=1}^m x_i A_i \preceq B &\iff \sum_{i=1}^m x_i M_i \preceq I_{\text{im } B} \\ \text{and } \sum_{i \in S} A_i \preceq \alpha B &\iff \sum_{i \in S} M_i \preceq \alpha I_{\text{im } B}. \end{aligned}$$

We assume that $A_i \preceq B$, so $\lambda_{\max}(M_i) \leq 1$. Apply Corollary 3.10 with $\delta = 4 \log n / \log \log n$, $\mu = 1$ and $R = 1$. A standard calculation shows that the right-hand side of (5) is less than $1/n$. \square

Chekuri, Vondrák and Zenklusen [19, 20] considered the problem of rounding a point in a matroid polytope to an extreme point, subject to additional packing constraints. Their result generalizes the low-congestion multi-path routing problem studied earlier by Srinivasan et al. [49, 26], but it is itself a special case of Theorem 4.1 where the matrices A_i and B are diagonal. The factor $\alpha = O(\log n / \log \log n)$ is optimal in Theorem 4.1 because it is optimal for rounding this low-congestion multi-path routing problem, and even for the congestion minimization problem [31].

4.2 Rounding an isotropic distribution to a nearly orthonormal basis

Let $w_1, \dots, w_m \in \mathbb{R}^n$ satisfy $\|w_i\| = 1$ for all i . Let p_1, \dots, p_m be a probability distribution on these vectors such that the covariance matrix is $\sum_i p_i w_i w_i^\top = I/n$. A random vector drawn from that distribution is said to be in *isotropic position*.

Theorem 4.2. *There is a polynomial time algorithm (either randomized or deterministic) to compute a subset $S \subseteq [m]$ such that $\{w_i : i \in S\}$ forms a basis of \mathbb{R}^n , and for which $\|\sum_{i \in S} w_i w_i^\top\| \leq \alpha$, where $\alpha = O(\log n / \log \log n)$.*

As is discussed in Appendix F, the recent breakthrough on the Kadison-Singer problem [34] implies the following existential result:

Theorem 4.3. *There exists $S \subseteq [m]$ such that $\{w_i : i \in S\}$ forms a basis of \mathbb{R}^n , and for which $\|\sum_{i \in S} w_i w_i^\top\| = O(1)$.*

We now prove Theorem 4.2 using Theorem 4.1. Let \mathbf{M} be the linear matroid corresponding to the vectors $\{w_1, \dots, w_m\}$. Let P be the base polytope of that linear matroid. Let $r : 2^{[m]} \rightarrow \mathbb{Z}_+$ be the rank function of that matroid, i.e., $r(S) = \dim(\text{span}\{w_i : i \in S\})$. Then

$$P := \left\{ x \in \mathbb{R}_+^m : x(J) \leq r(J) \ \forall J \subseteq [m] \text{ and } x([m]) = r([m]) \right\}.$$

Define $A_i = w_i w_i^\top$, $B = I$ and

$$Q = P \cap \left\{ x \in \mathbb{R}^m : \sum_i x_i A_i \preceq B \right\}.$$

Let $x = n \cdot p$. Then the following claim and the hypothesis that $\sum_i p_i w_i w_i^\top = I/n$ show that $x \in Q$.

Claim 4.4. $x \in P$.

Since $\|w_i\| = 1$, we have $A_i = w_i w_i^\top \preceq I = B$. Theorem 4.1 gives an algorithm to construct an extreme point $\chi(S)$ of P for which $\sum_{i \in S} A_i \preceq \alpha \cdot B$, with $\alpha = O(\log n / \log \log n)$. Since P is the base polytope of \mathbf{M} , $\{w_i : i \in S\}$ forms a basis of \mathbb{R}^n . Finally, $\sum_{i \in S} w_i w_i^\top \preceq \alpha \cdot I$. This completes the proof of Theorem 4.2, modulo the proof of Claim 4.4.

In Appendix C.1, we show that Theorem 4.2 can be generalized from a decomposition of the identity into rank-one matrices $w_i w_i^\top$ to a decomposition into matrices of arbitrary rank. The proof of Claim 4.4 is analogous to the proof of Claim C.2. We remark that Theorem 4.3 is not known to have a generalization to matrices of arbitrary rank.

4.3 Thin trees

Let $G = (V, E)$ be a graph. For convenience we assume that $V = [n]$. The *cut* defined by $U \subseteq V$ is

$$\delta_G(U) = \delta(U) = \{uv \in E : \text{exactly one of } u \text{ and } v \text{ is in } U\}.$$

For a subgraph T of G , let $\delta_T(U)$ denote all edges of T with exactly one endpoint in U .

Definition 4.5. A subgraph T of G is called ϵ -*thin* if $|\delta_T(U)| \leq \epsilon \cdot |\delta_G(U)|$ for all $U \subseteq V$.

Conjecture 4.6 (Goddyn [28]). *Every graph with connectivity at least k has an $f(k)$ -thin spanning subtree, for some function f that vanishes as k tends to infinity.*

The crucial detail in this conjecture is that the function f should not depend on the size of the graph. The best progress on this conjecture for general graphs is as follows.

Theorem 4.7 (Asadpour et al. [3]). *Let G be a graph with n vertices and connectivity k . Then G has a $O\left(\frac{\log n}{k \log \log n}\right)$ -thin spanning subtree. Moreover, there is a randomized, polynomial time algorithm to construct such a tree.*

Now we define spectrally-thin trees and prove an analog of this theorem. The Laplacian of G is the symmetric matrix L_G with rows and columns indexed by V defined by

$$L_G := \sum_{uv \in E} (e_u - e_v)(e_u - e_v)^\top.$$

Definition 4.8. Let T be a spanning subtree of G and let L_T be the Laplacian of T . The tree T is ϵ -spectrally-thin if $L_T \preceq \epsilon L_G$.

Any tree that is ϵ -spectrally-thin is also ϵ -thin, because

$$|\delta_T(U)| = \chi(U)^\top L_T \chi(U) \leq \epsilon \cdot \chi(U)^\top L_G \chi(U) = \epsilon \cdot |\delta_G(U)|.$$

The converse is not true. Moreover, the connectivity hypothesis in Theorem 4.7 does not suffice¹ to obtain a good spectrally-thin tree. The proof is in Appendix C.2.1.

Theorem 4.9. *For every $n, k \geq 1$, there exists a weighted graph with n vertices and connectivity k that does not have an $o(\sqrt{n}/k)$ -spectrally-thin spanning subtree.*

Nevertheless, if we strengthen the connectivity lower bound to a lower bound on the effective conductances, then we have the following construction of spectrally-thin trees. For an edge $e = uv \in E$, the *effective resistance* in G between u and v is $R_e := (e_u - e_v)^\top L_G^+(e_u - e_v)$. The *effective conductance* in G between u and v is $C_e := 1/R_e$.

Theorem 4.10. *Let G be a graph with n vertices such that $\kappa \leq C_e$ for every edge e . Then there is a polynomial time algorithm (either randomized or deterministic) to construct a $O\left(\frac{\log n}{\kappa \log \log n}\right)$ -spectrally-thin spanning subtree of G .*

Theorem 4.10 follows directly from Theorem 4.1, letting \mathbf{M} be the graphic matroid corresponding to G . It also follows from Theorem 4.2, as we show in Appendix C.2. That viewpoint is advantageous, since Theorem 4.3 then immediately implies

Theorem 4.11. *Let G be a graph with n vertices such that $\kappa \leq C_e$ for every edge e . Then G has a $O(1/\kappa)$ -spectrally-thin spanning subtree.*

We are not aware of any formal connection between Theorem 4.11 and Conjecture 4.6 or the traveling salesman problem.

Although Theorem 4.7 and Theorem 4.10 are formally incomparable, it is worth understanding their similarities and differences. Both results have a seemingly suboptimal factor of $\log n / \log \log n$. Theorem 4.7 requires only a connectivity lower bound, which is important in applications [28, 3], but the resulting tree is thin, not spectrally-thin; also, their algorithm is randomized. Theorem 4.10 requires a conductance lower bound (which is stronger than a connectivity lower bound), but the resulting tree is spectrally-thin (which is stronger than being thin); also, our algorithm can be

¹ This result was independently observed by M. de Carli Silva, N. Harvey and C. Sato, and by M. Goemans [29], using slightly different examples.

made deterministic. The use of randomization seems quite inherent in the algorithms [3, 20] for Theorem 4.7, as the thinness condition involves controlling exponentially many cuts, which seems difficult to accomplish by a deterministic, polynomial-time algorithm.

The quantities k and κ can be related in certain classes of graphs. We say that a family of graphs has *nearly equal resistances* if there is a constant c (independent of the number of vertices) such that $R_e \leq cR_f$ for all edges e, f . For example, any Ramanujan graph has nearly equal resistances. Edge-transitive graphs, such as hypercubes, have nearly equal (in fact, exactly equal) resistances.

Corollary 4.12. *Let G be a graph with n vertices, nearly equal resistances, and connectivity k . Then there is a deterministic, polynomial time algorithm to construct a $O\left(\frac{\log n}{k \log \log n}\right)$ -spectrally-thin tree of G .*

The proof is in Appendix C.2.

4.4 Column-subset selection

Column-subset selection is an important topic in numerical linear algebra [13, 51, 22, 12]. Similar questions are considered in operator theory [10, 11, 47, 51, 58]. In this section we prove a non-isotropic analog of Theorem 4.2, which gives a new result on column-subset selection. For a real matrix A , let $\|A\|_F = \sqrt{\text{tr } A^T A}$ denote its Frobenius norm. The *stable rank* of A is $\text{st.rank}(A) := \|A\|_F^2 / \|A\|^2$.

Theorem 4.13. *Let A be a real matrix of size $n \times m$ whose columns are denoted a_1, \dots, a_m . Suppose that $\|a_i\| = 1 \ \forall i$. Then there is a deterministic, polynomial time algorithm to compute $S \subseteq [m]$ of size $|S| \geq \lfloor \text{st.rank}(A) \rfloor$ such that $\{a_i : i \in S\}$ is linearly independent, and $\|\sum_{i \in S} a_i a_i^T\| \leq O(\log n / \log \log n)$.*

This is optimal with respect to $|S|$ as it can happen that $\text{st.rank}(A) = \text{rank}(A)$, in which case $\{a_i : i \in S\}$ is linearly dependent whenever $|S| > \text{st.rank}(A)$.

We now prove Theorem 4.13 using Theorem 4.1. Note that $\|A\|_F^2 = m$. Let $p \in \mathbb{R}^m$ be the vector with $p_i = \left[m / \|A\|^2 \right] / m$ for all i . Note that $\sum_i p_i = \lfloor \text{st.rank}(A) \rfloor$. We claim that p can be viewed as a “fractional set” of linearly independent vectors of size $\lfloor \text{st.rank}(A) \rfloor$. Formally, for any set $T \subseteq [m]$, let A_T denote the submatrix of A consisting of the columns in T . Define the following family of sets

$$\mathcal{B} = \{ I \subseteq [m] : \text{rank } A_I = |I| = \lfloor \text{st.rank}(A) \rfloor \}.$$

Then \mathcal{B} is the base family of the linear matroid corresponding to A , truncated to rank $\lfloor \text{st.rank}(A) \rfloor$. Let \mathbf{M} denote that matroid and let P denote its base polytope.

Claim 4.14. $p \in P$.

The proof is in Appendix C.3. Given this claim, all that remains is an easy application of Theorem 4.1. Define $A_i = a_i a_i^T$, $B = I$ and

$$Q = P \cap \left\{ x \in \mathbb{R}^m : \sum_i x_i A_i \preceq B \right\}.$$

We have $p \in Q$ by Claim 4.14 and the fact that

$$\sum_i p_i A_i \preceq \sum_i \frac{A_i}{\|A\|^2} = \frac{AA^T}{\|A\|^2} \preceq I = B.$$

Note that $A_i = a_i a_i^T \preceq I = B$. Theorem 4.1 gives a deterministic algorithm to construct an extreme point $\chi(S)$ of P for which $\sum_{i \in S} A_i \preceq \alpha \cdot B$, with $\alpha = O(\log n / \log \log n)$. Since S is a base of \mathbf{M} , the set $\{a_i : i \in S\}$ has rank equal to $|S| = \lfloor \text{st.rank}(A) \rfloor$. This completes the proof of Theorem 4.13.

Acknowledgements. N. Harvey thanks Joel Friedman and Mohit Singh for numerous enlightening discussions. We thank Isaac Fung for collaborating at a preliminary stage of this work. We also thank Christos Boutsidis, Joseph Cheriyan, Satoru Fujishige, Michel Goemans, Mary Beth Ruskai, Nikhil Srivastava, Joel Tropp, Roman Vershynin and Jan Vondrák for helpful discussions and suggestions.

References

- [1] A. A. Ageev and M. Sviridenko. Pipage rounding: A new method of constructing algorithms with proven performance guarantee. *Journal of Combinatorial Optimization*, 8, 2004.
- [2] R. Ahlswede and A. Winter. Strong converse for identification via quantum channels. *IEEE Transactions on Information Theory*, 48(3):569–579, Mar. 2002.
- [3] A. Asadpour, M. X. Goemans, A. Madry, S. Oveis Gharan, and A. Saberi. An $O(\log n / \log \log n)$ -approximation algorithm for the asymmetric traveling salesman problem. In *SODA*, 2010.
- [4] N. Bansal. Refinements of randomized rounding. [Invited lecture](#) at the Landelijk Netwerk Mathematische Beslikunde Lunteren Conference, 2012.
- [5] J. Batson, D. A. Spielman, and N. Srivastava. Twice-Ramanujan sparsifiers. In *STOC*, 2009.
- [6] R. Bhatia. *Matrix Analysis*. Springer, 1996.
- [7] R. Bhatia. The logarithmic mean. *Resonance*, 13, 2008.
- [8] I. Bjelaković and R. Siegmund-Schultze. Quantum Stein’s lemma revisited, inequalities for quantum entropies, and a concavity theorem of Lieb. arXiv:quant-ph/0307170.
- [9] B. Bollobás. *Modern Graph Theory*. Springer, 1998.
- [10] J. Bourgain and L. Tzafriri. Invertibility of “large” submatrices with applications to the geometry of Banach spaces and harmonic analysis. *Israel J. Math*, 57, 1987.
- [11] J. Bourgain and L. Tzafriri. On a problem of Kadison and Singer. *J. Reine. Angew. Math.*, 420:1–43, 1991.
- [12] C. Boutsidis, P. Drineas, and M. Magdon-Ismail. Near optimal column-based matrix reconstruction. In *FOCS*, 2011.
- [13] C. Boutsidis, M. W. Mahoney, and P. Drineas. An improved approximation algorithm for the column subset selection problem. In *SODA*, 2009.
- [14] S. C. Boyd and W. R. Pulleyblank. Optimizing over the subtour polytope of the travelling salesman problem. *Mathematical Programming*, 49:163–187, 1991.
- [15] G. Calinescu, C. Chekuri, M. Pál, and J. Vondrák. Maximizing a monotone submodular function subject to a matroid constraint. *SIAM Journal on Computing*, 40, 2011.
- [16] E. Carlen. Trace inequalities and quantum entropy: An introductory course. In R. Sims and D. Ueltschi, editors, *Entropy and the Quantum*, volume 529 of *Contemporary Mathematics*. American Mathematical Society, 2010.

- [17] B. C. Carlson. The logarithmic mean. *The American Mathematical Monthly*, 79(6):615–618, 1972.
- [18] P. G. Casazza and J. C. Tremain. The Kadison-Singer problem in mathematics and engineering. *Proceedings of the National Academy of Sciences*, 103(7):2032–2039, 2006.
- [19] C. Chekuri, J. Vondrák, and R. Zenklusen. Dependent randomized rounding for matroid polytopes and applications, 2009. arXiv:0909.4348.
- [20] C. Chekuri, J. Vondrák, and R. Zenklusen. Dependent randomized rounding via exchange properties of combinatorial structures. In *FOCS*, pages 575–584, 2010.
- [21] R. Y. Chen, A. Gittens, and J. A. Tropp. The masked sample covariance estimator: an analysis using matrix concentration inequalities. *Information and Inference*, 2012.
- [22] A. Deshpande and L. Rademacher. Efficient volume sampling for row/column subset selection. In *FOCS*, 2010.
- [23] E. G. Effros. A matrix convexity approach to some celebrated quantum inequalities. *Proceedings of the National Academy of Sciences*, 106(4):1006–1008, 2009.
- [24] H. Epstein. Remarks on Two Theorems of E. Lieb. *Communications in Mathematical Physics*, 31, 1973.
- [25] S. Fujishige. *Submodular Functions and Optimization*. Elsevier, second edition, 2005.
- [26] R. Gandhi, S. Khuller, S. Parthasarathy, and A. Srinivasan. Dependent rounding and its applications to approximation algorithms. *Journal of the ACM*, 53(3):324–360, 2006.
- [27] F. Gesztesy and E. Tsekanovskii. On matrix-valued Herglotz functions. *Math. Nachrichten*, 218:61–138, 2000.
- [28] L. A. Goddyn. Some open problems I like. Available at <http://www.math.sfu.ca/~goddyn/Problems/problems.html>.
- [29] M. Goemans. Thin spanning trees. Lecture at “Graph Theory @ Georgia Tech: Conference Honoring the 50th Birthday of Robin Thomas”, 2012.
- [30] R. A. Horn and C. R. Johnson. *Topics in Matrix Analysis*. Cambridge University Press, 1991.
- [31] T. Leighton, S. Rao, and A. Srinivasan. Multicommodity flow and circuit switching. In *Proc. Hawaii International Conference on System Sciences (HICSS)*, 1998.
- [32] E. H. Lieb. Convex trace functions and the Wigner-Yanase-Dyson conjecture. *Advances in Mathematics*, 11:267–288, 1973.
- [33] M. W. Mahoney. Randomized algorithms for matrices and data. arXiv:1104.5557.
- [34] A. Marcus, D. A. Spielman, and N. Srivastava. Interlacing families II: Mixed characteristic polynomials and the Kadison-Singer problem, June 2013. arXiv:1306.3969.
- [35] M. A. Nielsen and I. L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2000.
- [36] M. Ohya and D. Petz. *Quantum Entropy and Its Use*. Springer-Verlag, 1993.

- [37] R. I. Oliveira. Concentration of the adjacency matrix and of the Laplacian in random graphs with independent edges, 2010. arXiv:0911.0600.
- [38] S. Oveis Gharan and A. Saberi. Asymmetric traveling salesman problem on graphs with bounded genus. In *SODA*, 2011.
- [39] A. Panconesi and A. Srinivasan. Randomized distributed edge coloring via an extension of the Chernoff-Hoeffding bounds. *SIAM Journal on Computing*, 26(2):350–368, 1997.
- [40] P. Raghavan. Probabilistic construction of deterministic algorithms: Approximating packing integer programs. *Journal of Computer and System Sciences*, 37, 1988.
- [41] P. Raghavan and C. D. Thompson. Randomized rounding: a technique for provably good algorithms and algorithmic proofs. *Combinatorica*, 7:365–374, 1987.
- [42] B. Recht. A simpler approach to matrix completion. *Journal of Machine Learning Research*, 12, 2011.
- [43] M. Rudelson. Almost orthogonal submatrices of an orthogonal matrix. *Israel J. Math.*, 111(1):143–155, 1999.
- [44] M. Rudelson and R. Vershynin. Sampling from large matrices: An approach through geometric functional analysis. *Journal of the ACM*, 54(4), 2007.
- [45] W. Rudin. *Real and Complex Analysis*. McGraw-Hill, third edition, May 1986.
- [46] D. A. Spielman and N. Srivastava. Graph sparsification by effective resistances. In *STOC*, 2008.
- [47] D. A. Spielman and N. Srivastava. An elementary proof of the restricted invertibility theorem. *Israel J. Math*, 190(1):83–91, Aug. 2012.
- [48] A. Srinivasan. Approximation algorithms via randomized rounding: a survey. In M. Karonski and H. J. Promel, editors, *Lectures on Approximation and Randomized Algorithms*, pages 9–71. Polish Scientific Publishers PWN, 1999.
- [49] A. Srinivasan. Distributions on level-sets with applications to approximation algorithms. In *FOCS*, 2001.
- [50] A. Srinivasan. Pessimistic estimators, 2003. [Lecture notes](#) for CMSC858T “Randomized Algorithms”.
- [51] J. A. Tropp. Column subset selection, matrix factorization, and eigenvalue optimization. In *SODA*, 2009.
- [52] J. A. Tropp. User-friendly tail bounds for sums of random matrices. *Foundations of Computational Mathematics*, 2011.
- [53] J. A. Tropp. From the joint convexity of quantum relative entropy to a concavity theorem of Lieb. *Proc. Amer. Math. Soc.*, 140(5):1757–1760, 2012.
- [54] R. Vershynin. Introduction to the non-asymptotic analysis of random matrices. In Y. Eldar and G. Kutyniok, editors, *Compressed Sensing: Theory and Applications*, chapter 5, pages 210–268. Cambridge University Press, 2012.

- [55] N. Weaver. The Kadison-Singer problem in discrepancy theory. *Discrete Mathematics*, 278:227–239, 2004.
- [56] A. Wigderson and D. Xiao. Derandomizing the Ahlswede-Winter matrix-valued Chernoff bound using pessimistic estimators and applications. *Theory of Computing*, 4(3), 2008.
- [57] D. P. Williamson and D. B. Shmoys. *The Design of Approximation Algorithms*. Cambridge University Press, 2011.
- [58] P. Youssef. Restricted invertibility and the Banach-Mazur distance to the cube, 2012. arXiv:1206.0654v1.

A Pipage Rounding

Let p be a point in the matroid polytope P and assume that g satisfies (1). Delete all coordinates of p that are equal to zero and consider the residual problem. It is well-known that, for any such point p , there exists a chain of sets $\emptyset = C_0 \subseteq C_1 \subseteq \dots \subseteq C_k \subseteq [m]$ whose corresponding constraints of P span the constraints that are tight at p . If $|C_i \setminus C_{i-1}| = 1$ for every i then these give m linearly independent tight constraints, so the point p is an extreme point. Otherwise there is some set C_i , $i \geq 1$, for which $|C_i \setminus C_{i-1}| > 1$. In this case p is not an extreme point. To see this, let a and b be distinct elements of $C_i \setminus C_{i-1}$. Note that the point $p + z(e_a - e_b)$ satisfies all the constraints that are tight at p . So, for all z in some open neighborhood of 0, the point $p + z(e_a - e_b)$ is still feasible for P .

Define

$$\ell = \min \{ z \in \mathbb{R} : p + z(e_a - e_b) \in P \}$$

and

$$u = \max \{ z \in \mathbb{R} : p + z(e_a - e_b) \in P \}.$$

Define

$$p^\ell = p + \ell(e_a - e_b) \quad \text{and} \quad p^u = p + u(e_a - e_b).$$

Since $g(p + z(e_a - e_b))$ is concave, we must have either

$$g(p^\ell) \leq g(p) \quad \text{or} \quad g(p^u) \leq g(p).$$

Furthermore, both p^ℓ and p^u lie on a lower-dimensional face than p does. So starting from some initial $p^0 \in P$, m iterations suffice to find an extreme point \hat{p} of P with $g(\hat{p}) \leq g(p^0)$.

The randomized version of pipage rounding does not even need access to the function g . Instead, it simply chooses the next point p' to be p^ℓ with probability $\frac{u}{u-\ell}$, or p^u with probability $\frac{-\ell}{u-\ell}$. This ensures that $\mathbb{E}[p'] = p$, and the concavity of g yields $\mathbb{E}[g(p')] \leq g(p)$. Thus applying this procedure to some initial point $p_0 \in P$ until an extreme point \hat{p} is obtained, \hat{p} satisfies $\mathbb{E}[\hat{p}] = p_0$ and $\mathbb{E}[g(\hat{p})] \leq g(p_0)$.

B Proofs of concavity under swaps

Proof of Claim 3.4. We can rewrite

$$g(x) = e^{-\theta t} \cdot \prod_i (1 + x_i(e^{\theta w_i} - 1)).$$

Rewriting $g(x + z(e_a - e_b))$ in this way, all factors are non-negative and only two of them depend on z , so for some $c \geq 0$

$$\begin{aligned} & \frac{d^2}{dz^2} g(x + z(e_a - e_b)) \\ &= c \cdot \frac{d^2}{dz^2} \left((1 + (x_a + z)(e^{\theta w_a} - 1)) (1 + (x_b - z)(e^{\theta w_b} - 1)) \right) \\ &= c \cdot \left(-2(e^{\theta w_a} - 1)(e^{\theta w_b} - 1) \right). \end{aligned}$$

This is non-positive so g is concave under swaps. □

Proof of Claim 3.6. Recall that $\theta < 0$. Define $h : \{0, 1\}^m \rightarrow \mathbb{R}$ by $h(X) = e^{\theta f(X)}$. By Claim B.1, h is a supermodular function. Its multilinear extension is

$$H(x) = \mathbb{E}_{X \sim \mathcal{D}(x)} [h(X)].$$

Since $-h$ is submodular, it follows from results of Calinescu et al. [15] that $\frac{\partial^2 H}{\partial x_i \partial x_j} \geq 0$ for any $i, j \in [m]$. Since $g(x) = e^{-\theta t} \cdot H(x)$, the second derivative of

$$z \mapsto g(x + z(e_i - e_j))$$

is non-positive. Thus g is concave under swaps. \square

Claim B.1. *Let $f : 2^{[m]} \rightarrow \mathbb{R}$ be non-decreasing and submodular. Let $g : \mathbb{R} \rightarrow \mathbb{R}$ be non-increasing and convex. Then $g \circ f$ is supermodular.*

Proof. We require the following property of convex functions. Suppose a, b, c, d satisfy

$$a \leq \min\{b, c\} \leq \max\{b, c\} \leq d. \quad (8)$$

Then any function g that is convex on $[a, d]$ satisfies

$$\frac{g(d) - g(c)}{d - c} \geq \frac{g(b) - g(a)}{b - a}. \quad (9)$$

Fix any $A \subseteq B \subseteq [m]$, and an element $x \in [m] \setminus B$. Define

$$\begin{aligned} a &:= f(A), & b &:= f(A + x), & c &:= f(B), \\ d &:= f(B) + f(A + x) - f(A), & e &:= f(B + x). \end{aligned}$$

Since f is non-decreasing, (8) holds. Since f is submodular, $e \leq d$ holds. Since g is non-increasing, $g(e) \geq g(d)$. Combining that with (9) and the observation that $d - c = b - a$, we obtain

$$g(e) - g(c) \geq g(d) - g(c) \geq g(b) - g(a)$$

That is,

$$g(f(B + x)) - g(f(B)) \geq g(f(A + x)) - g(f(A)),$$

so $g \circ f$ is supermodular. \square

Proof of Lemma 3.9. We will show that

$$\forall x \in (0, 1)^m, \forall a, b \in [m], \quad z \mapsto g_{t, \theta}(x + z(e_a - e_b)) \text{ is concave.}$$

The boundary of $[0, 1]^m$ is handled by continuity. Note that

$$\mathbb{E}_{X \sim \mathcal{D}(x)} \left[e^{\theta X_i M_i} \right] = x_i \cdot e^{\theta M_i} + (1 - x_i) \cdot I =: C_i.$$

Adding z (sufficiently small) to the sampling probability of coordinate i , the expectation becomes

$$\begin{aligned} \mathbb{E}_{X \sim \mathcal{D}(x + ze_i)} \left[e^{\theta X_i M_i} \right] &= (x_i + z) \cdot e^{\theta M_i} + (1 - x_i - z) \cdot I \\ &= C_i + z \underbrace{\left(e^{\theta M_i} - I \right)}_{=: K_i}. \end{aligned}$$

Note that $C_i \succeq I$ and $K_i \succeq 0$ because $M_i \succeq 0$ and $\theta > 0$. Furthermore, the matrices C_i and K_i commute since any eigenbasis for M_i is also an eigenbasis of C_i and K_i .

To finish the proof we must show that, for distinct $a, b \in [m]$,

$$z \mapsto \text{tr} \exp \left(\log (C_a + zK_a) + \log (C_b - zK_b) + \underbrace{\sum_{i \notin \{a,b\}} \log \mathbb{E}_{X \sim \mathcal{D}(x)} \left[e^{\theta X_i M_i} \right]}_{=: L} \right)$$

is concave in a neighborhood of 0. This follows from Theorem 3.12. \square

C Proofs of Applications

C.1 Rounding decompositions of the identity

Here, we give a generalization of Theorem 4.2 to a decomposition of the identity into matrices of arbitrary rank.

Theorem C.1. *Let $X_1, \dots, X_m \in \mathbb{S}_+^n$ satisfy $\sum_{i=1}^m X_i = I$. Then there exists a subset $S \subseteq [m]$ with $|S| \leq n$ such that $\sum_{i \in S} X_i / \text{tr} X_i$ has full rank and maximum eigenvalue at most $\alpha = O(\log n / \log \log n)$.*

Proof. Let V_i be a matrix such that $X_i = V_i V_i^\top$. Define the function $r : 2^{[m]} \rightarrow \mathbb{Z}_+$ by

$$r(J) = \text{rank} \left(\sum_{j \in J} V_j V_j^\top \right) = \text{rank} V_J,$$

where V_J is the matrix obtained by concatenating in any order all columns from the matrices $\{V_j : j \in J\}$. It is well-known that such a function r is:

- *Normalized:* $r(\emptyset) = 0$,
- *Monotone:* $r(I) \leq r(J)$ whenever $I \subseteq J$, and
- *Submodular:* $r(I) + r(J) \geq r(I \cup J) + r(I \cap J)$ for all $I, J \subseteq [m]$.

For any normalized, monotone, submodular function $f : 2^{[m]} \rightarrow \mathbb{R}$, its *base polytope* is defined to be

$$B(f) := \left\{ x \in \mathbb{R}_+^m : x(J) \leq f(J) \ \forall J \subseteq [m], \text{ and } x([m]) = f([m]) \right\}.$$

Define the vector $p \in \mathbb{R}^m$ by $p_i = \text{tr} X_i$. Note that $p \geq 0$ and $\sum_i p_i = \text{tr}(\sum_i X_i) = n$, so we can think of p as defining a “fractional multiset” of n matrices. Intuitively, we want to “round” the coordinates of p to integers. To that end, define the polytope

$$P' := B(r) \cap \{x : \lfloor p \rfloor \leq x \leq \lceil p \rceil\},$$

where $\lfloor p \rfloor$ and $\lceil p \rceil$ respectively denote the component-wise floor and ceiling of the vector $p \in \mathbb{R}^m$. The polytope P' is not necessarily a matroid polytope; for example, a vector in P' could have a coordinate strictly greater than 1.

Claim C.2. $p \in P'$.

Claim C.3. $P := \{x - \lfloor p \rfloor : x \in P'\}$ is a matroid base polytope.

Claim C.2 is proven below. Claim C.3 is a folklore result that can be derived using reductions and contractions of submodular functions [25, §3.1(b)]; see also Fujishige’s remarks on crossing submodular functions [25, Eq. (3.97)].

Define $A_i = X_i / \text{tr } X_i$, $B = I$ and

$$Q := P \cap \left\{ x \in \mathbb{R}^m : \sum_i x_i A_i \preceq B \right\}.$$

Setting $x = p - \lfloor p \rfloor$, we have $x \in P$ by Claim C.2 and

$$\sum_i x_i A_i \preceq \sum_i p_i \frac{X_i}{\text{tr } X_i} = \sum_i X_i = B,$$

so $x \in Q$.

Since $\text{tr } A_i = 1$, we have $A_i \preceq B$. Applying Theorem 4.1, we obtain a vector $\hat{x} \in \{0, 1\}^n$ that is an extreme point of P , and for which $\sum_i \hat{x}_i A_i \preceq \alpha B$. Let S be the support of \hat{x} . Note that $\hat{x} + \lfloor p \rfloor \in P'$. So

$$|S| = \sum_{i=1}^m \hat{x}_i \leq \sum_{i=1}^m (\hat{x}_i + \lfloor p_i \rfloor) \leq r([m]) = n$$

and $\sum_{i \in S} X_i / \text{tr } X_i \preceq \alpha B$ as required. \square

Proof of Claim C.2. The box constraint $\lfloor p \rfloor \leq p \leq \lceil p \rceil$ is trivially satisfied. We have noted above that $\sum_i p_i = n$, so the constraint $p([m]) \leq r([m]) = n$ is also satisfied.

It remains to show that $\sum_{i \in I} p_i \leq r(I)$ for all I . For any positive semidefinite matrix, the average of the non-zero eigenvalues is a lower bound on the maximum eigenvalue, so

$$\frac{\text{tr}(\sum_{i \in I} X_i)}{\text{rank}(\sum_{i \in I} X_i)} \leq \left\| \sum_{i \in I} X_i \right\| \leq \left\| \sum_{i=1}^m X_i \right\| = 1.$$

Thus $\sum_{i \in I} p_i = \text{tr}(\sum_{i \in I} X_i) \leq \text{rank}(\sum_{i \in I} X_i) = r(I)$. This proves that $p \in P$. \square

C.2 Thin trees

Proof of Theorem 4.10. Recall the notation defined in Section 2. For $e = uv \in E$, define vectors $x_e = L_G^{+/2}(e_u - e_v)$ and $w_e = x_e / \|x_e\|$. Then $R_e = \|x_e\|^2$; let $p_e = R_e / (n - 1)$. It is well-known [9] that the vector of effective resistances describes the edge marginals of the uniform spanning tree, and hence that $\sum_e p_e = 1$. Then, following the argument of Spielman and Srivastava [46],

$$\begin{aligned} \sum_{e \in E} p_e w_e w_e^\top &= \frac{1}{n-1} \sum_{e \in E} x_e x_e^\top \\ &= \frac{1}{n-1} L_G^{+/2} \left(\sum_{e \in E} (e_u - e_v)(e_u - e_v)^\top \right) L_G^{+/2} \\ &= \frac{1}{n-1} I_{\text{im } L_G}. \end{aligned}$$

We view the vectors $\{w_e : e \in E\}$ as $(n - 1)$ -dimensional vectors in their linear span and apply Theorem 4.2. This gives a set $T \subseteq E$ of size $n - 1$ such that $\{w_e : e \in T\}$ is linearly independent and

$$\sum_{e \in T} w_e w_e^\top \preceq O(\log n / \log \log n) \cdot I_{\text{im } L_G}.$$

The first two conditions imply that the edges in T form a spanning tree on the vertex set V . Then since $R_e = \|x_e\|^2$, we have

$$\sum_{e \in T} \frac{x_e x_e^\top}{R_e} \preceq O(\log n / \log \log n) \cdot I_{\text{im } L_G}.$$

Equivalently,

$$\sum_{uv \in T} \frac{(e_u - e_v)(e_u - e_v)^\top}{R_{uv}} \preceq O(\log n / \log \log n) \cdot L_G.$$

Since we assume that $\kappa \leq C_e = 1/R_e$ for every edge e , we obtain

$$L_T = \sum_{uv \in T} (e_u - e_v)(e_u - e_v)^\top \preceq O\left(\frac{\log n}{\kappa \log \log n}\right) \cdot L_G.$$

So T is $O\left(\frac{\log n}{\kappa \log \log n}\right)$ -spectrally-thin. \square

Proof of Corollary 4.12. By the nearly equal resistances assumption, $R_e = O\left(\frac{n-1}{|E|}\right)$ for every edge e . On the other hand, the connectivity k is at most the average degree, which is $2|E|/n$. Thus $R_e = O(1/k)$ for every edge e . The result now follows from Theorem 4.10. \square

C.2.1 Proof of Theorem 4.9

Assume n is a multiple of 4. We define a graph that is related to an example of Boyd and Pulleyblank [14, p. 180]. There are two disjoint cycles, each of length $n/2$. Let us number the vertices in the first cycle as $1, \dots, n/2$ and the vertices in the second cycle as $n/2 + 1, \dots, n$. Add a matching where the i^{th} edge connects the i^{th} vertex in the first cycle and the i^{th} vertex in the second cycle. The edges in the cycles each have weight $w_c := k/2$ and the edges in the matching each have weight $w_m := 2k/n$. Obviously this weighted graph has connectivity at least k .

Let T be any subtree of G , without any weights on the edges of T .

Claim C.4. *Suppose that T uses only a single matching edge. There exists a vector z such that*

$$\frac{z^\top L_T z}{z^\top L_G z} = \Omega\left(\frac{\sqrt{n}}{k}\right).$$

Proof. Without loss of generality, $\{n/4, 3n/4\}$ be the matching edge used by T . Let $\alpha = n^{-0.5}$ and $c = 1 - \alpha$. Define the vector z where

$$z_i = \begin{cases} c^{|n/4-i|} & (i \leq n/2) \\ 0 & (i > n/2). \end{cases}$$

Numerator: The numerator is $z^\top L_T z = \sum_{uv \in E} (z_u - z_v)^2 \geq (z_{n/4} - z_{3n/4})^2 = 1$.

Denominator: To evaluate $z^\top L_G z$, we separately consider the cycle edges and matching edges. The contribution from the matching edges is

$$C_m := w_m \cdot \sum_{i=1}^{n/2} (z_i - z_{n/2+i})^2 < 2w_m \cdot \sum_{i \geq 0} c^{2i} < \frac{2w_m}{1-c} = \frac{2w_m}{\alpha}.$$

The contribution from the cycle edges is

$$\begin{aligned} C_c &:= w_c \sum_{i=2}^{n/2} (z_{i-1} - z_i)^2 + w_c (z_1 - z_{n/2})^2 < 2w_c \sum_{i \geq 1} (c^{i-1} - c^i)^2 \\ &= 2w_c (1-c)^2 \sum_{i \geq 0} c^{2i} = 2w_c \frac{(1-c)^2}{1-c^2} < 2w_c \frac{(1-c)^2}{1-c} = 2w_c \alpha. \end{aligned}$$

Since $\alpha = n^{-0.5}$, we get $C_m = O(k/\sqrt{n})$ and $C_c = O(k/\sqrt{n})$, so $z^\top L_G z = O(k/\sqrt{n})$. \square

Claim C.5. *Suppose that T uses $m > 1$ matching edges. There exists a vector z such that*

$$\frac{z^\top L_T z}{z^\top L_G z} = \Omega\left(\frac{\sqrt{n}}{k}\right).$$

Proof. Let the matching edges used by T be $\{a_1, b_1\}, \{a_2, b_2\}, \dots, \{a_m, b_m\}$. Define the vector z by

$$z_i = \begin{cases} c^{\min_j d_1(i,j)} & (i \leq n/2) \\ 0 & (i > n/2) \end{cases}$$

where d_1 denotes distance in the first cycle.

Numerator: As before, every matching edge used by T contributes at least 1, so $z^\top L_T z \geq m$.

Denominator: Obviously $z^\top L_G z$ is no more than m times what it would be if T used only a single matching edge. That is, $z^\top L_G z \leq O(mk/\sqrt{n})$. \square

C.3 Column-subset selection

Proof of Claim 4.14. The proof is analogous to Claim 4.4. As before, let $r : 2^{[m]} \rightarrow \mathbb{Z}_+$ be defined by

$$r(S) := \dim \text{span} \{a_i : i \in S\} = \text{rank } A_S A_S^\top.$$

Then

$$P := \left\{ x \in \mathbb{R}_+^m : x(J) \leq r(J) \quad \forall J \subseteq [m], \text{ and } x([m]) = \lfloor \text{st. rank}(A) \rfloor \right\}.$$

For any set $J \subseteq [m]$, we have

$$\begin{aligned} r(J) &= \text{rank } A_J A_J^\top \geq \frac{\text{tr } A_J A_J^\top}{\|A_J A_J^\top\|} \geq \frac{|J|}{\|A\|^2} \\ &\geq \left\lfloor \frac{m}{\|A\|^2} \right\rfloor \frac{|J|}{m} = p(J). \end{aligned}$$

Since $\sum_i p_i = \lfloor \text{st. rank}(A) \rfloor$, we have $p \in P$. \square

D Proof of Theorem 3.12

The outline of this proof follows a proof of Lieb's theorem presented by Epstein [24]. Epstein's proof proceeds via complex analytic techniques, and in particular makes use of some powerful results involving *Herglotz functions* (see, e.g., [6, 27]). While an effort has been made to make the treatment here accessible, a modicum of complex analysis will be assumed; a standard reference is [45].

For a complex number z , let $\Re z$ and $\Im z$ respectively denote the real and imaginary parts of z . Let $\mathbb{C}_{++} = \{z \in \mathbb{C} \mid \Im z > 0\}$ denote the open upper half-plane, and \mathbb{C}_+ the closed upper half-plane. Define \mathbb{C}_{--} and \mathbb{C}_- in the obvious corresponding way.

Definition D.1. A function $g : \mathbb{C}_{++} \rightarrow \mathbb{C}$ is called a *Herglotz function* (or *Pick function*) if it is analytic on \mathbb{C}_{++} and $g(\mathbb{C}_{++}) \subseteq \mathbb{C}_{++}$.

For example the map $z \mapsto az + b$ is Herglotz if $a \in \mathbb{R}_+$ and $b \in \mathbb{C}_+$. The maps $z \mapsto -1/z$ and $z \mapsto \log z$ are also Herglotz.

A key reason that Herglotz functions will be useful is the following classical theorem (see, e.g., [6, Eq. V.42] or [30, p. 542]).

Theorem D.2 (Herglotz-Nevalinna-Riesz representation theorem). *For any Herglotz function g , there exists $a \in \mathbb{R}$, $b \in \mathbb{R}_+$ and a positive Borel measure μ on \mathbb{R} , with $\int_{\mathbb{R}} \frac{1}{t^2+1} d\mu(t) < \infty$, s.t.*

$$g(z) = a + bz + \int_{\mathbb{R}} \left(\frac{1}{t-z} - \frac{t}{1+t^2} \right) d\mu(t) \quad \forall z \in \mathbb{C}_{++}. \quad (10)$$

Roughly speaking, this provides a description of a Herglotz function through its boundary (the real line); since the function may diverge as it approaches the real line, the generality of a measure (which may have atoms) is needed.

The relevance of this theorem to our purposes comes from the following:

Lemma D.3 (Implicit in [24]). *Let D be a domain² in \mathbb{C} containing $\mathbb{C}_{--} \cup \{0\}$. Suppose $f : D \rightarrow \mathbb{C}$ is analytic, its restriction to $D \cap \mathbb{R}$ is real-valued, and moreover the function g on \mathbb{C}_{++} defined by $g(z) = zf(1/z)$ is a Herglotz function. Then the restriction of f to $D \cap \mathbb{R}$ is concave in some neighborhood of the origin.*

Proof. Since $0 \in D$ and D is open, there exists some $\tau > 0$ so that the interval $[-\tau, \tau] \subset D$. Let D' be the image of D under the map $z \mapsto 1/z$; so D' contains $\mathbb{C}_{++} \cup [\tau^{-1}, \infty) \cup (-\infty, -\tau^{-1}]$. We may think of g as being defined on all of D' . Let μ be the positive Borel measure associated with g by Theorem D.2. This measure can be thought of as the limit of $\Im g(z)$ as z approaches the real line, in the appropriate distributional sense: this is known as the Stieltjes inversion formula; see, e.g., [6, Thm. V.4.12], [27, Thm. 2.2]. We will use only the following consequence:

If for some open interval $I \subseteq \mathbb{R}$,

$$\lim_{\epsilon \downarrow 0} \Im g(w + i\epsilon) = 0 \quad \text{for all } w \in I,$$

then $\mu(I) = 0$.

We deduce that μ is supported on $[-\tau^{-1}, \tau^{-1}]$, since $\lim_{\epsilon \downarrow 0} \Im g(w + i\epsilon) = \Im g(w) = 0$ for all $w \in D' \cap \mathbb{R}$.

Expressing f in terms of the Herglotz-Nevalinna-Riesz representation of g , we have that

$$f(z) = az + b + \int_{-\tau^{-1}}^{\tau^{-1}} \frac{z^2}{zt-1} d\mu(t).$$

(Note that the final term of (10) can be folded into the constant a — since μ is Borel and has bounded support, it is finite.) Now calculate the second derivative of f , considered as a real-valued function on $D \cap \mathbb{R}$:

$$f''(x) = \int_{-\tau^{-1}}^{\tau^{-1}} \frac{d^2}{dx^2} \left(\frac{x^2}{xt-1} \right) d\mu(t) = \int_{-\tau^{-1}}^{\tau^{-1}} \frac{2}{(xt-1)^3} d\mu(t).$$

So for all $x \in (-\tau, \tau)$, $f''(x) < 0$, and so f (as a real-valued function on $D \cap \mathbb{R}$) is concave in the neighborhood of 0. \square

²Recall that a domain is an open, connected set.

We will apply Lemma D.3 with f as in the statement of Theorem 3.12:

$$f(z) = \text{tr} \exp \left(L + \log(C_1 + zK_1) + \log(C_2 - zK_2) \right).$$

In order to extend our definition of \log beyond symmetric matrices, we use (again following [24]) the Cauchy integral description

$$\log C = \int_0^\infty \frac{1}{t+1} - (t+C)^{-1} dt;$$

this is well-defined as long as C has no nonpositive eigenvalues. As our domain D , we take $\mathbb{C}_{--} \cup B_\epsilon$, where B_ϵ is an open ball around the origin of radius $\epsilon := \frac{1}{2} \min\{\lambda_{\min}(C_1)\|K_1\|^{-1}, \lambda_{\min}(C_2)\|K_2\|^{-1}\}$. This ensures that

Lemma D.4. *The function f is well-defined and analytic on D .*

For convenience, we withhold the proof until the end of this section.

To deduce that f is concave by Lemma D.3, we must show that g defined by $g(z) = zf(1/z)$ is Herglotz. We have

$$\begin{aligned} g(z) &= z \cdot f(1/z) \\ &= z \text{tr} \exp \left(L + \log(C_1 + K_1/z) + \log(C_2 - K_2/z) \right) \\ &= \text{tr} \exp \left(\log(zI) + L + \log(C_1 + K_1/z) + \log(C_2 - K_2/z) \right) \\ &= \text{tr} \exp \left(L + \log(C_1 z + K_1) + \log(C_2 + (-1/z)K_2) \right). \end{aligned}$$

We will work with complex matrices for the remainder of this section, so let $M_n(\mathbb{C})$ denote the space of $n \times n$ complex matrices, and \mathbb{H}^n the space of $n \times n$ Hermitian matrices. We will make use of operator formalism on occasion; in particular, the identity I will generally be omitted, and so for a scalar $w \in \mathbb{C}$, wI will be written as simply w .

An arbitrary matrix $C \in M_n(\mathbb{C})$ has a unique decomposition $C = P + iQ$, where $P, Q \in \mathbb{H}^n$. This is obtained by taking $P = \frac{1}{2}(C + C^*)$ and $Q = \frac{1}{2i}(C - C^*)$, where C^* denotes the adjoint (conjugate transpose) of C . By analogy with the scalar ($n=1$) case, we say that P is the ‘‘real part’’ of C , denoted by $\Re C$, and that Q is the ‘‘imaginary part’’ of C , denoted by $\Im C$. (Note that this has nothing to do with the entry-wise real and imaginary parts of the matrix.)

This analogy to the scalar case provides a lot of helpful intuition, and so at this point we will sketch a version of the proof for $n = 1$. The full argument will follow the same essential steps, though the generalization is not completely straightforward. The scalar analog of a Hermitian matrix is a real number, and the scalar analog of a positive definite matrix is a positive number; so we consider the function $h : \mathbb{C} \rightarrow \mathbb{C}$ defined by

$$h(z) = \exp \left(l + \log(c_1 z + k_1) + \log(c_2 + (-1/z)k_2) \right),$$

with real parameters $l \in \mathbb{R}$, $k_1, k_2 \geq 0$ and $c_1, c_2 > 0$. Then

$$\Re \log(c_1 z + k_1) = \arg(c_1 z + k_1) \in (0, \arg z].$$

Similarly,

$$\Im \log(c_2 + (-1/z)k_2) = \arg(c_2 + (-1/z)k_2) \in [0, \arg(-1/z)).$$

Since $\arg(-1/z) = \pi - \arg z$, we obtain

$$\Im(l + \log(c_1 z + k_1) + \log(c_2 + (-1/z)k_2)) \in (0, \pi).$$

Since $\Im e^{a+ib} = e^a \sin b$ for $a, b \in \mathbb{R}$, we deduce that $\Im h(z) > 0$, as required.

We now resume the argument for the case $n > 1$. Define

$$\begin{aligned} \mathcal{I}_{++} &= \{C \in M_n(\mathbb{C}) : \Im C \succ 0\} \\ \mathcal{I}_+ &= \{C \in M_n(\mathbb{C}) : \Im C \succeq 0\}. \end{aligned}$$

Much of the argument revolves around noting that \mathcal{I}_{++} is closed under various operations. For example, if $C, A \in \mathcal{I}_{++}$ then clearly $A + C \in \mathcal{I}_{++}$. The following is less straightforward:

Lemma D.5 ([24, pp. 318–319]). *For any $C \in \mathcal{I}_{++}$,*

- (i) $-C^{-1} \in \mathcal{I}_{++}$, and
- (ii) $0 \prec \Im \log C \prec \pi$.

We refer to [24] for the proofs, but we again note the intuition by analogy with the $n = 1$ case, where C is just an element of \mathbb{C}_{++} . Then $C = re^{i\theta}$ for some $r > 0$ and $0 < \theta < \pi$; so $-C^{-1} = r^{-1}e^{i(\pi-\theta)} \in \mathbb{C}_{++}$ and $\log C = \log r + i\theta$.

A crucial lemma will be the following:

Lemma D.6. *Let $A, B \in \mathbb{H}^n$ satisfy $A, B \succeq 0$, where in addition at least one of A and B are strictly positive definite. Then for any $z \in \mathbb{C}_{++}$, $\log(A + Bz)$ is defined and*

$$0 \preceq \Im \log(A + Bz) \preceq \arg z.$$

Moreover, if $A \succ 0$, then the left inequality is strict, and if $B \succ 0$, the right inequality is strict.

Proof. We first observe that the conditions imply that $A + Bz$ has no nonpositive real eigenvalues, and hence that the logarithm is well defined. It suffices to show that $A + Bz$ is nonsingular, since we can apply the same argument to $A' + Bz$, where $A' = A + t$ for any $t \geq 0$.

If $B \succ 0$, then $B^{1/2}$ exists and is positive definite. Thus

$$A + Bz = B^{1/2}(\underbrace{B^{-1/2}AB^{-1/2}}_{=: Q} + z)B^{1/2}.$$

But Q is Hermitian (as can be seen since $B^{-1/2}$ and A are Hermitian) and so it has real spectrum; thus since $\Im z > 0$, 0 is not in the spectrum of $Q + z$. Hence $Q + z$ and so also $A + Bz$ are invertible.

If instead $A \succ 0$, then

$$A + Bz = zA^{1/2}(1/z + A^{-1/2}BA^{-1/2})A^{1/2},$$

and similar reasoning applies.

Suppose first that $B \succ 0$. Then $A + Bz \in \mathcal{I}_{++}$, and so by Lemma D.5 (ii) we immediately have that $\Im \log(A + Bz) \succ 0$. Now if $B \succeq 0$ but is not positive definite, then $B + \epsilon \succ 0$ for any $\epsilon > 0$, and so $\Im \log(A + (B + \epsilon)z) \succ 0$. Since $\log(A + Bz)$ is well defined, we have by continuity that

$$\Im \log(A + Bz) = \lim_{\epsilon \downarrow 0} \Im \log(A + (B + \epsilon)z) \succeq 0.$$

This completes the proof of the left inequality.

For the right inequality, suppose first that $A \succ 0$. Since $\arg z = \Im \log z$, our goal is to show that

$$\Im(\log z - \log(A + Bz)) \succ 0,$$

or equivalently (using that $A + Bz$ is nonsingular)

$$\Im \log((A/z + B)^{-1}) \succ 0.$$

Now since $-1/z \in \mathbb{C}_{++}$, it follows that $\Im(-A/z) \succ 0$. Since $\Im B = 0$, we obtain that $-A/z - B \in \mathcal{I}_{++}$. Thus $(A/z + B)^{-1} \in \mathcal{I}_{++}$ by Lemma D.5 (i), and so the result follows by Lemma D.5 (ii). If $A \succeq 0$ but A is not positive definite, we apply a limiting argument as before to deduce that $\Im \log(A + Bz) \succeq 0$. \square

We will also need the following result:

Lemma D.7 ([24]). *If $0 \prec \Im C \prec \pi$, then $\operatorname{tr} \exp C \in \mathbb{C}_{++}$.*

We omit the proof, which proceeds by first showing that the spectrum of C is contained in the strip $\{z \in \mathbb{C} : 0 < \Im z < \pi\}$, and then using the spectral mapping theorem to deduce that the spectrum of $\exp C$ lies in \mathbb{C}_{++} .

Lemma D.8. *The function g is Herglotz.*

Proof. Take any $z \in \mathbb{C}_{++}$. By Lemma D.6, we have that

$$\begin{aligned} 0 &\preceq \Im \log(C_1 z + K_1) \prec \arg z \\ \text{and} \quad 0 &\prec \Im \log(C_2 + (-1/z)K_2) \preceq \arg(-1/z). \end{aligned}$$

Since $\arg(-1/z) = \pi - \arg z$, we obtain that

$$0 \prec \Im(L + \log(C_1 z + K_1) + \log(C_2 + (-1/z)K_2)) \prec \pi.$$

Thus by Lemma D.7, $g(z) \in \mathbb{C}_{++}$. Hence g is indeed Herglotz. \square

Applying Lemma D.3, and observing the proof of Lemma D.4 below, Theorem 3.12 has been proved.

Proof of Lemma D.4. Firstly, if $z \notin \mathbb{R}$, then either $C_1 + zK_1 \in \mathcal{I}_{++}$, or $-(C_1 + zK_1) \in \mathcal{I}_{++}$. Thus, as observed by Epstein, $\log(C_1 + zK_1)$ is defined; indeed, we already proved more in Lemma D.6. The same is true for $\log(C_2 - zK_2)$.

Now suppose $z \in (-\epsilon, \epsilon)$. Then

$$C_1 + zK_1 \succeq C_1 - \epsilon \|K_1\| \succeq C_1 - \frac{1}{2} \lambda_{\min}(C_1) \succ 0.$$

Similarly $C_2 - zK_2 \succ 0$. \square

E Weaker Proof of Theorem 3.12

In this appendix we prove Theorem 3.12, under the additional hypothesis that C_i & K_i commute. This suffices to prove Lemma 3.9. The argument builds on Lieb's original proof [32] of Theorem 3.11.

Theorem E.1. *Let $L \in \mathbb{S}^n$, $C_1, C_2 \in \mathbb{S}_{++}^n$ and $K_1, K_2 \in \mathbb{S}_+^n$ be such that C_1 & K_1 commute, and that C_2 & K_2 commute. Then*

$$f(z) := \operatorname{tr} \exp \left(L + \log(C_1 + zK_1) + \log(C_2 - zK_2) \right) \quad (11)$$

is concave in a neighborhood of 0.

First we need some preliminary definitions. For $x, y \geq 0$, define the logarithmic mean and binomial mean as follows:

$$\begin{aligned} \operatorname{LM}(x, y) &= \begin{cases} \frac{x-y}{\log x - \log y} & (x \neq y) \\ x & (\text{otherwise}) \end{cases} \\ \operatorname{BM}(x, y) &= \left(\frac{x+y}{2} + \sqrt{xy} \right) / 2 = \left(\frac{\sqrt{x} + \sqrt{y}}{2} \right)^2. \end{aligned}$$

Theorem E.2 (Carlson [17], Bhatia [7]). *For $x, y \geq 0$,*

$$\sqrt{xy} \leq \operatorname{LM}(x, y) \leq \operatorname{BM}(x, y) \leq (x+y)/2.$$

For any $X \in \mathbb{S}_{++}^n$, define the operators $T_X, R_X : \mathbb{S}^n \rightarrow \mathbb{S}^n$ by

$$\begin{aligned} T_X(Y) &:= \int_0^\infty (X+tI)^{-1} Y (X+tI)^{-1} dt \\ R_X(Y) &:= 2 \int_0^\infty (X+tI)^{-1} Y (X+tI)^{-1} Y (X+tI)^{-1} dt. \end{aligned}$$

Claim E.3. *Let $X \in \mathbb{S}_{++}^n$ and $Y \in \mathbb{S}^n$.*

- **(P1):** *If X and Y commute then $T_X(Y) = YX^{-1}$ and $R_X(Y) = Y^2X^{-2}$.*
- **(P2):** *The inverse of T_X is the operator T_X^{-1} where $T_X^{-1}(Y) = \int_0^1 X^t Y X^{1-t} dt$.*
- **(P3):** *In a basis in which X is diagonal, we have $(T_X^{-1}(Y))_{i,j} = Y_{i,j} \cdot \operatorname{LM}(X_{i,i}, X_{j,j})$.*
- **(P4):** *T_X is a positive map, i.e., $T_X(Y) \in \mathbb{S}_+^n$ whenever $Y \in \mathbb{S}_+^n$.*

Proof. See Lieb [32] p. 277, and Ohya and Petz [36] Eq. (3.7) and p. 49. □

Claim E.4. *For any $C \in \mathbb{S}_{++}^n$, $K \in \mathbb{S}^n$ and $x \in \mathbb{R}$,*

$$\log(C + xK) = \log C + xT_C(K) - \frac{1}{2}x^2R_C(K) + O(x^3).$$

Proof. See Lieb [32] equations (3.6) and (3.9), and Ohya and Petz [36, p. 53]. □

Claim E.5. Let $L \in \mathbb{S}^n$, $C_1, C_2 \in \mathbb{S}_{++}^n$ and $K_1, K_2 \in \mathbb{S}^n$. Define $M = \exp(L + \log C_1 + \log C_2)$. Then

$$\begin{aligned} & \exp\left(L + \log(C_1 + zK_1) + \log(C_2 - zK_2)\right) \\ &= M + z \int_0^1 M^{1-s} \left(T_{C_1}(K_1) - T_{C_2}(K_2)\right) M^s ds \\ &+ z^2 \left(-\frac{1}{2} \int_0^1 M^{1-s} \left(R_{C_1}(K_1) + R_{C_2}(K_2)\right) M^s ds \right. \\ &\left. + \int_0^1 \int_0^s M^{1-s} \left(T_{C_1}(K_1) - T_{C_2}(K_2)\right) M^{s-u} \left(T_{C_1}(K_1) - T_{C_2}(K_2)\right) M^u du ds \right) + O(z^3). \end{aligned}$$

Proof. Similar to Ohya and Petz [36, p. 53]. \square

Proof of Theorem E.1. The theorem is equivalent to $0 \leq \frac{d^2 f}{dz^2} \Big|_{z=0}$ (assuming that this derivative exists). From Claim E.5 we have

$$\begin{aligned} & \frac{d^2 f}{dz^2} \Big|_{z=0} \tag{12} \\ &= -\operatorname{tr} M(R_{C_1}(K_1) + R_{C_2}(K_2)) + \operatorname{tr} \int_0^1 (T_{C_1}(K_1) - T_{C_2}(K_2)) M^y (T_{C_1}(K_1) - T_{C_2}(K_2)) M^{1-y} dy \\ &= -\operatorname{tr} M(R_{C_1}(K_1) + R_{C_2}(K_2)) + \operatorname{tr} (T_{C_1}(K_1) - T_{C_2}(K_2)) T_M^{-1} (T_{C_1}(K_1) - T_{C_2}(K_2)). \end{aligned}$$

From **(P1)** and the assumption that C_i and K_i commute we have $R_{C_i}(K_i) = T_{C_i}(K_i)^2$. So the assertion of the theorem is equivalent to

$$\operatorname{tr} M T_{C_1}(K_1)^2 + \operatorname{tr} M T_{C_2}(K_2)^2 \geq \operatorname{tr} \left(T_{C_1}(K_1) - T_{C_2}(K_2) \right) T_M^{-1} \left(T_{C_1}(K_1) - T_{C_2}(K_2) \right). \tag{13}$$

We will prove the more general statement that for all $M \in \mathbb{S}_{++}^n$ and $X, Y \in \mathbb{S}_+^n$,

$$\operatorname{tr} M X^2 + \operatorname{tr} M Y^2 \geq \operatorname{tr} (X - Y) T_M^{-1} (X - Y). \tag{14}$$

This implies (13) by our assumption that $K_1, K_2 \in \mathbb{S}_+^n$ and **(P4)**.

The preceding discussion is basis-independent. It is now convenient to fix a basis in which M is diagonal and to view M, X and Y as matrices in that basis. Let us denote the diagonal entries of M by $\lambda_i = M_{i,i}$; these are positive since we assume $M \in \mathbb{S}_{++}^n$. By **(P3)**, the right-hand side of (14) is

$$\begin{aligned} \operatorname{tr} (X - Y) T_M^{-1} (X - Y) &= \sum_{i,j} \operatorname{LM}(\lambda_i, \lambda_j) \cdot (X_{i,j} - Y_{i,j})^2 \\ &\leq \sum_{i,j} \operatorname{BM}(\lambda_i, \lambda_j) \cdot (X_{i,j} - Y_{i,j})^2, \end{aligned} \tag{15}$$

by Theorem E.2. We may rewrite the right-hand side as

$$\begin{aligned} & \sum_{i,j} \left(\frac{\lambda_i}{4} + \frac{\lambda_j}{4} + \frac{\sqrt{\lambda_i \lambda_j}}{2} \right) ((X_{i,j})^2 + (Y_{i,j})^2 - 2X_{i,j}Y_{i,j}) \\ &= \frac{\operatorname{tr} M X^2}{2} + \frac{\operatorname{tr} M^{1/2} X M^{1/2} X}{2} + \frac{\operatorname{tr} M Y^2}{2} + \frac{\operatorname{tr} M^{1/2} Y M^{1/2} Y}{2} - \operatorname{tr} M X Y - \operatorname{tr} M^{1/2} X M^{1/2} Y \end{aligned} \tag{16}$$

by repeatedly using the observation

$$\sum_{i,j} D_{i,i} P_{i,j} Q_{i,j} E_{j,j} = \text{tr} DPEQ = \text{tr} EPDQ$$

for all $D, E \in \mathbb{D}^n$, $P, Q \in \mathbb{S}^n$.

Thus, combining (14), (15) and (16), it suffices to prove

$$\text{tr} MX^2 - \text{tr} M^{1/2} X M^{1/2} X + \text{tr} MY^2 - \text{tr} M^{1/2} Y M^{1/2} Y \geq -2 \text{tr} MXY - 2 \text{tr} M^{1/2} X M^{1/2} Y$$

for every $M, X, Y \in \mathbb{S}_+^n$.

Since that inequality is invariant under choice of orthonormal basis, and since $\text{tr} M^{1/2} X M^{1/2} Y \geq 0$, it suffices to prove

$$\text{tr} XD^2X - \text{tr} XDXD + \text{tr} YD^2Y - \text{tr} YD YD \geq -2 \text{tr} XD^2Y \quad \forall D \in \mathbb{D}^n, \forall X, Y \in \mathbb{S}_+^n. \quad (17)$$

Denote the diagonal entries of D by $d_i = D_{i,i}$. Then

$$\text{tr} XD^2X - \text{tr} XDXD = \frac{1}{2} \sum_{i,j} X_{i,j}^2 (d_i^2 + d_j^2) - \sum_{i,j} X_{i,j}^2 d_i d_j = \frac{1}{2} \sum_{i,j} X_{i,j}^2 (d_i - d_j)^2.$$

So the left-hand side of (17) equals

$$\sum_{i,j} \frac{X_{i,j}^2 + Y_{i,j}^2}{2} (d_i - d_j)^2 \geq \sum_{i,j} |X_{i,j} Y_{i,j}| \cdot (d_i - d_j)^2,$$

by the arithmetic-mean geometric-mean (AM-GM) inequality. The right-hand side of (17) is

$$-2 \text{tr}(XD^2Y) = -\text{tr}(XD^2Y) - \text{tr}(D^2XY) = -\sum_{i,j} X_{i,j} Y_{i,j} (d_i^2 + d_j^2).$$

So, to prove (17), it suffices to prove that

$$\sum_{i,j} |X_{i,j} Y_{i,j}| \cdot (d_i - d_j)^2 \geq -\sum_{i,j} X_{i,j} Y_{i,j} (d_i^2 + d_j^2). \quad (18)$$

We will prove the more general inequality

$$\sum_{i,j} |Z_{i,j}| \cdot (d_i - d_j)^2 \geq -\sum_{i,j} Z_{i,j} (d_i^2 + d_j^2) \quad \forall d \in \mathbb{R}^n, \forall Z \in \mathbb{S}_+^n. \quad (19)$$

This implies (18) by letting $Z = X \circ Y$ (the Hadamard product of X and Y), which is positive semidefinite by the Schur product theorem [6, p. 23]. Rearranging, (19) becomes

$$\frac{1}{2} \sum_{i,j} (|Z_{i,j}| + Z_{i,j}) (d_i^2 + d_j^2) \geq \sum_{i,j} |Z_{i,j}| d_i d_j. \quad (20)$$

Since $|Z_{i,j}| + Z_{i,j} \geq 0$, the AM-GM inequality implies that the left-hand side is at least

$$\sum_{i,j} (|Z_{i,j}| + Z_{i,j}) d_i d_j = \sum_{i,j} |Z_{i,j}| d_i d_j + d^\top Z d.$$

Since $Z \in \mathbb{S}_+^n$, this implies (20). \square

F Connections to the Kadison-Singer Problem

The Kadison-Singer problem, which dates back to 1959, is an important, and until very recently unsolved, question in operator theory. The importance of this question has become increasingly apparent in recent years as it is now known to be equivalent, or closely related, to numerous conjectures in disparate areas of mathematics [18]. In a very recent breakthrough, Marcus, Spielman and Srivastava [34] positively resolved the Kadison-Singer problem. More precisely, they proved the following strong form of Weaver's conjecture [55, Conjecture KS₂ and Theorem 2]:

Theorem F.1 ([34]). *Let $\epsilon > 0$, and $u_1, \dots, u_m \in \mathbb{C}^n$ such that $\|u_i\| \leq \epsilon$ for all i , and $\sum_i u_i u_i^\top = I$. Then there exists a partition of $[m]$ into S_1, S_2 such that for each $j \in \{1, 2\}$,*

$$\sum_{i \in S_j} u_i u_i^\top \leq \frac{1}{2}(1 + \sqrt{2}\epsilon)^2. \quad (21)$$

It is well-known that, given a strong discrepancy result such as (21), an iterative argument yields a sparse object that gives a good approximation. See, e.g., Rudelson [43]. For the sake of completeness, we include here a detailed argument that Theorem F.1 implies the existence of $O(1/\kappa)$ -spectrally-thin trees.

First, the following corollary of Theorem F.1 will be convenient for induction purposes.

Corollary F.2. *There exists a constant $C \geq 1$ such that the following is true. Let $v_1, \dots, v_m \in \mathbb{R}^n$ be such that $\alpha I \preceq \sum_i v_i v_i^\top \preceq \beta I$ and $\|v_i\|^2 = \delta := n/m$ for all i . Suppose that $\alpha \in [1/2, 1]$ and $\beta \in [1, 2]$. Then there exists $S \subseteq [m]$ satisfying*

$$(\alpha - C\sqrt{\delta})I \preceq 2 \sum_{i \in S} v_i v_i^\top \preceq (\beta + C\sqrt{\delta})I. \quad (22)$$

Proof. Let $\alpha, \beta, \delta, v_1, \dots, v_m$ be as in the statement of Corollary F.2. Note that $\delta \leq 1$, since $m \geq n$. Letting $M = \sum_i v_i v_i^\top$, we see that $\|M^{-1}\| \leq \alpha^{-1}$. Define $u_i = M^{-1/2}v_i$. Then

$$\sum_i u_i u_i^\top = M^{-1/2} \left(\sum_i v_i v_i^\top \right) M^{-1/2} = I$$

and

$$\|u_i\|^2 \leq \|M^{-1}\| \|v_i\|^2 \leq \alpha^{-1} \delta =: \epsilon.$$

Applying Theorem F.1 of Marcus et al. [34], we deduce (21), and hence (since $\epsilon \leq 2$)

$$2 \sum_{i \in S_j} u_i u_i^\top \preceq 1 + 4\sqrt{2}\epsilon \quad \text{for } j \in \{1, 2\}.$$

Consequently,

$$2 \sum_{i \in S_1} v_i v_i^\top \preceq (1 + 4\sqrt{2}\epsilon)M \preceq (1 + 4\sqrt{2}\epsilon)\beta \preceq \beta + 16\sqrt{\delta}$$

by the hypotheses $\alpha \in [1/2, 1]$ and $\beta \in [1, 2]$.

Observing that

$$2 \sum_{i \in S_1} u_i u_i^\top = 2I - 2 \sum_{i \in S_2} u_i u_i^\top \succeq 1 - 4\sqrt{2}\epsilon,$$

we similarly obtain

$$2 \sum_{i \in S_1} v_i v_i^\top \succeq (1 - 4\sqrt{2}\epsilon)M \succeq (1 - 4\sqrt{2}\epsilon)\alpha \succeq \alpha - 4\sqrt{2}\delta.$$

Thus taking $S = S_1$, we see that (22) holds with $C = 16$. \square

We may now prove Theorem 4.3 by an application of Corollary F.2. By an argument similar to the proof of Theorem 4.10, this implies the existence of $O(1/\kappa)$ -spectrally-thin trees.

Claim F.3. *Corollary F.2 implies Theorem 4.3.*

Proof. As in Theorem 4.3, let $w_1, \dots, w_m \in \mathbb{R}^n$ satisfy $\|w_i\| = 1$ for all i . Let p_1, \dots, p_m be a probability distribution on these vectors such that the covariance matrix is $\sum_i p_i w_i w_i^\top = I/n$.

Without loss of generality, we may assume $p_i = 1/m \forall i$. To see this, suppose that p_1, \dots, p_m are rational numbers of the form q_i/M where q_1, \dots, q_m, M are nonnegative integers. Then we may replace each w_i with q_i copies of itself. The uniform distribution on the resulting vectors still has covariance matrix I/n . Proving Theorem 4.3 for the resulting vectors establishes the theorem for the original vectors under distribution p . If p_1, \dots, p_m are irrationals, we may approximate them by rationals while introducing vanishing error.

Define $v_i = \sqrt{n/m} \cdot w_i$, so that $\|v_i\|^2 = n/m =: \delta_0$ for all i . We will iteratively construct sets $S_t \subseteq [m]$, with $S_0 = [m]$. Let C be as in Corollary F.2. Define $\alpha_0 = \beta_0 = m$, and then inductively

$$\begin{aligned}\alpha_{t+1} &= \alpha_t - C(2^t n)^{1/2} (\beta_t)^{1/2} \\ \beta_{t+1} &= \beta_t + C(2^t n)^{1/2} (\beta_t)^{1/2}.\end{aligned}$$

Let $\epsilon \in (0, 1]$ be a small constant to be chosen in a moment, and let

$$T = \max \left\{ t : C \sum_{j=0}^{t-1} (2^j n/m)^{1/2} \leq \epsilon/2 \right\}.$$

This choice of T is motivated by the following:

Claim F.4. *For all $t \leq T$, $\beta_t \leq m(1 + \epsilon)$ and $\alpha_t \geq m(1 - \epsilon)$.*

Proof. For $0 \leq t < T$,

$$\beta_{t+1} = \beta_t (1 + C(2^t n/\beta_t)^{1/2}) \leq \beta_t (1 + C(2^t n/m)^{1/2}).$$

So

$$\begin{aligned}\beta_t &\leq m \prod_{j=0}^{t-1} (1 + C(2^j n/m)^{1/2}) \\ &\leq m \exp \left(C \sum_{j=0}^{t-1} (2^j n/m)^{1/2} \right) \\ &\leq m \exp(\epsilon/2) \\ &\leq m(1 + \epsilon).\end{aligned}$$

Note that

$$\alpha_0 - \alpha_t = \beta_t - \beta_0$$

and so since $\beta_t \leq m(1 + \epsilon)$, $\alpha_t \geq m(1 - \epsilon)$. □

Note that since $\sum_{j=0}^{T-1} (2^j n/m)^c = \Theta((2^T n/m)^{1/2})$, we have that

$$2^T = \Theta \left(\frac{m}{n} \epsilon^2 \right). \tag{23}$$

So we may choose $\epsilon \in (0, 1/3]$ to be a constant sufficiently small so that

$$(1 - \epsilon)^{-1} 2^T (n/m) \leq d. \quad (24)$$

Our first goal will be to show inductively that for all $t \leq T$, there exists a set $S_t \subseteq [m]$ so that

$$\alpha_t \preceq m 2^t \sum_{i \in S_t} v_i v_i^\top \preceq \beta_t. \quad (25)$$

Note that this is true for $t = 0$ by assumption.

It will be convenient to define $\gamma_t = 2^t |S_t|$. Suppose (25) holds for some particular $t < T$. Define

$$v_i^{(t)} = v_i \cdot \sqrt{m 2^t / \gamma_t},$$

so that $\|v_i^{(t)}\|^2 = \frac{n}{|S_t|} =: \delta_t$ for all t . Then just by scaling,

$$\alpha_t / \gamma_t \preceq \sum_{i \in S_t} v_i^{(t)} (v_i^{(t)})^\top \preceq \beta_t / \gamma_t.$$

Taking a trace yields $n \alpha_t / \gamma_t \leq n \leq n \beta_t / \gamma_t$, i.e.,

$$\alpha_t \leq \gamma_t \leq \beta_t. \quad (26)$$

By (26), Claim F.4 and (24), we have

$$\begin{aligned} 1/2 &\leq \frac{1 - \epsilon}{1 + \epsilon} \leq \frac{\alpha_t}{\beta_t} \leq \frac{\alpha_t}{\gamma_t} \leq 1 \\ 1 &\leq \frac{\beta_t}{\gamma_t} \leq \frac{\beta_t}{\alpha_t} \leq \frac{1 + \epsilon}{1 - \epsilon} \leq 2 \\ \delta_t &= 2^t n / \gamma_t \leq 2^T n / \alpha_t \leq 2^T n / (m(1 - \epsilon)) \leq d. \end{aligned}$$

Now apply Corollary F.2 with S_t instead of $[m]$, $v_i^{(t)}$ instead of v_i , α_t / γ_t instead of α , β_t / γ_t instead of β , and δ_t instead of δ . The hypotheses of Corollary F.2 are satisfied, so it follows that there is a set $S_{t+1} \subseteq S_t$ with

$$\alpha_t / \gamma_t - C \delta_t^{1/2} \preceq 2 \sum_{i \in S_{t+1}} v_i^{(t)} (v_i^{(t)})^\top \preceq \beta_t / \gamma_t + C \delta_t^{1/2}.$$

Rewriting in terms of the original v_i 's, we obtain

$$\alpha_t - C \gamma_t \delta_t^{1/2} \preceq 2^{t+1} m \sum_{i \in S_{t+1}} v_i v_i^\top \preceq \beta_t + C \gamma_t \delta_t^{1/2}.$$

Now

$$\gamma_t \delta_t^{1/2} = \gamma_t (n / |S_t|)^{1/2} = (2^t n)^{1/2} (\gamma_t)^{1/2} \leq (2^t n)^{1/2} (\beta_t)^{1/2}.$$

Hence

$$\alpha_{t+1} \preceq 2^{t+1} m \sum_{i \in S_{t+1}} v_i v_i^\top \preceq \beta_{t+1},$$

and the inductive step is achieved.

From Claim F.4 and (25) for $t = T$, we deduce that

$$m(1 - \epsilon) \preceq 2^T n \sum_{i \in S_T} w_i w_i^\top \preceq m(1 + \epsilon);$$

hence by (23) and since ϵ is a constant,

$$\Theta((1 - \epsilon)\epsilon^{-2}) \preceq \sum_{i \in S_T} w_i w_i^\top \preceq \Theta((1 + \epsilon)\epsilon^{-2}) = \Theta(1).$$

The left inequality implies that $\{w_i : i \in S_T\}$ spans \mathbb{R}^n . To conclude, select an arbitrary basis $B \subseteq S_T$; then

$$\sum_{i \in B} w_i w_i^\top \preceq \sum_{i \in S_T} w_i w_i^\top \preceq \Theta(1),$$

and so Theorem 4.3 holds. □