

ℓ_p -Spread and Restricted Isometry Properties of Sparse Random Matrices

Venkatesan Guruswami*
venkatg@berkeley.edu
UC Berkeley

Peter Manohar†
pmanohar@cs.cmu.edu
Carnegie Mellon University

Jonathan Mosheiff‡
jmosheiff@cs.cmu.edu
Carnegie Mellon University

Abstract

Random subspaces X of \mathbb{R}^n of dimension proportional to n are, with high probability, well-spread with respect to the ℓ_2 -norm. Namely, every nonzero $x \in X$ is “robustly non-sparse” in the following sense: x is $\varepsilon \|x\|_2$ -far in ℓ_2 -distance from all δn -sparse vectors, for positive constants ε, δ bounded away from 0. This “ ℓ_2 -spread” property is the natural counterpart, for subspaces over the reals, of the minimum distance of linear codes over finite fields, and corresponds to X being a Euclidean section of the ℓ_1 unit ball. Explicit ℓ_2 -spread subspaces of dimension $\Omega(n)$, however, are unknown, and the best known explicit constructions (which achieve weaker spread properties), are analogs of *low density parity check* (LDPC) codes over the reals, i.e., they are kernels of certain *sparse* matrices.

Motivated by this, we study the spread properties of the kernels of *sparse random matrices*. We prove that with high probability such subspaces contain vectors x that are $o(1) \cdot \|x\|_2$ -close to $o(n)$ -sparse with respect to the ℓ_2 -norm, and in particular are *not* ℓ_2 -spread. This is strikingly different from the case of random LDPC codes, whose distance is asymptotically almost as good as that of (dense) random linear codes.

On the other hand, for $p < 2$ we prove that such subspaces are ℓ_p -spread with high probability. The spread property of sparse random matrices thus exhibits a threshold behavior at $p = 2$. Our proof for $p < 2$ moreover shows that a random sparse matrix has the stronger restricted isometry property (RIP) with respect to the ℓ_p norm, and in fact this follows solely from the unique expansion of a random biregular graph, yielding a somewhat unexpected generalization of a similar result for the ℓ_1 norm [BGI⁺08]. Instantiating this with suitable explicit expanders, we obtain the first explicit constructions of ℓ_p -RIP matrices for $1 \leq p < p_0$, where $1 < p_0 < 2$ is an absolute constant.

Keywords: Spread Subspaces, Euclidean Sections, Restricted Isometry Property, Sparse Matrices

*Supported in part by NSF grants CCF-1908125 and CCF-2210823, and a Simons Investigator Award.

†Supported in part by an ARCS Scholarship, NSF Graduate Research Fellowship (under grant numbers DGE1745016 and DGE2140739), and NSF CCF-1814603.

‡Supported in part by NSF CCF-1814603.

Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

Contents

1	Introduction	1
1.1	Kernels of sparse matrices	2
1.2	Our results	3
2	Proof overview	8
2.1	Theorem 1: $\ker(A)$ is not ℓ_2 -spread	8
2.2	Theorem 4: ℓ_p -RIP for $p < 2$ from vertex expansion	10
2.3	Theorem 2: bounds on the singular values of A	12
2.4	Organization	14
3	Preliminaries	14
3.1	Basic notation	14
3.2	Biregular bipartite graphs and sparse matrices	15
3.3	ℓ_p norms, compressibility, spread subspaces, and distortion	16
4	Properties of random biregular graphs	17
4.1	A negative association lemma	17
4.2	Proof of Proposition 4.1: scarcity of cycles	19
4.3	Proof of Proposition 4.2: unique expansion	21
5	Limitations of ℓ_2-spread	22
5.1	Proof of Lemma 5.1: constructing x	24
5.2	Proof of Lemma 5.2: rounding x to $y \in \ker(A)$	25
6	Positive results for ℓ_p-RIP and ℓ_p-spread for $p \in [1, 2)$	25
6.1	Proof of Corollary 5: ℓ_p -RIP and ℓ_p -spread of sparse random matrices	26
6.2	Proof of Corollary 6: explicit construction of ℓ_p -RIP matrices	27
6.3	Proof of Lemma 6.1: ℓ_p -RIP from unique expansion	28
7	Singular values of random biregular matrices	31
7.1	Counting hikes when $s \leq \text{polylog}(n)$	35
7.2	Counting hikes when $s = \omega(\text{polylog}(n))$	37
A	A weak positive bound on ℓ_2-spread	41
B	General relations between spread, distortion and restricted isometry properties	43
B.1	ℓ_p -spread implies ℓ_q -spread for $q < p$	43
B.2	ℓ_p -RIP implies ℓ_p -spread	44
B.3	Equivalence between ℓ_p -spread and (ℓ_q, ℓ_p) -distortion	45

1 Introduction

Classical results in asymptotic geometric analysis on the Gelfand/Kolmogorov widths of ℓ_2 balls [FLM77, Kas77, GG84] show that random subspaces X of \mathbb{R}^n of dimension proportional to n (say, defined as the kernel of random $n/2 \times n$ matrices with i.i.d. Gaussian or ± 1 entries) are *good Euclidean sections* of ℓ_1^n : namely, $\|x\|_1 \geq \Omega(\sqrt{n}) \|x\|_2$ for every $x \in X$. An elementary proof of this fact also follows from the Johnson-Lindenstrauss (JL) property of random matrices, its connection to the restricted isometry property (RIP) and compressed sensing, and their relationship to the Euclidean sections property [BDDW07].

The condition $\|x\|_1 \geq \Omega(\sqrt{n}) \|x\|_2$ can equivalently¹ be expressed as a “well-spreadness” criterion satisfied by every nonzero vector $x \in X$: the largest δn entries of x have at most $1 - \varepsilon$ of its ℓ_2 mass, for some positive constants δ, ε bounded away from 0 as $n \rightarrow \infty$. Equivalently, this means that all nonzero vectors $x \in X$ are *incompressible*—there is no sparse vector that approximates x well in ℓ_2 norm (in other words, $\|x - y\|_2 \geq \varepsilon \|x\|_2$ for all δn -sparse vectors y). This can be naturally viewed as a robust analog, for subspaces of \mathbb{R}^n , of the distance property of linear error-correcting codes.

The above well-spreadness criterion can naturally be imposed with respect to any ℓ_p metric: a subspace X is said to be ℓ_p -spread if every nonzero vector $x \in X$ is $\varepsilon \|x\|_p$ -far in ℓ_p -distance from all δn -sparse vectors. The ℓ_p -spread property is a more stringent requirement for larger p (see Proposition 3.7). For $p > 2$, the optimal asymptotic dimension of ℓ_p -spread subspaces is at most $O_p(n^{2/p})$ and thus $o(n)$ [Glu83]. In this work, we therefore focus on $p \in [1, 2]$ where it is possible to have ℓ_p -spread subspaces of dimension proportional to n .

For a subspace X of \mathbb{R}^n , define its ℓ_p -distortion $\Delta_p(X)$ to be the following quantity:

$$\Delta_p(X) := \sup_{x \in X \setminus \{0\}} \frac{n^{1-\frac{1}{p}} \|x\|_p}{\|x\|_1}.$$

Note that $1 \leq \Delta_p(x) \leq n^{1-1/p}$. Good ℓ_p -spread of X can be captured by the condition that $\Delta_p(X)$ is bounded by a fixed constant independent of n ; this generalizes the aforementioned equivalence¹ of ℓ_2 -spread and the Euclidean section property. The term *distortion* is used because the natural inclusion of X in \mathbb{R}^n induces a bi-Lipschitz embedding of X , taken with the ℓ_p norm, into ℓ_1^n , with distortion $\Delta_p(X)$. The distortion/spread property of subspaces with respect to different ℓ_p norms has been extensively studied, owing to its connections to width properties in convex geometry [Glu83, KT07], embeddings between metric spaces [Ind06], compressed sensing [Don06, CRT06, KT07], error-correction over the reals [CT05, GLW08], and the restricted isometry (RIP) and dimensionality-reduction/Johnson-Lindenstrauss (JL) properties [KT07, BDDW07, AGR15].

Despite a lot of interest and the abundance of probabilistic constructions, an outstanding question is to construct an *explicit* subspace $X \subseteq \mathbb{R}^n$ of dimension $\Omega(n)$ that is ℓ_2 -spread, or equivalently has $\Delta_2(X) \leq O(1)$. By explicit, we mean deterministically constructing a basis for the subspace (or its dual) in $\text{poly}(n)$ time.² This is a counterpart, for subspaces of \mathbb{R}^n , of the problem of constructing asymptotically good binary linear codes $C \subseteq \{0, 1\}^n$: namely, codes whose dimension and minimum distance are both proportional to n . In addition to being a natural and basic challenge in pseudorandomness, explicit constructions are also valuable in applications of spread subspaces

¹See Proposition 3.11.

²Explicit constructions of ℓ_p -spread spaces of dimension $\Omega(n)$ are given in [BGI⁺08] ($p = 1$) and [Kar11] ($1 \leq p < 2$).

such as compressed sensing, as they provide a *guarantee* that the matrix will have the stipulated properties. This is particularly important since there are no known methods to efficiently certify the ℓ_2 -spread (or even ℓ_p -spread) of random subspaces.

1.1 Kernels of sparse matrices

In the case of $p = 2$, the best known explicit constructions of subspaces $X \subseteq \mathbb{R}^n$ with $\dim(X) \geq \Omega(n)$, in terms of their distortion $\Delta_2(X)$, are due to [GLR10]. They give a construction analogous to Tanner codes from coding theory [Tan81], combining appropriately chosen unbalanced bipartite expanders and local subspaces, to produce X with $\dim(X) \geq n - o(n)$ and $\Delta_2(X) \leq (\log n)^{O(\log \log \log n)}$ (so almost poly-logarithmic).³ A simpler construction analogous to Sipser-Spielman codes [SS96], using s -regular spectral expanders and local well-spread subspaces of \mathbb{R}^s , was given in [GLW08] and achieves⁴ $\Delta_2(X) \leq n^{O(1/\log s)}$. An alternate probabilistic construction achieving similar parameters to [GLW08] based on tensor products was given in [IS10]. The approach of [IS10] can further achieve distortion approaching 1 at the expense of making $\dim(X)$ smaller, but still $\Omega(n)$.

One notable attribute of the constructions above is that the subspace X can be expressed as the kernel of a matrix that is *sparse*. For instance, the construction of [GLW08] picks a matrix where each row is s -sparse with ± 1 entries (that are chosen randomly for a probabilistic construction), and the construction in [IS10] defines the subspace $X \subseteq \mathbb{R}^n$ as the k -fold tensor product of another subspace, and so X can be defined as the kernel of an $n^{1/k}$ -sparse matrix. Moreover, known explicit constructions of ℓ_p -spread subspaces for $1 \leq p < 2$ [BGI⁺08, Kar11] are also kernels of sparse matrices.

The sparsity of these constructions is inherited from the “underlying constructions” for codes; the constructions of [GLR10, GLW08, IS10] come from “lifting” constructions of linear codes (namely, Tanner codes [Tan81], Sipser-Spielman codes [SS96], and tensor product codes, respectively) to this setting, and these constructions (for linear codes) are known to give good low density parity check (LDPC) codes: namely, codes that are the kernels of sparse matrices.

In light of these works, a natural question (and indeed one explicitly posed in [GLW08]), is the following.

Question 1. Does there exist an $m \times n$ matrix A with $n - m \geq \Omega(n)$ whose rows are s -sparse for $s \leq O(1)$ (or even $s \leq \text{polylog}(n)$) such that $\Delta_2(\ker(A)) \leq O(1)$?

The approaches of [GLW08, IS10] show that one can achieve $\Delta_2(\ker(A)) \leq \exp(O(1/\delta))$ when $s = n^\delta$. A positive answer to Question 1, even via random matrices, would likely yield good progress towards explicit constructions, as $O(1)$ -sparse matrices are likely easier to derandomize than dense random ones. A negative answer to Question 1 would likely rule out explicit constructions based on the current state-of-the-art approaches of [GLR10, GLW08, IS10].

In addition to exploring the potential of the approaches behind the current best constructions, sparsity is desirable from a computational efficiency standpoint. Sparse matrices lead to faster algorithms, for example when used as measurement matrices in compressed sensing or to compute a sparse JL transform for dimensionality-reduction.

³For sublinear dimension, an explicit construction of $X \subseteq \mathbb{R}^n$ with distortion $\Delta_2(X) \leq 1 + o(1)$ and $\dim(X) \geq n/2^{O((\log \log n)^2)}$ was given in [Ind07].

⁴This construction is not explicit except for very small s , as the local subspace of \mathbb{R}^s is either constructed by brute force or drawn at random.

Motivated by these considerations, we study the ℓ_2 -spread, and, more generally, ℓ_p -spread ($1 \leq p \leq 2$) of subspaces defined as the kernel of *sparse random matrices*. Such subspaces are the continuous analogues of random low density parity check (LDPC) codes. Random LDPC codes have been studied in coding theory since Gallager’s seminal work [Gal63], with a renaissance since the mid 1990s [RU08] due to their fast iterative decoding algorithms and performance close to capacity.

Random LDPC codes are known to achieve rate vs. distance trade-offs approaching that of random (dense) linear codes [Gal63]. Recently, even the list-decodability, and indeed any “local” property, of random LDPC codes was shown to be similar to that of random linear codes [MRR⁺20]. Given that random subspaces are well-spread and that random LDPC codes achieve similar properties to random (dense) codes, one might naturally expect, by analogy, that the kernels of sparse random matrices are also well-spread.

1.2 Our results

Our results paint a precise picture of the ℓ_p -spread of kernels X of sparse random matrices. Before stating our results, we first define ℓ_p -spread and state the random matrix model that we use.

Definition 1.1 (ℓ_p -spread). Fix $p \in [1, \infty]$, $\varepsilon \in [0, 1]$ and $k \leq n \in \mathbb{N}$. A vector $y \in \mathbb{R}^n$ is k -sparse if $|\text{supp}(y)| \leq k$. A vector $x \in \mathbb{R}^n \setminus \{0^n\}$ is said to be (k, ε) - ℓ_p -compressible if there exists a k -sparse $y \in \mathbb{R}^n$ such that $\|x - y\|_p \leq \varepsilon \|x\|_p$. Otherwise, we say that x is (k, ε) - ℓ_p -spread.

A subspace $X \subseteq \mathbb{R}^n$ is (k, ε) - ℓ_p -spread if every $x \in X \setminus \{0^n\}$ is (k, ε) - ℓ_p -spread.

The random matrix model. A matrix $A \in \{0, 1, -1\}^{m \times n}$ is said to be (s, t) -biregular if every row and column of A has exactly s and t nonzero entries, respectively. Let $\mathcal{M}_{m,n,s,t}$ denote the set of all (s, t) -biregular matrices in $\{0, 1, -1\}^{m \times n}$.

All of our theorems for random matrices will be for a matrix A drawn uniformly at random from $\mathcal{M}_{m,n,s,t}$, where $\alpha = \frac{m}{n} = \frac{t}{s} \in (0, 1)$ is a fixed constant and $n \rightarrow \infty$; for this exposition, we will use A to denote a random matrix from $\mathcal{M}_{m,n,s,t}$, and B to denote an arbitrary matrix in $\{0, 1, -1\}^{m \times n}$. We additionally assume that $s := s(n) \leq n^c$ for some absolute constant $0 < c < 1$, and $t = \alpha s \geq 3$. An event \mathcal{E} is said to hold *with high probability* if $\lim_{n \rightarrow \infty} \Pr[\mathcal{E}] = 1$. All asymptotic notation refers to the regime of $n \rightarrow \infty$ and constant α . The constants implied by asymptotic notation are universal, unless stated otherwise. The symbols c, c', c_1 and c_2 always stand for positive universal constants, which may differ across different lemma and theorem statements. We use the phrase “in particular” in theorem statements to refer to an implication that follows by either of the generic reductions of Proposition 3.11 (ℓ_p -spread implies ℓ_p -distortion) or Proposition 3.8 (ℓ_p -RIP implies ℓ_p -spread).

1.2.1 Poor ℓ_2 -spread of sparse random matrices

Our first theorem shows that, surprisingly, $\ker(A)$ is, with high probability, *not* ℓ_2 -spread.

Theorem 1 (Poor ℓ_2 -spread of $\ker(A)$). *With high probability over A , there exists an $(m^c, \frac{n^{-\Omega(\log(1/\alpha)/\log s)}}{1-\sqrt{\alpha}})$ - ℓ_2 -compressible vector $x \in \ker(A)$, where $c < 1$ is an absolute constant. In particular,*

$$\Delta_2(\ker(A)) \geq (1 - \sqrt{\alpha}) \cdot n^{\Omega(\log(1/\alpha)/\log s)} .$$

Moreover, there is a poly(n)-time algorithm that, on input A , outputs such an x .

Choosing $s = O(1)$ in Theorem 1 (and letting α be bounded away from 1) implies⁵ that $\Delta_2(\ker(A)) \geq n^{\Omega(1)}$ with high probability, and choosing $s = \text{polylog}(n)$ implies $\Delta_2(\ker(A)) \geq n^{\Omega(\log(1/\alpha)/\log \log n)}$. We always trivially have $\Delta_2(\ker(A)) \leq \sqrt{n}$, so not only does Theorem 1 answer Question 1 in the negative for sparse random matrices, but it also does so in a very strong sense. For instance, when $s = O(1)$, Theorem 1 shows that $\Delta_2(\ker(A))$ is “maximally bad”, up to a constant factor in the exponent.

Another point of interest is the choice $s = n^\delta$ for some fixed δ . This yields the tradeoff of $\Delta_2(\ker(A)) \geq (\frac{1}{\alpha})^{\Omega(\frac{1}{\delta})}$, which precisely matches the tradeoff (in terms of δ) achieved by both [GLW08, IS10]. While our matrix ensemble is “more random” compared to those in [GLW08, IS10], Theorem 1 can nonetheless be interpreted as giving evidence that this $\exp(O(\frac{1}{\delta}))$ tradeoff from [GLW08, IS10] is tight and inherent to sparse constructions.

Our proof of Theorem 1 is *constructive*, in the sense that we give a very simple, efficient algorithm to find such an $x \in \ker(A)$. This moreover shows that for sparse random matrices, one can efficiently *refute* the claim that $\Delta_2(\ker(A)) = O(1)$, as the vector x is a refutation witness. Our algorithm provides an interesting counterpoint to the work of [BBH⁺12], who gave an algorithm based on the sum-of-squares SDP hierarchy to certify that $\Delta_2(\ker(A)) \leq O(1)$ with high probability for *dense* matrices A where $\dim(\ker(A)) \leq O(\sqrt{n})$. In contrast, our algorithm succeeds when $\dim(\ker(A)) = \Omega(n)$ and the matrix A is *sparse*. The two results taken together suggest an interesting relationship between the density and $\dim(\ker(A))$ of matrices A for which we can efficiently certify or refute bounds on $\Delta_2(\ker(A))$.

We also note that, by the well-known duality formula relating Kolmogorov and Gelfand widths (see [KT07] and the references therein), Theorem 1 implies that the row span of A is far from approximating the ℓ_2 -sphere in ℓ_∞ distance. Concretely, with high probability over A there exists $x \in \mathbb{R}^n$ with $\|x\|_2 = 1$ that is $(1 - \sqrt{\alpha}) \cdot n^{\Omega(\log(1/\alpha)/\log s)}/\sqrt{n}$ -far in ℓ_∞ norm from all vectors of the form $A^\top y$, where $y \in \mathbb{R}^m$.

The proof of Theorem 1 requires the following strong bound that we show on the singular values of A .

Theorem 2 (Singular value bound). *With high probability, the set of singular values $\sigma(A)$ of A satisfy*

$$\sigma(A) \subseteq \left[\sqrt{s-1} - (1+o(1)) \cdot \sqrt{t-1}, \sqrt{s-1} + (1+o(1)) \cdot \sqrt{t-1} \right].$$

Moreover, the above bound holds even without our (otherwise global) assumption that $s \leq n^c$ for some absolute constant $c < 1$.

Theorem 2 should not be surprising, especially given the recent works of [BDH18, Bor19, BC19, MOP20a, MOP20b, OW20, Zhu20], and indeed our proof follows the same overall blueprint of these works. Most of these papers, however, only handle the case when the degree of the graph is *constant* as $n \rightarrow \infty$;⁶ this corresponds to the case of $s = O(1)$ in Theorem 2. Theorem 2 thus differs as it allows for $s = \omega(1)$, and indeed we can even take $s = n^c$ for some absolute constant $c < 1$. On the other hand, most of the aforementioned works deal with the case of *unsigned* adjacency matrices, whereas we only prove Theorem 2 for *randomly signed* adjacency matrices. We note that proving

⁵Note that since $\alpha s = t \geq 3$, we must have $\log s \geq \log \frac{1}{\alpha}$.

⁶The exceptions are [MOP20a], which handles $\text{polylog}(n)$ degree, and [Zhu20], which handles n^c degree but does not obtain as sharp bounds.

the analogue of Theorem 2 for unsigned adjacency matrices (where $\sigma(A)$ now denotes the set of singular values, excluding the trivial value of \sqrt{st}) and for all (non-constant) s, t remains open.

The singular value bound in Theorem 2 is challenging to prove because it is so sharp. Indeed, it is not too difficult to show that $\sigma(A) \subseteq [\sqrt{s} - O(\sqrt{t}), \sqrt{s} + O(\sqrt{t})]$ with high probability via black-box applications of known results, e.g., [BV16]. However, this does not suffice for our use in the proof of Theorem 1, as the aforementioned weaker bound would only suffice to prove Theorem 1 provided that $\alpha \leq c$ for some absolute constant c , where c depends on the absolute constant c' hidden in the “ $O(\sqrt{t})$ ”. We need the sharp bound of Theorem 2 in order to allow for α to be an arbitrary constant in $(0, 1)$.

As a counterpart to Theorem 1, we give the following partial converse, which shows that $\ker(A)$ is (k, ε) - ℓ_2 -spread for a weak choice of parameters k and ε .

Theorem 3 (Converse to Theorem 1). *Assume that $t \geq 9$. Then, with high probability over A , the space $\ker(A)$ is $(\Omega(\alpha^2 n / t^4), \alpha^{O(\log n / \log t)})$ - ℓ_2 -spread.*

We note that in Theorem 3, the parameter k is $\Omega(\alpha^2 n / t^4) = m^{\Omega(1)}$ and the parameter ε is $\alpha^{O(\log n / \log s)}$.⁷ Theorem 3 thus shows that the parameters in Theorem 1 are tight up to the universal constants in the exponent. Our proof of Theorem 3 is an adaptation of the proof of [BR17, Lemma 3.4].

1.2.2 ℓ_p -RIP and ℓ_p -spread for $p < 2$

We next focus on the ℓ_p norm for $p < 2$. For $p < 2$, there are known explicit constructions of ℓ_p -spread subspaces [Kar11]. Because of this, for $p < 2$ we focus on the stronger,⁸ well-studied *Restricted Isometry Property (RIP)*. We also note that the constructions of [Kar11] are highly structured, and so even though they also come from sparse matrices, they do not tell us anything about the ℓ_p -spread of sparse *random* matrices.

We prove that sparse random matrices are not only ℓ_p -spread, but are also ℓ_p -RIP, and, moreover, this follows merely from the expansion of the underlying bipartite graph of the random matrix A . In particular, we prove that *any* signed adjacency matrix B of a left-regular bipartite expander graph G is ℓ_p -RIP, provided that the maximum right degree s_{\max} is above a small threshold independent of n .

The RIP is a well-studied property of matrices from the compressed sensing literature, defined as follows.

Definition 1.2 (ℓ_p -RIP). Let $B \in \mathbb{R}^{m \times n}$ be a matrix. We say that B is (k, ε) - ℓ_p -RIP if there exists $K > 0$ such that for every k -sparse $x \in \mathbb{R}^n$, it holds that⁹

$$K(1 - \varepsilon) \|x\|_p \leq \|Bx\|_p \leq K(1 + \varepsilon) \|x\|_p \quad .$$

We note that ℓ_p -RIP implies ℓ_p -spread,⁸ and in fact it is a strictly stronger property [KT07].

RIP matrices have been studied extensively in the context of compressed sensing, as they yield a polynomial-time algorithm based on linear programming for the *robust sparse recovery problem*.

⁷This follows since $t = \alpha s \leq s$, $m = \alpha n$, and $s \leq n^c$ for some absolute constant c .

⁸See Proposition 3.8.

⁹We note that the standard definition of RIP typically appears without the normalization factor K above. We include the parameter K for convenience, as the random sparse matrices we consider are not normalized.

Namely, given a “noisy measurement sketch” $y = Bx + e$ of a vector x , where B is (k, ε) - ℓ_p -RIP and $\|e\|_p \leq \eta$, there is a polynomial-time algorithm to recover an estimate \hat{x} for x with the so-called “ ℓ_p - ℓ_1 guarantee,” namely the estimate \hat{x} satisfies $\|\hat{x} - x\|_p \leq O\left(k^{-(1-\frac{1}{p})} \|x - x^*\|_1 + \eta\right)$, where x^* is a k -sparse vector minimizing $\|x - x^*\|_1$ (see Appendix A in [AGR15] for details). We note that if B is merely ℓ_p -spread, then B suffices for the (non-robust) sparse recovery problem, i.e., when there is no noise e .

We now turn to formally stating our results. We first recall the definition of a (unique) bipartite expander.

Definition 1.3 (Unique expanders). A bipartite graph $G = (V_L = [n], V_R = [m], E)$ is a t -left-regular (γ, μ) -unique expander if (1) $\deg(u) = t$ for all $u \in V_L$, and (2) for all $S \subseteq V_L$, $|S| \leq \gamma n$, there are at least $t(1 - \mu)|S|$ vertices $v \in V_R$ which each have exactly one neighbor in S .

A matrix $B \in \{0, 1, -1\}^{m \times n}$ is a *signed adjacency matrix* of a bipartite graph $G = (V_L = [n], V_R = [m], E)$ if

$$B_{r,u} \neq 0 \iff (u, r) \in E$$

for all $u \in V_L, r \in V_R$.

Theorem 4 (ℓ_p -RIP of expander graphs). Let G be a bipartite t -left-regular (γ, μ) -unique expander with maximum right degree s_{\max} , and let B be any signed adjacency matrix of G . Let $0 < \varepsilon \leq 1$ and $1 \leq p < 2$ such that $\varepsilon^2 \geq 9\mu s_{\max}^{p-1}$. Then, B is $(\gamma n, \varepsilon)$ - ℓ_p -RIP, i.e., for every γn -sparse $x \in \mathbb{R}^n$,

$$t^{\frac{1}{p}}(1 - \varepsilon) \|x\|_p \leq \|Bx\|_p \leq t^{\frac{1}{p}}(1 + \varepsilon) \|x\|_p .$$

Theorem 4 generalizes a result of [BGI⁺08], which shows that any signed adjacency matrix B of G is ℓ_1 -RIP, provided that G is an expander. This is somewhat surprising, as the proof in [BGI⁺08] makes heavy use of properties specific to the ℓ_1 norm.¹⁰

The ℓ_p -RIP of matrices for general p has been studied in other contexts, most notably in [AGR15]. As is typical when studying RIP matrices, they view the sparsity parameter k as a fixed function of n , and determine m as a function of k, n . However, the results in [AGR15] are incomparable to ours, as they hold only for the low-sparsity case of $k = O(n^{1/p})$ (so $k = o(n)$ if $p > 1$), but we are concerned with the case of $k = \Omega(n)$, when the sparsity is a small constant fraction of n .

As a random t -left-regular bipartite graph is a good expander with high probability, we obtain the following corollary of Theorem 4, which shows that $\ker(A)$ for $A \leftarrow \mathcal{M}_{m,n,s,t}$ achieves very good ℓ_p -spread for every $p \in [1, 2)$. Thus, the poor ℓ_2 -spread of $\ker(A)$ is in fact specific to the case of $p = 2$.

Corollary 5 (Good ℓ_p -RIP and ℓ_p -spread of A). Fix $p \in [1, 2)$, $0 < \varepsilon < \frac{1}{2}$, and suppose that $s \geq \left(\frac{18}{\alpha \varepsilon^2}\right)^{\frac{1}{2-p}}$. Then, with high probability over A , the matrix A is $(\Omega(\gamma n), \varepsilon)$ - ℓ_p -RIP for $\gamma = \frac{\alpha^2}{t^4}$: for every $\Omega(\gamma n)$ -sparse $x \in \mathbb{R}^n$, it holds that

$$t^{\frac{1}{p}}(1 - \varepsilon) \|x\|_p \leq \|Ax\|_p \leq t^{\frac{1}{p}}(1 + \varepsilon) \|x\|_p .$$

In particular, the subspace $\ker(A)$ is $(\Omega(\gamma n), \Omega\left(\gamma^{1-\frac{1}{p}}\right))$ - ℓ_p -spread and $\Delta_p(\ker(A)) \leq O\left(1/\gamma^{2-\frac{2}{p}}\right)$.

¹⁰They also show that their proof for ℓ_1 -RIP extends to ℓ_p -RIP for $p \leq 1 + O\left(\frac{1}{\log n}\right)$, because the “Hölder factor” of $n^{1-\frac{1}{p}}$ is $O(1)$, but it does not extend to ℓ_p for any constant $p > 1$.

Fixing p, α, ε to be constants and taking s to be a large enough constant, this shows that $\ker(A)$ is $(\Omega(n), \Omega(1))$ - ℓ_p -spread with high probability, and therefore $\Delta_p(\ker(A)) = O(1)$. Together with Theorem 1, this shows that the ℓ_p -spread property of $\ker(A)$ exhibits an interesting *threshold phenomenon* at $p = 2$.

We also combine Theorem 4 with the explicit constructions of expander graphs of [CRVW02] to obtain the following corollary, which gives an explicit construction of ℓ_p -RIP matrices for all $p \in [1, p_0)$, where $1 < p_0 < 2$ is an absolute constant. We thus obtain the first explicit construction of a matrix B achieving the “ ℓ_p/ℓ_1 guarantee” for the robust sparse recovery problem, and our matrices are for the regime $k = \Theta(n)$ and any $p \in [1, p_0)$. Previously, such constructions were only known for $p \leq 1 + O\left(\frac{1}{\log n}\right)$ [BGI⁺08]. Unlike Corollary 5, our explicit constructions only extend up to some threshold $p_0 < 2$. This is because the expanders of [CRVW02] achieve weaker expansion than random graphs. Concretely, the “expansion error” μ of the [CRVW02] expanders is $\mu = O(1/t)^\tau$ for some constant $\tau < 1$, which yields the threshold of $p_0 = 1 + \tau$, whereas random graphs achieve $\mu = O(1/t)$, allowing for $p_0 = 2$.

Corollary 6 (Explicit construction of ℓ_p -RIP matrices). *Let $0 < \varepsilon < \frac{1}{2}$, $\alpha \in (0, 1)$, and let $n \in \mathbb{N}$ be sufficiently large. For some universal constant $1 < p_0 < 2$, there exists a deterministic algorithm which, given $p \in [1, p_0)$, ε, α and n , outputs in time $\text{poly}(n/\delta) + 2^{O(1/\delta)}$ a matrix $B \in \{0, 1\}^{m \times n}$, for some $m \leq \alpha n$, such that B is $(\gamma n, \varepsilon)$ - ℓ_p -RIP, for some $\delta, \gamma = \text{poly}(\varepsilon, \alpha)^{\frac{1}{p_0 - p}}$. In particular, $\ker(B)$ is $(\gamma n, \gamma^{1 - \frac{1}{p}})$ - ℓ_p -spread and $\Delta_p(\ker(B)) \leq 1/\gamma^{2 - \frac{2}{p}}$.*

Note that as ε, α and p are constants, the matrix B in Corollary 6 is $(\Omega(n), O(1))$ - ℓ_p -RIP, $\ker(B)$ is $(\Omega(n), \Omega(1))$ - ℓ_p -spread, and $\Delta_p(\ker(B)) \leq O(1)$.

As noted earlier, [Kar11] gives explicit constructions of $(\Omega(n), \Omega(1))$ - ℓ_p -spread subspaces, for all $1 \leq p < 2$. This is incomparable to Corollary 6: on one hand, [Kar11] obtains the full range of $1 \leq p < 2$, but on the other hand, his matrices only are ℓ_p -spread and do not satisfy the (strictly stronger) ℓ_p -RIP. Our construction is moreover the “simplest” black-box reduction to expansion: we show that the mere adjacency matrix of a bipartite expander is ℓ_p -RIP. While the constructions in [Kar11] are themselves not too complicated, we think that this is nonetheless an interesting conceptual contribution of our work.

Finally, we also prove the following partial converse to Corollary 5, which shows that when $s^{2-p} \lesssim \frac{1}{\alpha}$ (i.e., s^{2-p} is a constant factor below the threshold in Corollary 5), then A is *not* ℓ_p -RIP.

Theorem 7 (Partial converse to Corollary 5). *Let $p \in [1, 2)$, $\varepsilon > 0$. If $s - 1 \leq \left(\frac{1}{(1+\varepsilon)\alpha}\right)^{\frac{1}{2-p}}$, then with high probability over A , there exists an n^c -sparse vector $x \in \mathbb{R}^n \setminus \{0^n\}$ such that*

$$\frac{\|Ax\|_p}{\|x\|_p} \leq t^{\frac{1}{p}} \cdot m^{-\Omega\left(\frac{\varepsilon}{\log s}\right)}.$$

Note that $\|Ae_1\|_p / \|e_1\|_p = t^{\frac{1}{p}}$ always holds, so Corollary 5 demonstrates that, given small enough s , the ratio $\frac{\|Ax\|_p}{\|x\|_p}$ has a large range over different choices of n^c -sparse x .

2 Proof overview

We outline the proofs of our results. For the purposes of this exposition, we will adopt the same convention as in Section 1.2 and use A and B to denote a uniformly sampled matrix from $\mathcal{M}_{m,n,s,t}$ and arbitrary matrix from $\{0, 1, -1\}^{m \times n}$, respectively. Recall that $\mathcal{M}_{m,n,s,t}$ denotes the set of (s, t) -biregular matrices with entries in $\{0, 1, -1\}$, and that $\frac{m}{n} = \frac{t}{s} = \alpha$ for some constant α , and $n \rightarrow \infty$. For simplicity of this exposition, in this section we restrict ourselves to the regime $s = O(1)$ unless stated otherwise.

We naturally associate with B the bipartite graph $G = G_B = (V_L, V_R, E)$ with $n = |V_L|$ left vertices, $m = |V_R|$ right vertices, and an edge between $u \in V_L$ and $r \in V_R$ if $B_{r,u} \neq 0$. We view the rows and columns of B as indexed by V_R and V_L , respectively, and identify \mathbb{R}^n with \mathbb{R}^{V_L} , and \mathbb{R}^m with \mathbb{R}^{V_R} . In addition, we define the function $\text{sign} = \text{sign}_B : E \rightarrow \{1, -1\}$, which maps an edge $\{u, r\}$ as above to $B_{r,u}$. We note that the combination of G_B and sign_B completely describes B .

2.1 Theorem 1: $\ker(A)$ in not ℓ_2 -spread

For simplicity, we only sketch here why $\ker(A)$ is likely to contain an $(o(n), o(1))$ -compressible vector, and leave the more refined parameter setting for the actual proof of Theorem 1 in Section 5.

The proof of Theorem 1 consists of two steps. In the first step we find an $o(n)$ -sparse vector $x \in \mathbb{R}^n$ with $\|x\|_2 \geq 1$ and $\|Ax\|_2 \leq o(1)$ (Lemma 5.1). In the second step we find a vector $y \in \ker(A)$ with $\|y - x\|_2 \leq o(1) \cdot \|y\|_2$ (Lemma 5.2 and Theorem 2). In particular, y is $(o(n), o(1))$ -compressible, so $\ker(A)$ cannot be ℓ_2 -spread.

Below, we outline these two steps. It is straightforward to see, given the construction described below, that both x and y can be computed in polynomial time given A .

We also note that an ℓ_p analog of Step 1 is the main technical component in the proof of Theorem 7, and is also proven in Section 5.

Step 1: constructing a sparse x with small $\|Ax\|_2$. To obtain the vector x , we first prove (Proposition 4.1) that G is highly likely to contain a vertex $v^* \in V_L$ such that the ball of radius $2\ell + 1$ about v^* , for some $\ell \leq O(\log n)$, contains no cycles. That is, the radius- ℓ neighborhood of v^* is a complete (t, s) -biregular tree T rooted at v^* . Recall that a rooted tree is (t, s) -biregular if the even depth (resp. odd depth) inner vertices have degree t (resp. s). The existence of such a vertex v^* is the only random property of A needed in this step of the proof. In particular, assuming that G has the aforementioned property, our construction of x is always possible, regardless of the sign function.

To describe the construction of x itself, we assume for simplicity that $\text{sign}(e) = 1$ for all $e \in E$. Namely, all the non-zero entries of A are 1. In this setting, let $v \in V_L \cap T$ be a vertex of depth $2k$ in the tree for some $k \geq 0$ (note that a vertex in V_L must have even depth), and set $x_v = (-(s-1))^{-k}$. For any $x \in V_L \setminus T$, set $x_v = 0$. Note that $\text{supp}(x) \subseteq T$. We choose ℓ above to be as large as possible, i.e., $O(\log n)$, so that the size of T is roughly n^c for some $c < 1$. In particular, x is $\approx n^c$ -sparse. Also, note that $\|x\|_2^2 \geq x_{v^*}^2 = 1$. We informally refer to the vector x produced by this construction as a *tree vector*.

Our construction guarantees that $(Ax)_r = 0$ for every internal node $r \in T \cap V_R$. Indeed, suppose that r is of depth $2k+1$. Then, it has one neighbor of depth $2k$, and $s-1$ neighbors of depth $2k+2$. As $(Ax)_r$ is the sum of x_v over neighbors v of r , we have $(Ax)_r = (-(s-1))^{-k} + (s-1) \cdot (-(s-1))^{-(k+1)} = 0$.

To compute $\|Ax\|_2$, it thus suffices to compute $|(Ax)_r|$ when r is one of the $t(t-1)^\ell (s-1)^\ell$ leaves

of T . It is not hard to see that in this case $|(Ax)_r| = (s - 1)^\ell$, and so

$$\|Ax\|_2^2 = t(t - 1)^\ell (s - 1)^\ell \cdot (s - 1)^{-2\ell} = e^{-\Omega(\ell)} = o(1) .$$

We note that our tree vector construction is similar in spirit to a construction by Noga Alon [GLW08, Theorem 8], which demonstrates the limitations of expander-based analysis of the spread property. In [GLW08], however, they *choose* their graph G so that (their analog of) the tree vector x will lie in (their analog of) $\ker(A)$ *by design*. Our graph is random and not up to our choice, so we cannot simply orchestrate the graph so that our tree vector x to belong to $\ker(A)$. This necessitates that we perform the nontrivial step of rounding x to some $y \in \ker(A)$, which we discuss next.

Step 2: finding $y \in \ker(A)$ close to x . Our main goal in this step is to establish the following lemma:

Lemma 2.1 (Informal). *With high probability over A , it holds that every $x \in \mathbb{R}^n$ with $\|Ax\|_2 \leq o(1)$ is $o(1)$ -close to some vector $y \in \ker(A)$.*

Indeed, let x be the tree vector constructed in Step 1. Then x is $o(n)$ -sparse with $\|x\|_2 \geq 1$ and $\|Ax\|_2 \leq o(1)$. By Lemma 2.1, there exists a vector $y \in \ker(A)$, which is $o(1)$ -close to x . This vector y is $(o(n), o(1))$ -compressible, which yields Theorem 1 in the present parameter setting.

One may naively try to prove Lemma 2.1 by locally perturbing x to try to make $Ax = 0^m$, e.g. by designing a greedy algorithm for this task. This approach, however, seems difficult to execute, especially given that Lemma 2.1 is in fact not true in general. For example, it could be the case that x is a (unit norm) right singular vector of A with singular value $o(1)$. Then, $\|Ax\|_2 = o(1)$, but $\|x - y\|_2 \geq 1$ for all $y \in \ker(A)$, and in fact the closest vector in $\ker(A)$ to x is 0^n .

Instead, we set y to be the orthogonal projection of x onto $\ker(A)$. In hindsight, this is the obvious choice for y , as then $y \in \ker(A)$ is the vector that minimizes $\|x - y\|_2$. How large can $\|x - y\|_2$ be? Intuitively, we would like to say that $\|Ax\|_2$ being small implies that $\|x - y\|_2$ is small as well. As the earlier example shows, this is not true for a general matrix A , as A could have small singular values. However, the implication *does* hold provided that all singular values of A are $\Omega(1)$.¹¹ Indeed, the singular value decomposition of A implies that

$$\|Ax\|_2 = \|A(x - y)\|_2 \geq \sigma_{\min}(A) \|x - y\|_2 ,$$

where $\sigma_{\min}(A)$ is the minimum singular value of A and the inequality holds as $x - y$ is orthogonal to $\ker(A)$. Hence, $\|x - y\|_2 \leq \frac{\|Ax\|_2}{\sigma_{\min}(A)}$.

The main technical component of Step 2 is therefore the lower bound on $\sigma_{\min}(A)$, given by Theorem 2. As we have argued above, the crude lower bound of $\sigma_{\min}(A) \geq \Omega(1)$ suffices to yield Lemma 2.1. Indeed, if $\sigma_{\min}(A) \geq \Omega(1)$, then $\|x - y\|_2 \leq o(1)$, and so $\ker(A)$ contains an $(o(n), o(1))$ -compressible vector. The precise high-probability lower bound on $\sigma_{\min}(A)$ established in Theorem 2 implies a finer quantitative version of Lemma 2.1, which yields the full Theorem 1. The latter gives a much sharper bound on the $o(1)$ term, and also applies to sparsity all the way up to $O(n^c)$ for some $c > 0$.

We remark that one can easily show that $\sigma_{\min}(A) \geq \sqrt{s} - O(\sqrt{t})$ via “off-the-shelf” methods, such as [BV16]. However, this would only allow us to prove Theorem 1 provided that $\alpha \leq c$ for

¹¹Technically, what matters is the minimum *nonzero* singular value. However, with high probability the matrix A will be full rank (i.e., rank m), so that $\sigma_{\min}(A) > 0$. Indeed, this is trivially implied by Theorem 2.

some absolute constant $c < 1$ (related to the $O(1)$ factor in front of \sqrt{t} above), and thus would not allow us to take α to be *any* constant in $(0, 1)$, e.g., $\alpha = 0.999$. Our sharper bound also highlights the difficulty in lower bounding the minimum singular value when $\alpha = m/n$ is close to 1.

We postpone our discussion of the proof of Theorem 2 to Section 2.3, and turn next to our positive result for ℓ_p -spread for $p < 2$.

2.2 Theorem 4: ℓ_p -RIP for $p < 2$ from vertex expansion

We sketch the proof of Theorem 4. For simplicity, we will assume that $B \in \mathcal{M}_{m,n,s,t}$, i.e., that the bipartite graph G_B is t -left-regular and s -right-regular (and hence $s_{\max} = s$) and also that G_B is a (γ, μ) -unique expander. For this exposition, we only discuss the claimed lower bound on $\|Bx\|_p$ stated in the theorem, namely,

$$\|Bx\|_p \geq t^{\frac{1}{p}}(1 - \varepsilon) \|x\|_p \quad (1)$$

for all γn -sparse $x \in \mathbb{R}^n$, as the upper bound is obtained via a variation on the same method.

Theorem 4 for tree vectors. As a warm-up for the proof of Theorem 4, we show why the $o(n)$ -sparse tree vector x constructed in Section 2.1 does not yield a counterexample to Eq. (1). Let $x^{(k)}$ ($0 \leq k \leq \ell$) denote the restriction of x to the vertices in the $2k$ -th level of the tree T . Then,

$$\|x^{(k)}\|_p^p = t(t-1)^{k-1}(s-1)^{(1-p)k} .$$

For $p = 2$, this expression decreases exponentially in k , and thus the ℓ_2 -mass of x is concentrated at the top of the tree. For $p < 2$, however, $\|x^{(k)}\|_p^p$ actually *grows* exponentially in k provided that s is large enough (concretely, one needs $s^{2-p} \gtrsim \frac{1}{\alpha}$).¹² In this case, the ℓ_p -mass is concentrated towards the bottom of the tree. Moreover, one can take s large enough so that all but an $\frac{\varepsilon}{2}$ -fraction of the mass lies in the bottom layer. Then,

$$\frac{\|Bx\|_p^p}{\|x\|_p^p} \geq \left(1 - \frac{\varepsilon}{2}\right) \cdot \frac{\|Bx\|_p^p}{\|x^{(\ell)}\|_p^p} = \left(1 - \frac{\varepsilon}{2}\right) \cdot \frac{t(t-1)^\ell (s-1)^{(1-p)\ell}}{t(t-1)^{\ell-1}(s-1)^{(1-p)\ell}} = \left(1 - \frac{\varepsilon}{2}\right) \cdot (t-1) \geq (1 - \varepsilon)t .$$

Hence, x is not a counterexample to Eq. (1).

Theorem 4 for general vectors with tree-shaped support. Fix a set $S \subseteq V_L$ such that the subgraph induced by $T := S \cup N(S)$ is a (t, s) -biregular tree. We generalize the above discussion of tree vectors by explaining why Eq. (1) holds for any vector $x \in \mathbb{R}^n$ supported on S .

Given $r \in N(S)$, let v_r denote the parent of r in the tree T . In an overly optimistic scenario, if we could show that $|(Bx)_r| \approx |v_r|$ for all $r \in N(S)$, then we would be done, as each vertex $v \in S$ has $t-1$ children.¹³ Each of these children then contributes $\approx |x_v|^p$ mass to $\|Bx\|_p^p$, so that $\|Bx\|_p^p \approx (t-1) \cdot \sum_{v \in S} |x_v|^p = (t-1) \cdot \|x\|_p^p$, implying Eq. (1). As the tree vector case shows, one cannot, in fact, hope to guarantee $|(Bx)_r| \approx |v_r|$ for all $r \in N(S)$. Indeed, for a tree vector x we have $(Bx)_r = 0$ for any non-leaf $r \in N(S)$. Thus, a more delicate analysis is required.

For intuition, let us consider the viewpoint of an adversary seeking to construct an x supported on S such that $\|Bx\|_p^p$ is small. We shall think of the adversary as assigning values to $\{x_v\}_{v \in S}$ starting

¹²And, indeed, if instead $s^{2-p} \lesssim \frac{1}{\alpha}$, then we have $\|Bx\|_p = o(1)$, and this gives us Theorem 7.

¹³Except for the root, which has t children.

from the root, and then moving down the tree.

For each non-leaf $r \in N(S)$, let W_r denote the set of $s-1$ children of r . Recall that v_r is the parent of r . Note that $|(Bx)_r| \geq |x_{v_r}| - \sum_{u \in W_r} |x_u|$ due to the triangle inequality. Hence, when assigning values to the vertices in W_r , the adversary morally has two choices: (1) either make $\sum_{u \in W_r} |x_u| \ll |x_{v_r}|$, in which case $|(Bx)_r| \approx |x_{v_r}|$, or (2) make $|(Bx)_r| \approx 0$, in which case $\sum_{u \in W_r} |x_u| \approx |x_{v_r}|$.¹⁴ Let us fix some $\beta < 1$, and suppose that the adversary chooses values for $\{x_u\}_{u \in W_r}$ such that $\sum_{u \in W_r} |x_u| \leq \beta |x_{v_r}|$, i.e., the adversary chooses Case (1). We then have that

$$|(Bx)_r| \geq \left(|x_{v_r}| - \sum_{u \in W_r} |x_u| \right) \geq (1 - \beta) |x_{v_r}| ,$$

so $|(Bx)_r| \approx |x_{v_r}|$, which is what we wanted. Next, suppose that the adversary makes $\sum_{u \in W_r} |x_u| \geq \beta |x_{v_r}|$, i.e., the adversary chooses Case (2). Then, $|(Bx)_r|$ can be small, but applying Hölder's inequality, we have

$$\sum_{u \in W_r} |x_u|^p \geq \frac{\beta^p}{(s-1)^{p-1}} \cdot |x_{v_r}|^p .$$

Now, suppose that all children r of v_r have this property. Then, the total ℓ_p^p mass of all of the grandchildren of v_r must be at least $\frac{(t-1)\beta^p}{(s-1)^{p-1}} \cdot |x_{v_r}|^p \gg |x_{v_r}|^p$. We thus see that, intuitively, the adversary has merely pushed its task down to the grandchildren of v_r , and in doing so has not made any progress towards its overall goal. Indeed, this is precisely what happens in the case of a tree vector!

The above informal argument shows that the adversary does not “win” in either case. We can concretely capture this intuition via the following potential function:

$$a_r(x) := |(Bx)_r|^p + \Theta(1) \cdot \frac{(s-1)^{p-1}}{\varepsilon^{p-1}} \sum_{u \in W_r} |x_u|^p .$$

In the actual proof, this choice of $a_r(x)$ allows us to cleanly express the intuition that either $|(Bx)_r|$ is large or $\sum_{u \in W_r} |x_u|$ is large, and further extends beyond the “toy case” of tree-supported vectors.

Using expansion when $S \cup N(S)$ is not a tree. We now turn to the general case, where the subgraph induced by $S \cup N(S)$ (where $S = \text{supp}(x)$) is not necessarily a tree. We observe that above, we are only using the tree structure to show that the rooted tree $S \cup N(S)$ trivially has a 1-to- $(t-1)$ “matching” with the following properties: (1) every vertex $r \in N(S)$ is matched with exactly one vertex $v \in S$, and (2) every vertex $v \in S$ is matched with at least $t-1$ vertices in $N(S)$. Indeed, when $S \cup N(S)$ is a tree, such a matching exists by matching each vertex $r \in N(S)$ with its parent v_r .

To generalize the above, we use the (unique) expansion of G to construct a similar matching that suffices for the proof. Recall that G is a (γ, μ) -unique expander, meaning that every set $S \subseteq V_L$ of size $\leq \gamma n$ has at least $t(1-\mu)|S|$ unique neighbors, i.e., neighbors of a unique element of S . We construct the matching by “peeling off” vertices one at a time from S , each time matching a vertex with $\geq t(1-\mu)$ vertices in $N(S)$, namely its neighbors that are not neighbors of any of the remaining “unpeeled” vertices in S .

The above step can be viewed as extracting a “tree-like” subgraph from $S \cup N(S)$, where each vertex $v \in S$ has at least $t(1-\mu)$ “children” (the vertices it was matched with), and at most μt “parents” (its neighbors that it was *not* matched with). Each vertex $r \in N(S)$ still has exactly one

¹⁴Note that the adversary has the third choice of setting $\sum_{u \in W_r} |x_u| \gg |x_{v_r}|$, but this is worse for the adversary.

“parent” and $\leq s - 1$ “children”. Once we have the above “tree-like” subgraph, the argument for trees goes through with only minor modifications, so this finishes the proof.

We note that the existence of this “tree-like” subgraph for any set S with $|S| \leq \gamma n$ immediately implies that G is a (γ, μ) -vertex expander, and hence a $(\gamma, 2\mu)$ -unique expander. Thus, the existence of such a subgraph for every S of size at most γn is equivalent to unique expansion, up to a factor of 2 loss in the parameter μ .

Comparison with [BGI⁺08]. We briefly summarize the proof in [BGI⁺08] for the case of $p = 1$, and explain why their proof does not extend to the case of $p > 1$.

The proof in [BGI⁺08] proceeds as follows. For a vector x supported on S , let E_0 denote the set of edges between S and $N(S)$. First, they match each $r \in N(S)$ to its neighbor $v \in S$ with $|x_v|$ maximized. Let E_1 be the set of edges in this matching, and let $E_2 = E_0 \setminus E_1$. For any $r \in N(S)$, it then follows that $|(Bx)_r| \geq |x_v| - \sum_{u \in W_r} |x_u|$, where $(v, r) \in E_1$ and $W_r = \{u : (u, r) \in E_2\}$. Hence, $\|Bx\|_1 \geq \sum_{(v,r) \in E_1} |x_v| - \sum_{(u,r) \in E_2} |x_u|$. We observe that this step of the proof is specific to the ℓ_1 norm, and does not generalize to larger ℓ_p norms.

The main step in the proof is to argue that $\sum_{(v,r) \in E_2} |x_v| \leq t\varepsilon \|x\|_1$ using expansion. With this in hand, it immediately follows that $\sum_{(v,r) \in E_1} |x_v| \geq t(1 - \varepsilon) \|x\|_1$, because $\sum_{(v,r) \in E_0} |x_v| = t \|x\|_1$ by regularity. It then follows that $\|Bx\|_1 \geq t(1 - 2\varepsilon) \|x\|_1$. Note that the upper bound $\|Bx\|_1 \leq t \|x\|_1$ is trivial, so this shows that B is ℓ_1 -RIP.

One may attempt to generalize this proof to $p > 1$ by replacing $|x_v|$ with $|x_v|^p$. For example, using expansion it follows that $\sum_{(v,r) \in E_2} |x_v|^p \leq t\varepsilon \|x\|_p^p$, and as $\sum_{(v,r) \in E_0} |x_v|^p = t \|x\|_p^p$, we then have $\sum_{(v,r) \in E_1} |x_v|^p \geq t(1 - \varepsilon) \|x\|_p^p$. But this is not enough to complete the proof, as it does *not* follow that $\|Bx\|_p^p \geq \sum_{(v,r) \in E_1} |x_v|^p - \sum_{(v,r) \in E_2} |x_v|^p$. Indeed, this is a fundamental barrier, and is the reason why our analysis for $p \geq 1$ proceeds by analyzing the “local” potential function $a_r(x)$, rather than the two “global” sums over E_1 and E_2 above.

2.3 Theorem 2: bounds on the singular values of A

We give a brief overview of the proof of Theorem 2. First, we observe that in order to bound the singular values of A , it suffices to bound the spectrum of $M := AA^\top - s \cdot I$, as each singular value of A is the square root of an eigenvalue of AA^\top . Note that M is a square matrix with an all-0 diagonal, by regularity of A .

Step 1: reducing to the nomadic walk matrix via a modified Ihara–Bass formula. The first step in the proof is to relate bounds on the spectrum of M to the spectral radius (i.e., maximum eigenvalue in absolute value) $\rho(B)$ of B , the *nomadic walk matrix* introduced in [MOP20b].¹⁵ The nomadic walk matrix B is indexed by pairs of edges¹⁶ (e_1, e_2) in G that form a length 2 non-backtracking walk in G , and its $((e_1, e_2), (e'_1, e'_2))$ -th entry is $\text{sign}(e'_1)\text{sign}(e'_2)$ if $e_1 \rightarrow e_2 \rightarrow e'_1 \rightarrow e'_2$ forms a non-backtracking walk of length 4 in G , and is 0 otherwise. Note that B is *not* symmetric.

Theorem 2.2 (Modified Ihara–Bass formula, Theorem 3.1 of [MOP20b], informal). *If $\rho(B) \leq$*

¹⁵We note that one could most likely also prove Theorem 2 using the standard nonbacktracking walk matrix and Ihara–Bass formula, e.g., with similar methods as in [BDH18].

¹⁶In [MOP20b], the nomadic walk matrix is indexed by *directed* edges. In our context, this is equivalent to a length 2 oriented walk $e_1 \rightarrow e_2$, which is equivalent to a pair (e_1, e_2) of *undirected* edges, as the ordering in the pair gives the unique orientation $e_1 \rightarrow e_2$ in the walk.

$(1 + o(1))\sqrt{(s-1)(t-1)}$, then the spectrum $\text{Spec}(M)$ of M satisfies:

$$\text{Spec}(M) \subseteq [t - 2 - 2(1 + o(1))\sqrt{(s-1)(t-1)}, t - 2 + 2(1 + o(1))\sqrt{(s-1)(t-1)}] .$$

The above theorem thus shows that it suffices to prove that $\rho(B) \leq (1 + o(1))\sqrt{(s-1)(t-1)}$ with high probability.

We remark that bounds on the spectra of matrices of the form of M were studied in [MOP20b] for the case of $s = O(1)$. Unfortunately, this is insufficient to prove Theorem 2, as we wish to allow s to be any function of n (provided that $s \leq n^c$ for some absolute constant c). However, [MOP20b, Theorem 3.1] is a general statement that holds regardless of s , so we can make use of it in our setting.

Step 2: bounding $\rho(B)$ via the trace method by counting hikes. The natural approach to bound $\rho(B)$ is by applying the trace method. As the matrix B is not symmetric, we compute:

$$T := \mathbb{E}_{\text{sign}}[\text{tr}(B^\ell (B^\top)^\ell)] ,$$

where the expectation is taken over the function sign that determines the signs of the entries of A . By carefully expanding this expectation, one can show that the nonzero contributions to T roughly come from length $4(\ell - 1)$ closed walks in G where (1) each edge in the walk appears an even number of times, and (2) the walk is non-backtracking, except possibly at the middle step in the walk. Such walks (of length 4ℓ) are commonly referred to as (2ℓ) -*hikes* [MOP20a].

To finish the proof of Theorem 2, we thus turn to obtaining a careful bound on the number of such walks.

Counting these walks requires extra care in our setting as our graph is bipartite, and so the bound needs to be sensitive to the difference in right/left degree. The counting of such hikes also differs greatly depending on whether $s \leq \text{polylog}(n)$ or $s = \omega(\text{polylog}(n))$.

Step 3: counting the number of hikes when $s \leq \text{polylog}(n)$. [MOP20a, Section 3] counts the number of such hikes when $s = O(1)$, provided that G is bicycle-free at radius $O(\log n)$. Namely, any vertex v participates in at most one cycle of length $O(\log n)$. By repeating their proof, one can show that their bounds can be extended to the case when $s \leq \text{polylog}(n)$. However, we still cannot use their bound on the number of such hikes naively, as their counting is for non-bipartite graphs and thus yields a bound of $m(1 + o(1))^\ell (s-1)^{2\ell}$, simply because it treats left and right vertices the same, and the maximum degree of a vertex is s . We refine their approach to ensure that right and left vertices contribute roughly equally, which will yield the desired bound. One may, at first glance, be tempted to assume that this is trivial because a closed walk in a bipartite graph has an equal number of left and right vertices, but this is not the case, as we shall see.

We adopt the bookkeeping approach of [MOP20a]. We think of a hike as discovering the graph G as one traverses the hike. A step in a hike is *fresh* if it uses an edge for the first time and ends at a previously undiscovered vertex; it is *boundary* if it uses an edge for the first time but ends up at an old vertex; finally, it is *stale* if it uses an old edge.

Because each edge must appear an even number of times, a hike can have at most 2ℓ fresh steps. Each fresh step “pays” a factor of $(s-1)$ (if we move from a right to a left vertex) or $(t-1)$ (if we move from a left to a right vertex) in our bound in the number of hikes, as this is the number of choices for the next vertex that the hike moves to. [MOP20a, Theorem 2.13] implies that, since G is bicycle-free, the number of boundary steps is $\ll \ell$; they also show that one can bound the “number

of choices” for the stale steps by some small factor, which we ignore here.

We need to augment the argument of [MOP20a] with the following addition: if a hike has c fresh steps, then the number of fresh steps c_R that start at a right vertex is $\approx \frac{c}{2}$, and similarly for c_L for left vertices. Note that by definition, $c = c_R + c_L$. The key observation is that a fresh step from a right (resp. left) vertex must be followed by either a fresh step from a left (resp. right) vertex, or by a boundary step. Indeed, after we take a fresh step we are at a previously unvisited vertex, so the next step must use a new edge; in particular, it cannot be *stale*. This implies that the deviation of each of c_L, c_R from $\frac{c}{2}$ is bounded by the number of boundary steps, which is $\ll \ell$.

This implies a bound of $m(1 + o(1))^\ell (s - 1)^\ell (t - 1)^\ell$ on the number of hikes, provided that $s \leq \text{polylog}(n)$. The m comes from the number of start vertices in the hike, and the $(1 + o(1))^\ell$ comes from the stale and boundary steps, as well as our new deviation term. This yields the desired bound for sparse s .

Step 4: counting the number of hikes when $s = \omega(\text{polylog}(n))$. For s this large, the graph G is “dense”, and so it will not be bicycle-free at radius $O(\log n)$. This rules out the approach of [MOP20a], which relies on G being bicycle-free. Instead, we adapt a standard counting approach (for bounding the operator norm of a random $n \times n$ Gaussian matrix) given in [Tao12, Section 2.3.6] to our bipartite setting. Our crucial observation here is to note that any hike can have at most ℓ distinct left vertices. As we pay a factor of $(s - 1)$ every time we move to a left vertex, it then follows that the “power” of $(s - 1)$ in our bound can only be at most ℓ . The standard counting argument of [Tao12, Section 2.3.6] for Gaussian matrices then goes through, yielding the desired bound.

2.4 Organization

The rest of the paper is organized as follows. Section 3 introduces notation and definitions that we use in our proofs. In Section 4, we state and prove several claims about properties of a certain random bipartite graph naturally associated with the random matrix A . In Section 5, we prove our negative results, namely, Theorems 1 and 7. In Section 6 we prove Theorem 4 and Corollaries 5 and 6, our positive results about ℓ_p for $p < 2$. Our proofs in Section 5 rely on the singular value bounds stated in Theorem 2, which we prove in Section 7. Finally, in Appendix A we prove Theorem 3, our positive result for ℓ_2 -spread.

3 Preliminaries

3.1 Basic notation

For an integer $k \in \mathbb{N}$, let $[k] := \{1, \dots, k\}$.

For a vector $x \in \mathbb{R}^n$, we let $\text{supp}(x) := \{i \in [n] : x_i \neq 0\}$. We say x is k -sparse if $|\text{supp}(x)| \leq k$ and $x \neq 0^n$.

Given a graph G , a vertex u , and $\ell \in \mathbb{N}$, we let $B_G(u, \ell)$ denote the induced subgraph on vertices of distance at most ℓ from u . For a set of vertices S , we let $B_G(S, \ell) := \bigcup_{u \in S} B_G(u, \ell)$.

Let $\sigma(A)$ denote the set of singular values of a matrix A . We also denote $\sigma_{\min}(A) = \min \sigma(A)$ and $\sigma_{\max} = \max \sigma(A)$. If A is square, we denote the set of its eigenvalues by $\text{Spec}(A)$, and its spectral radius by $\rho(A) := \max_{\lambda \in \text{Spec}(A)} |\lambda|$. We remark that unlike in Sections 1.2 and 2, we will now use A to denote arbitrary $m \times n$ matrices, not just uniformly random ones from $\mathcal{M}_{m,n,s,t}$.

3.2 Biregular bipartite graphs and sparse matrices

Let $G = (V_L, V_R, E)$ denote a bipartite graph with $n = |V_L|$ left vertices and $m = |V_R|$ right vertices. We say that G is t -left-regular if $\deg(u) = t$ for all $u \in V_L$. If, additionally, $\deg(v) = s$ for all $v \in V_R$, we say that G is (t, s) -biregular.

For a vertex v , we let $N_G(v)$ denote the set of neighbors of v in G . For a set of vertices S , we let $N_G(S) := \cup_{v \in S} N_G(v)$ denote the set of neighbors of S in G , and we let $U_G(S)$ denote the set of unique neighbors of S , i.e., $U_G(S)$ is the set of $v \in N_G(S)$ such that $v \in N_G(u)$ for some $u \in S$ and $v \notin N(u')$ for all other $u' \in S$ with $u \neq u'$. We omit the subscript G when it is clear from the context. For $S \subseteq V_L$ and $T \subseteq V_R$, we let $E(S, T) = E_G(S, T) = \{\{u, v\} \mid u \in S \wedge v \in T\}$.

Let $\mathcal{G}_{m,n,s,t}$ denote the set of all (t, s) -biregular graphs with $|V_L| = n$ and $|V_R| = m$. Namely, $\mathcal{G}_{m,n,s,t} = \{G_A \mid A \in \mathcal{M}_{m,n,s,t}\}$ (see the beginning of Section 2 for a reminder of the definition of G_A). We note that sampling a uniformly random element of $\mathcal{M}_{m,n,s,t}$ is equivalent to sampling a uniformly random graph G_A from $\mathcal{G}_{m,n,s,t}$ and then sampling a uniformly random edge signing $\text{sign}_A : E \rightarrow \{\pm 1\}$.

We also define the following properties.

Definition 3.1. A graph G is said to be *bicycle-free* at radius $\ell \in \mathbb{N}$ if for all $v \in V$, the ball $B_G(v, \ell)$ contains at most one cycle.

Definition 3.2 (Expansion and unique expansion). Let G be a t -left-regular graph. Fix $\gamma, \mu \in (0, 1)$. We say that G is:

- (1) a (γ, μ) -vertex expander if for all $S \subseteq V_L$ with $|S| \leq \gamma n$, it holds that $|N(S)| \geq t(1 - \mu)|S|$,
- (2) a (γ, μ) -unique expander if for all $S \subseteq V_L$ with $|S| \leq \gamma n$, it holds that $|U(S)| \geq t(1 - \mu)|S|$.

We note the following proposition, which shows that expansion implies unique expansion, up to a factor 2 in μ .

Proposition 3.3 (Expansion implies unique expansion). *Let G be a t -left-regular (γ, μ) -vertex expander. Then, G is a $(\gamma, 2\mu)$ -unique expander.*

Proof. Fix $S \subseteq V_L$ with $|S| \leq \gamma n$. Let $T = N(S)$, $T_1 = U(S)$, and $T_2 = T \setminus T_1$. We have $|T_1| + 2|T_2| \leq |E(S, T)|$, as each $v \in T$ must have at least two edges to S , and $|E(S, T)| \leq t|S|$, as G is t -left-regular. As $|T_1| + |T_2| = |T|$, we thus have $2|T| - |T_1| \leq t|S|$. As $|T| \geq t(1 - \mu)|S|$, it follows that $|T_1| \geq t(1 - 2\mu)|S|$, which finishes the proof. \square

We recall the formal guarantees of the explicit expander construction of [CRVW02].

Theorem 3.4 ([CRVW02, Theorem 7.1], rephrased). *For some universal $c \geq 1$, there is a deterministic algorithm which, given $\beta \in (0, 1)$ an inverse power of 2, $\mu > 0$ and $n \in \mathbb{N}$ a sufficiently large power of 2, outputs in time $\text{poly}\left(n, \frac{1}{\beta}, \frac{1}{\mu}\right) + \text{poly}\left(2^{\frac{1}{\beta}}, 2^{\frac{1}{\mu}}\right)$ a bipartite graph $G = (V_L, V_R, E)$ with $|V_L| = n$, $|V_R| = m := \beta n$, such that G is a t -left-regular $(\Omega(\frac{\mu m}{t}), \mu)$ -unique expander, for $t \leq O\left(\left(\frac{1}{\beta \mu}\right)^c\right)$.*

We note that [CRVW02] only constructs vertex expanders, not unique expanders. However, as Proposition 3.3 shows, these are equivalent up to a factor of 2 in the parameter μ , so this does not matter.

3.3 ℓ_p norms, compressibility, spread subspaces, and distortion

Let $p \in \mathbb{R}$ satisfy $1 \leq p \leq \infty$. Given a matrix $A \in \mathbb{R}^{m \times n}$ we denote its $\ell_p \rightarrow \ell_p$ operator norm by

$$\|A\|_p = \sup_{x \in \mathbb{R}^n \setminus \{0\}} \frac{\|Ax\|_p}{\|x\|_p}.$$

We recall the following two facts about ℓ_p norms.

Lemma 3.5 (Hölder's inequality). *Let $x, y, z \in \mathbb{R}^n$ be vectors, with $z_i = x_i y_i$ for all $i \in [n]$. Then $\|z\|_1 \leq \|x\|_p \|y\|_{\frac{p}{p-1}}$.*

Furthermore, let $1 \leq q < p$. Then $\|x\|_p \leq \|x\|_q \leq \|x\|_p n^{1/q-1/p}$ for any $x \in \mathbb{R}^n$.

Next, we recall the definition of compressible and spread vectors, as well as spread subspaces.

Definition 3.6 (ℓ_p -spread). Fix $p \in [1, \infty]$, $\varepsilon \in [0, 1]$ and $k \leq n \in \mathbb{N}$. A vector $y \in \mathbb{R}^n$ is k -sparse if $|\text{supp}(y)| \leq k$. A vector $x \in \mathbb{R}^n \setminus \{0^n\}$ is said to be (k, ε) - ℓ_p -compressible if there exists a k -sparse $y \in \mathbb{R}^n$ such that $\|x - y\|_p \leq \varepsilon \|x\|_p$. Otherwise, we say that x is (k, ε) - ℓ_p -spread.

A subspace $X \subseteq \mathbb{R}^n$ is (k, ε) - ℓ_p -spread if every $x \in X \setminus \{0^n\}$ is (k, ε) - ℓ_p -spread.

We note that ℓ_p -spread implies ℓ_q -spread, for any $q \leq p$, up to a small change in parameters. Formally, the following proposition holds.

Proposition 3.7 (ℓ_p -spread implies ℓ_q -spread). *Suppose that $X \subseteq \mathbb{R}^n$ is $(2k, \varepsilon)$ - ℓ_p -spread. Then for every $1 \leq q < p$, X is (k, ε_q) - ℓ_q -spread for $\varepsilon_q = \varepsilon^2 \left(\frac{k}{n}\right)^{\frac{1}{q}}$.*

In particular, if X is $(\Omega(n), \Omega(1))$ - ℓ_p -spread, then X is also $(\Omega(n), \Omega(1))$ - ℓ_q -spread for every $1 \leq q < p$.

We prove Proposition 3.7 in Appendix B.1.

Being ℓ_p -RIP (Definition 1.2) is a stronger property than being ℓ_p -spread. This is formalized by the following proposition, which we prove in Appendix B.2.

Proposition 3.8 (ℓ_p -RIP implies ℓ_p -spread). *Let $p \in [1, \infty]$, and let $A \in \mathbb{R}^{m \times n}$ be a (k, ε) - ℓ_p -RIP matrix. Then, $\ker(A)$ is (k, ε') - ℓ_p -spread for $\varepsilon' = \frac{1-\varepsilon}{2+\varepsilon(1+(\frac{2n}{k})^{1-\frac{1}{p}})}$.*

Definition 3.9 (Distortion of a vector). Let $1 \leq q < p$. The (ℓ_q, ℓ_p) -distortion of a vector $x \in \mathbb{R}^n \setminus \{0\}$ is

$$\Delta_{q,p}(x) := \frac{\|x\|_p \cdot n^{1/q-1/p}}{\|x\|_q}.$$

By Lemma 3.5, $1 \leq \Delta_{q,p}(x) \leq n^{1/q-1/p}$ always holds. Note that $\Delta_{q,p}(x) = 1$ iff all entries of x are equal, and $\Delta_{q,p}(x) = n^{1/q-1/p}$ iff x is supported on a single coordinate. So the distortion is a measure of the well-spreadness of the vector, with smaller values corresponding to more spread.

Definition 3.10 (Distortion of a subspace). Given $n \in \mathbb{N}$ and a subspace $X \subseteq \mathbb{R}^n$, the (ℓ_q, ℓ_p) -distortion of X is

$$\Delta_{q,p}(X) := \sup\{\Delta_{q,p}(x) \mid x \in X \setminus \{0\}\}.$$

We remark that $\Delta_p(X)$ (as defined in Section 1) is simply $\Delta_{1,p}(X)$.

Finally, we observe the following equivalence between ℓ_p -spread and (ℓ_q, ℓ_p) -distortion, which we prove in Appendix B.3.

Proposition 3.11 (Compressibility and distortion). *The following holds for all $1 \leq q < p$, $k \in \mathbb{N}$ and $x \in \mathbb{R}^n$.*

1. Let $\varepsilon > 0$. If x is (k, ε) - ℓ_p -compressible then $\Delta_{q,p}(x) \geq \frac{1}{\left(\frac{k}{n}\right)^{\frac{1}{q}-\frac{1}{p}} + \varepsilon}$.
2. The vector x is $\left(k, \frac{\left(\frac{n}{k}\right)^{\frac{1}{q}}}{\Delta_{q,p}(x)}\right)$ - ℓ_p -compressible.

In particular, if a subspace X is (k, ε) - ℓ_p -spread, then $\Delta_{q,p}(X) \leq \frac{1}{\varepsilon} \left(\frac{n}{k}\right)^{\frac{1}{q}}$ for all $1 \leq q < p$.

4 Properties of random biregular graphs

In this section, $G^* = (V_L, V_R, E^*)$ is a random graph sampled uniformly from $\mathcal{G}_{m,n,s,t}$. We state Propositions 4.1 and 4.2 and 4.2, which show that G^* satisfies certain properties with high probability.

Proposition 4.1 (Scarcity of cycles). *Suppose that $1 \leq \ell = \ell(n) \leq c \cdot \frac{\log(m)}{\log(ts)}$ for a small enough universal constant $c > 0$. Then, with high probability, the graph G^* has the following properties.*

1. There exists $v^* \in V_L$ such that $B_{G^*}(v^*, 2\ell + 1)$ does not contain a cycle.
2. The graph G^* is bicycle-free at radius $2\ell + 1$. (See Definition 3.1)

Note that the above proposition is vacuously true when $ts > m^c$.

Proposition 4.2 (Unique expansion). *For some $c > 0$, the graph G^* is a $\left(\frac{ca^2}{t^4}, \frac{2}{t}\right)$ -unique expander, with high probability.*

4.1 A negative association lemma

Our main technical tool for proving Propositions 4.1 and 4.2 is Lemma 4.4. This lemma applies in the setting where a certain subset of edges $F \subseteq E^*$ of the random graph G^* has been revealed. We fix some vertex u and sample a random (yet unrevealed) edge, originating from u to some random vertex $v \in N_{G^*}(u)$. The claim gives a natural lower bound for the probability that $\deg_F(v) = 0$, namely, that v is “new,” in the sense that it does not touch any of the edges in F .

In order to state Lemma 4.4 we introduce some notation. Given a set of edges $F \subseteq \binom{V_L \cup V_R}{2}$ and a vertex $v \in V_L \cup V_R$, we denote $N_F(v) = \{u \in V_L \cup V_R \mid \{v, u\} \in F\}$ and $\deg_F(v) = |N_F(v)|$.

Definition 4.3. A set of edges $F \subseteq \binom{V_L \cup V_R}{2}$ is called *viable* if it can be completed to an (s, t) -biregular bipartite graph, namely, if $F \subseteq E(G)$ for some $G \in \mathcal{G}_{m,n,s,t}$.

Lemma 4.4 (A Negative Association like property). *The following holds for every viable set of edges F .*

1. Let $u \in V_L$ such that $\deg_F(u) < t$. Let v be uniformly sampled from $N_{G^*}(u) \setminus N_F(u)$. Then,

$$\Pr_{G^*,v} [\deg_F(v) > 0 \mid F \subseteq E^*] \leq \frac{|\{x \in V_R \mid \deg_F(x) > 0\}|}{|V_R|}.$$

2. Let $u \in V_R$ such that $\deg_F(u) < s$. Let v be uniformly sampled from $N_{G^*}(u) \setminus N_F(u)$. Then,

$$\Pr_{G^*,v} [\deg_F(v) > 0 \mid F \subseteq E^*] \leq \frac{|\{x \in V_L \mid \deg_F(x) > 0\}|}{|V_L|} .$$

Proof. We only prove the first claim. The second claim follows by symmetry.

Let $V'_R = V_R \setminus N_F(u)$ and $V''_R = \{x \in V_R \mid \deg_F(x) = 0\}$. Fix two vertices $z \in V'_R$, $y \in V''_R$. We claim that

$$\Pr_{G^*} [y \in N_{G^*}(u) \mid F \subseteq E^*] \geq \Pr_{G^*} [z \in N_{G^*}(u) \mid F \subseteq E^*] . \quad (2)$$

This suffices to prove the Lemma, since

$$\begin{aligned} \Pr_{G^*,v} [\deg_F(v) = 0 \mid F \subseteq E^*] &= \sum_{x \in V''_R} \Pr_{G^*,v} [v = x \mid F \subseteq E^*] \\ &= \frac{1}{t - \deg_F(u)} \cdot \sum_{x \in V''_R} \Pr_{G^*} [x \in N_{G^*}(u) \mid F \subseteq E^*] \\ &\geq \frac{1}{t - \deg_F(u)} \cdot \frac{|V''_R|}{|V'_R|} \cdot \sum_{x \in V'_R} \Pr_{G^*} [x \in N_{G^*}(u) \mid F \subseteq E^*] \\ &= \frac{|V''_R|}{|V'_R|} \geq \frac{|V''_R|}{|V_R|} , \end{aligned}$$

where the first inequality follows from Eq. (2).

We turn to proving Eq. (2). For every set $W \subseteq V_L$ such that $u \in W$, and for every $1 \leq k \leq \frac{|W|}{2}$, fix some injective map

$$g_{W,k}: \{U \subseteq W \mid |U| = k \text{ and } u \in U\} \rightarrow \{U \subseteq W \mid |U| = k \text{ and } u \notin U\} .$$

Since $k \leq \frac{|W|}{2}$, the right set is at least as large as the left one, and thus such a map exists.

Let

$$\mathcal{Y} = \{G \in \mathcal{G}_{m,n,s,t} \mid F \cup \{\{u, y\}\} \subseteq E(G)\}$$

and

$$\mathcal{Z} = \{G \in \mathcal{G}_{m,n,s,t} \mid F \cup \{\{u, z\}\} \subseteq E(G)\} .$$

Since the ratio between the left and right sides of Eq. (2) is $\frac{|\mathcal{Y}|}{|\mathcal{Z}|}$, it suffices to show that $|\mathcal{Y}| \geq |\mathcal{Z}|$. We do so by proving the existence an injection $f: \mathcal{Z} \setminus \mathcal{Y} \rightarrow \mathcal{Y} \setminus \mathcal{Z}$.

We define the injection f by fixing a graph $G = (V_L, V_R, E) \in \mathcal{Z} \setminus \mathcal{Y}$ and describing its image $f(G)$. Let $S = N_G(y) \cap N_G(z)$ and $T = N_F(z) \setminus S$. Denote $U = N_G(z) \setminus (S \cup T)$, $V = N_G(y) \setminus S$, and $W = (N_G(y) \cup N_G(z)) \setminus (S \cup T)$. We can write the relevant neighbor sets as disjoint unions:

$$N_G(y) = S \sqcup V , \quad N_G(z) = S \sqcup T \sqcup U , \quad \text{and} \quad W = U \sqcup V .$$

Our assumption that $G \in \mathcal{Z} \setminus \mathcal{Y}$ yields $u \in N_G(z) \setminus N_G(y) = U \sqcup T$. Since $u \notin N_F(z) \supseteq T$, it follows that $u \in U$. Let $k = |U| = s - |S| - |T|$. Note that $|V| = s - |S| \geq k$, and so $|W| = |U| + |V| \geq 2k$. Hence, $g_{W,k}(U)$ is well-defined.

To describe the graph $f(G)$ we specify the neighbor sets, with regard to $f(G)$, of every vertex $x \in V_R$. Let

$$N_{f(G)}(x) = \begin{cases} S \sqcup T \sqcup g_{W,k}(U) & \text{if } x = z \\ S \sqcup (W \setminus g_{W,k}(U)) & \text{if } x = y \\ N_G(x) & \text{if } x \notin \{y, z\} . \end{cases}$$

Note that G and $f(G)$ agree on all edges, except for some of the edges between W and $\{y, z\}$. In particular, observe that no edge of F connects W to $\{y, z\}$. Indeed, $U \cap N_F(z)$ by definition, and $\deg_F(y) = 0$ due to our assumption that $y \in V_R''$. Thus, $F \subseteq E(f(G))$.

Note that, in both G and $f(G)$, every vertex in W is connected to exactly one of y and z . Hence, the left degrees of G and $f(G)$ are identical. Since $|g_{W,k}(U)| = |U|$, the vertices y and z also have the same degree under G and $f(G)$, implying that all the right degrees are identical. Consequently, $f(G) \in \mathcal{G}_{m,n,s,t}$. By definition of $g_{W,k}$, we have $u \in W \setminus g_{W,k}(U) \subseteq N_{f(G)}(y) \setminus N_{f(G)}(z)$, and so $f(G) \in \mathcal{Y} \setminus \mathcal{Z}$.

It remains to prove that f is an injection. Indeed, suppose that $f(G) = f(G')$. It suffices to show that $G = G'$. Clearly, G and G' agree on every edge except, perhaps, for edges between W and $\{y, z\}$. In particular, the set

$$S = N_G(y) \cap N_G(z) = N_{G'}(y) \cap N_{G'}(z) = N_{f(G)}(y) \cap N_{f(G)}(z)$$

is well-defined in terms of $f(G)$. Since $g_{W,k}$ is injective, we have

$$N_G(z) \cap W = N_{G'}(z) \cap W = g_{W,k}^{-1} \left(N_{f(G)}(z) \setminus (S \sqcup T) \right)$$

and

$$N_G(y) \cap W = N_{G'}(y) \cap W = W \setminus N_G(z) ,$$

so that $G = G'$. □

4.2 Proof of Proposition 4.1: scarcity of cycles

To prove Proposition 4.1, we show the following lemma.

Lemma 4.5. *Let $\ell \in \mathbb{N}$ and $v^* \in V_L$. Then,*

$$\Pr \left[\mathbf{B}_{G^*}(v^*, 2\ell + 1) \text{ contains a cycle} \right] \leq \frac{8\ell^2 t^{2\ell+1} s^{2\ell}}{m} . \quad (3)$$

and

$$\Pr \left[\mathbf{B}_{G^*}(v^*, 2\ell + 1) \text{ contains two (possibly intersecting) cycles} \right] \leq \frac{\ell^3 (ts)^{O(\ell)}}{m^2} . \quad (4)$$

Proposition 4.1 follows readily from Lemma 4.5. Indeed, the first property holds with high probability for some arbitrarily chosen vertex v^* due to Eq. (3), provided that c is small enough. The second property holds by Eq. (4), via a union bound over all vertices in V_L .

Proof of Lemma 4.5.

Proof of Eq. (3): Consider the following randomized algorithm.

Algorithm 1 Detect a cycle near v^*

- 1: Set $v_0 = v^*$ and set v_1 be a uniformly random neighbor of v_0 .
 - 2: **for** $i = 2, \dots, 4\ell + 1$ **do**
 - 3: Sample v_i uniformly from $N_{G^*}(v_{i-1}) \setminus \{v_{i-2}\}$
 - 4: **if** $v_i \in \{v_0, \dots, v_{i-3}\}$ **then**
 - 5: accept.
 - 6: reject.
-

Let T_i denote the event that Algorithm 1 reaches the i -th step and accepts on that step. Let $T = \bigcup_{i=1}^{4\ell+1} T_i$ denote the event that the algorithm eventually accepts.

Lemma 4.4, applied to the edge set $\{\{v_{j-1}, v_j\} \mid 1 \leq j \leq i-1\}$ and to the vertex v_{i-1} , yields

$$\Pr \left[T_i \mid \bigcap_{j=2}^{i-1} \overline{T_j} \right] \leq \begin{cases} \frac{i}{2|V_L|} & \text{if } i \text{ is even} \\ \frac{i-1}{2|V_R|} & \text{if } i \text{ is odd} \end{cases} .$$

Thus,

$$\Pr[T] \leq \sum_{i=2}^{4\ell+1} \Pr \left[T_i \mid \bigcap_{j=2}^{i-1} \overline{T_j} \right] \leq \frac{4\ell^2}{|V_L|} + \frac{4\ell^2}{|V_R|} .$$

Let C denote the event that $B_{G^*}(v^*, 2\ell + 1)$ contains a cycle. If such a cycle exists, at least one run of Algorithm 1 will detect it. Since there are at most $t^{2\ell+1} \cdot s^{2\ell}$ possible runs,

$$\Pr[T \mid C] \geq \frac{1}{t^{2\ell+1}s^{2\ell}} .$$

Therefore,

$$\frac{4\ell^2}{|V_L|} + \frac{4\ell^2}{|V_R|} \geq \Pr[T] \geq \Pr[C] \cdot \Pr[T \mid C] \geq \frac{\Pr[C]}{t^{2\ell+1}s^{2\ell}} ,$$

and so,

$$\Pr[C] \leq \ell^2 t^{2\ell+1} s^{2\ell} \cdot \left(\frac{4}{m} + \frac{4}{n} \right) \leq \frac{8\ell^2 t^{2\ell+1} s^{2\ell}}{m} .$$

Proof of Eq. (4): We prove Eq. (4) along similar lines to Eq. (3). We define another algorithm, which attempts to detect two cycles near v^* .

Algorithm 2 Detect two cycles near v^*

- 1: Use Algorithm 1 to seek a simple cycle in $B_{G^*}(v^*, 2\ell + 1)$.
 - 2: **if** a cycle u_0, \dots, u_{r-1}, u_0 ($r < 4\ell + 1$) has been found **then**
 - 3: Sample a random $0 \leq j \leq r - 1$.
 - 4: Let $w_0 = u_j$.
 - 5: Uniformly sample w_1 from $N_{G^*}(u_j) \setminus \{u_{j-1}, u_{j+1}\}$ (where $j \pm 1$ is taken mod r).
 - 6: In a similar fashion to Algorithm 1, seek a path w_0, w_1, \dots, w_q ($q \leq 4\ell + 1$), where $w_q \in \{w_0, \dots, w_{q-3}\} \cup \{u_0, \dots, u_{r-1}\}$.
 - 7: **if** such a path has been found **then**
 - 8: accept.
 - 9: **else**
 - 10: reject.
 - 11: **else**
 - 12: reject.
-

As before, we denote by T the event that the call to Algorithm 1 in Line 1 accepts. We denote by T' the event that Algorithm 2 accepts.

We bound $\Pr[T' | T]$ similarly to our bound of $\Pr[T]$. By Lemma 4.4, the probability that $w_i \in \{w_0, \dots, w_{i-3} \cup \{u_0, \dots, u_{r-1}\}$ is at most $\frac{i+r}{2|V_L|}$ (if $w_{i-1} \in V_R$) or $\frac{i+r}{2|V_R|}$ (if $w_{i-1} \in V_L$). Hence,

$$\Pr[T' | T] \leq \frac{32\ell}{|V_L|} + \frac{32\ell}{|V_R|} ,$$

and so

$$\Pr[T'] = \Pr[T' | T] \cdot \Pr[T] \leq \left(\frac{32\ell}{|V_L|} + \frac{32\ell}{|V_R|} \right) \cdot \left(\frac{4\ell^2}{|V_L|} + \frac{4\ell^2}{|V_R|} \right) \leq O\left(\frac{\ell^3}{|V_L|^2} \right) .$$

Let C' denote the event that $B_{G^*}(v^*, 2\ell + 1)$ contains two cycles. As in the proof of Eq. (3), C' implies that at least one run of Algorithm 2 accepts. Thus,

$$\Pr[T' | C'] \geq \frac{1}{(4\ell + 1) \cdot t^{4\ell+2} s^{4\ell+2}} .$$

We conclude that

$$\Pr[C'] \leq \frac{\Pr[T']}{\Pr[T' | C']} \leq \frac{\ell^3 (ts)^{O(\ell)}}{m^2} . \quad \square$$

4.3 Proof of Proposition 4.2: unique expansion

To prove Proposition 4.2, we show the following lemma.

Lemma 4.6. *Let $S \subseteq V_L$ and denote $k = |S|$. Then*

$$\Pr_{G^*} [S \text{ has at most } (t-2)|S| \text{ unique neighbors}] \leq \left(\frac{et^2k}{2m} \right)^{2k} .$$

Our proof of Lemma 4.6 is an adaptation of the proof of [Vad12, Theorem 4.4].

Proposition 4.2 then follows from Lemma 4.6 by the union bound. Let $\gamma = \frac{c\alpha^2}{t^4}$. Then,

$$\begin{aligned} & \Pr_{G^*} \left[G^* \text{ is not a } \left(\gamma, \frac{2}{t} \right)\text{-unique expander} \right] \\ & \leq \sum_{k=1}^{\gamma n} \sum_{\substack{S \subseteq [n] \\ |S|=k}} \Pr_{G^*} [S \text{ has less than } (t-2)|S| \text{ unique neighbors}] \\ & \leq \sum_{k=1}^{\gamma n} \binom{n}{k} \cdot \left(\frac{et^2k}{2m} \right)^{2k} \leq \sum_{k=1}^{\gamma n} \left(\frac{ne}{k} \right)^k \cdot \left(\frac{et^2k}{2m} \right)^{2k} = \sum_{k=1}^{\gamma n} \left(\frac{e^3 t^4 k}{\alpha^2 n} \right)^k. \end{aligned} \quad (5)$$

Taking $\gamma \leq \frac{\alpha^2}{2e^3 t^4}$, we have $\frac{e^3 t^4 k}{\alpha^2 n} < \frac{1}{2}$ whenever $k \leq \gamma n$. Hence, the sum on the right-hand side of Eq. (5) is dominated by its first term, namely, it is $O\left(\frac{t^4}{\alpha^2 n}\right)$.

Proof of Lemma 4.6. Write $S = \{v_1, \dots, v_k\}$. Consider a process in which the tk edges touching $|S|$ are revealed in tk steps, where the first t steps reveal the neighbors of v_1 , the next t steps reveal those of v_2 , and so on.

Let $V_i \subseteq V_R$ ($0 \leq i \leq tk$) denote the set of neighbors of S revealed by the i -th step. Note that $V_i = V_{i-1}$ if and only if the neighbor revealed in the i -th step has already been revealed in a previous step. Otherwise, $|V_i| = |V_{i-1}| + 1$. By Lemma 4.4, the probability of the event $V_i = V_{i-1}$, conditioned on all previous steps, is at most $\frac{|V_{i-1}|}{m} \leq \frac{tk}{m}$. Thus,

$$\begin{aligned} & \Pr [S \text{ has at most } (t-2)k \text{ unique neighbors}] = \Pr [\text{at least } 2k \text{ steps } i \text{ have } V_i = V_{i-1}] \\ & \leq \binom{tk}{2k} \left(\frac{tk}{m} \right)^{2k} \leq \left(\frac{etk}{2k} \right)^{2k} \cdot \left(\frac{tk}{m} \right)^{2k} = \left(\frac{et^2k}{2m} \right)^{2k}. \quad \square \end{aligned}$$

5 Limitations of ℓ_2 -spread

Our main goal in this section is to prove Theorem 1 given Theorem 2. We recall Theorem 1.

Theorem 1 (Poor ℓ_2 -spread of $\ker(A)$). *With high probability over A , there exists an $(m^c, \frac{n^{-\Omega(\log(1/\alpha)/\log s)}}{1-\sqrt{\alpha}})$ - ℓ_2 -compressible vector $x \in \ker(A)$, where $c < 1$ is an absolute constant. In particular,*

$$\Delta_2(\ker(A)) \geq (1 - \sqrt{\alpha}) \cdot n^{\Omega(\log(1/\alpha)/\log s)}.$$

Moreover, there is a poly(n)-time algorithm that, on input A , outputs such an x .

The proof of Theorem 1 relies on the following two lemmas.

Lemma 5.1. *Let $A \in \mathcal{M}_{m,n,s,t}$. Suppose that G_A satisfies Property 1 of Proposition 4.1 with regard to some $\ell \in \mathbb{N}$. Then, there is a vector $x \in \mathbb{R}^n \setminus \{0^n\}$ with $|\text{supp}(x)| \leq 1 + 2t(t-1)^{\ell-1}(s-1)^\ell$ and*

$$\frac{\|Ax\|_p^p}{\|x\|_p^p} \leq t(t-1)^\ell (s-1)^{(1-p)\ell}$$

for all $p \geq 1$. In particular,

$$\frac{\|Ax\|_2^2}{\|x\|_2^2} \leq \frac{t(t-1)^\ell}{(s-1)^\ell}. \quad (6)$$

Furthermore, x can be computed in polynomial time given A .

Lemma 5.2. *Let $A \in \mathbb{R}^{m \times n}$ be a matrix, and let $x \in \mathbb{R}^n$ with $\|x\|_2 = 1$. Then, $\ker(A)$ contains a $(|\text{supp}(x)|, \frac{\varepsilon}{1-\varepsilon})$ - ℓ_2 -compressible vector y , where $\varepsilon = \frac{\|Ax\|_2}{\sigma_{\min}(A)}$. Furthermore, y can be computed in polynomial time given A and x .*

Lemmas 5.1 and 5.2 are proven later in this section. We first use them to prove Theorem 1.

Proof of Theorem 1 given Theorem 2 and Lemmas 5.1 and 5.2. Let c denote the constant from Proposition 4.1. Let $\ell = \left\lceil c' \frac{\log(m)}{\log(ts)} - 1 \right\rceil$, where $c' = \min(c, \frac{1}{2})$.

With high probability, we sample a matrix A such that the events in Proposition 4.1 and Theorem 2 all occur. For the remainder of the proof, we shall view A as fixed, assuming that these events hold.

In particular, our assumption implies the hypothesis of Lemma 5.1. Hence, there is a vector $x \in \mathbb{R}^n$, computable in polynomial time, which satisfies Eq. (6). Applying Lemma 5.2 to $\frac{x}{\|x\|_2}$ yields a $(2(ts)^\ell, \frac{\varepsilon}{1-\varepsilon})$ - ℓ_2 -compressible vector in $y \in \ker(A)$, where $\varepsilon = \frac{\|Ax\|_2}{\sigma_{\min}(A)} \leq \frac{\sqrt{t} \cdot \alpha^{\frac{\ell}{2}}}{\sigma_{\min}(A)}$ due to Eq. (6). As the event in Theorem 2 occurs,

$$\sigma_{\min}(A) \geq \sqrt{s-1} - (1+o(1)) \cdot \sqrt{t-1} \geq \sqrt{s} - (1+o(1)) \sqrt{t} = \sqrt{s} \cdot \left(1 - \sqrt{\alpha} - o(1)\right),$$

and so

$$\varepsilon \leq \frac{\alpha^{\frac{\ell+1}{2}}}{1 - \sqrt{\alpha} - o(1)} \leq \frac{n^{-\Omega\left(\frac{\log(1/\alpha)}{\log s}\right)}}{1 - \sqrt{\alpha}}.$$

Also, $(ts)^\ell \leq m^{c'}$, so y is $\left(m^{c'}, \frac{n^{-\Omega\left(\frac{\log(1/\alpha)}{\log s}\right)}}{1 - \sqrt{\alpha}}\right)$ - ℓ_2 -compressible. By Proposition 3.11,

$$\Delta_{1,2}(\ker A) \geq \frac{1}{n^{-\Omega(1)} + \frac{n^{-\Omega\left(\frac{\log(1/\alpha)}{\log s}\right)}}{1 - \sqrt{\alpha}}} \geq (1 - \sqrt{\alpha}) \cdot n^{\Omega\left(\frac{\log(1/\alpha)}{\log s}\right)},$$

where we use that $n^{-\Omega(1)} \leq n^{-\Omega\left(\frac{\log(1/\alpha)}{\log s}\right)}$ as $t = \alpha s \geq 3$. Finally, the last part of the theorem follows since y can be computed from A, x in polynomial time. \square

In addition, Lemma 5.1 suffices to prove Theorem 7.

Theorem 7 (Partial converse to Corollary 5). *Let $p \in [1, 2)$, $\varepsilon > 0$. If $s - 1 \leq \left(\frac{1}{(1+\varepsilon)\alpha}\right)^{\frac{1}{2-p}}$, then with high probability over A , there exists an n^c -sparse vector $x \in \mathbb{R}^n \setminus \{0^n\}$ such that*

$$\frac{\|Ax\|_p}{\|x\|_p} \leq t^{\frac{1}{p}} \cdot m^{-\Omega\left(\frac{\varepsilon}{\log s}\right)}.$$

Proof of Theorem 7 given Lemma 5.1. Let $\ell = \left\lfloor c \frac{\log m}{\log(ts)} \right\rfloor$, where c is as in Proposition 4.1. The graph G_A satisfies Property 1 of Proposition 4.1 with high probability, and so we proceed assuming that this is the case. Lemma 5.1 now yields a vector $x \in \mathbb{R}^n \setminus \{0\}$ with

$$\begin{aligned} \frac{\|Ax\|_p}{\|x\|_p} &\leq \left(t(t-1)^\ell (s-1)^{(1-p)\ell} \right)^{\frac{1}{p}} = t^{\frac{1}{p}} \left(\frac{t-1}{s-1} \cdot (s-1)^{2-p} \right)^{\frac{\ell}{p}} \leq t^{\frac{1}{p}} (\alpha \cdot (s-1)^{2-p})^{\frac{\ell}{p}} \\ &\leq t^{\frac{1}{p}} (1+\varepsilon)^{-\frac{\ell}{p}} \leq t^{\frac{1}{p}} \cdot m^{-\Omega\left(\frac{\varepsilon}{\log s}\right)}. \end{aligned} \quad \square$$

5.1 Proof of Lemma 5.1: constructing x

Write $G_A = (V_L, V_R, E)$ as in Section 3.2. By our assumption of Proposition 4.1 – Property 1, there is a vertex $v^* \in V_L$ such that $\mathbf{B}_G(v^*, 2\ell + 1)$ is isomorphic to a (t, s) -biregular tree of depth $2\ell + 1$, rooted at v^* . Recall that a (t, s) -biregular tree is a tree in which a non-leaf vertex has degree t (resp. s) if it belongs to V_L (resp. V_R).

We define the vector $x \in \mathbb{R}^n$ as follows. Given $v \in V_L$, if $v \notin \mathbf{B}_G(v^*, 2\ell + 1)$, let $x_v = 0$. Otherwise, we have $v \in V_L \cap \mathbf{B}_G(v^*, 2\ell + 1)$. Let $v^* = u_0, r_1, u_1, r_2, u_2, \dots, r_k, u_k = v$ ($0 \leq k \leq \ell$) denote the unique simple path in $\mathbf{B}_G(v^*, 2\ell + 1)$ from v^* to v . Let

$$x_v = \prod_{i=1}^k \left(\text{sign}(u_{i-1}, r_i) \cdot \text{sign}(u_i, r_i) \cdot \frac{-1}{s-1} \right).$$

Clearly, x is computable in polynomial time given A .

We compute $(Ax)_r = \sum_{v \in N_G(r)} \text{sign}(v, r) \cdot x_v$ for every $r \in V_R$ by considering three cases:

- If r is an inner vertex in $\mathbf{B}_G(v^*, 2\ell + 1)$, denote its parent in the tree by $v' \in V_L$ and its children by $v_1, \dots, v_{s-1} \in V_L$. Note that

$$x_{v_i} = \frac{-1}{s-1} \cdot x_{v'} \cdot \text{sign}(v', r) \cdot \text{sign}(v_i, r)$$

for $1 \leq i \leq s-1$. Hence,

$$\begin{aligned} (Ax)_r &= \text{sign}(v', r) \cdot x_{v'} + \sum_{i=1}^{s-1} \text{sign}(v_i, r) \cdot x_{v_i} \\ &= \text{sign}(v', r) \cdot x_{v'} + \sum_{i=1}^{s-1} \frac{-1}{s-1} \cdot x_{v'} \cdot \text{sign}(v_i, r)^2 \cdot \text{sign}(v', r) = 0, \end{aligned}$$

where we used the fact that $\text{sign}(v_i, r) \in \{1, -1\}$, and so $\text{sign}(v_i, r)^2 = 1$.

- If r lies outside of $\mathbf{B}_G(v^*, 2\ell + 1)$ then $(Ax)_r = 0$ since every neighbor $v \in V_L$ of r has $x_v = 0$.
- Finally, if r is a leaf of the tree $\mathbf{B}_G(v^*, 2\ell + 1)$, then r has a unique neighbor $v \in V_R$ such that $x_v \neq 0$ (namely, v is the parent of r in the tree). This vertex v is distance 2ℓ from v^* , so $|x_v| = (s-1)^{-\ell}$. Thus, $|(Ax)_r| = (s-1)^{-\ell}$.

Since the tree has $t \cdot (s - 1)^\ell \cdot (t - 1)^\ell$ leaves,

$$\|Ax\|_p^p = t \cdot (s - 1)^\ell \cdot (t - 1)^\ell \cdot (s - 1)^{-p\ell} = t \cdot (s - 1)^{(1-p)\ell} \cdot (t - 1)^\ell.$$

Clearly, since $x_{v^*} = 1$, we have $\|x\|_p^p \geq 1$, and the lemma follows.

5.2 Proof of Lemma 5.2: rounding x to $y \in \ker(A)$

Let Π be the orthogonal projection onto $\ker(A)$, and let Π^\perp be the projection onto the subspace orthogonal to $\ker(A)$. Let $y = \Pi x \in \ker(A)$ be the vector such that $\|x - y\|_2 = \min_{z \in \ker(A)} \|x - z\|_2$. By using the SVD of A , we observe that $\|Ax\|_2 \geq \sigma_{\min}(A) \cdot \|\Pi^\perp x\|_2 = \sigma_{\min}(A) \cdot \|x - y\|_2$, as $\Pi^\perp x = x - \Pi x = x - y$. Hence, $\|x - y\|_2 \leq \frac{\|Ax\|_2}{\sigma_{\min}(A)} =: \varepsilon$. Therefore,

$$\|x - y\|_2 \leq \frac{\varepsilon}{\|y\|_2} \cdot \|y\|_2 \leq \frac{\varepsilon}{\|x\|_2 - \|x - y\|_2} \cdot \|y\|_2 \leq \frac{\varepsilon}{\|x\|_2 - \varepsilon} \cdot \|y\|_2.$$

We conclude that y is $\left(|\text{supp}(x)|, \frac{\varepsilon}{\|x\|_2 - \varepsilon}\right)$ - ℓ_2 -compressible.

6 Positive results for ℓ_p -RIP and ℓ_p -spread for $p \in [1, 2)$

In this section, we prove Theorem 4 and Corollaries 5 and 6. We restate Theorem 4 below.

Theorem 4 (ℓ_p -RIP of expander graphs). *Let G be a bipartite t -left-regular (γ, μ) -unique expander with maximum right degree s_{\max} , and let B be any signed adjacency matrix of G . Let $0 < \varepsilon \leq 1$ and $1 \leq p < 2$ such that $\varepsilon^2 \geq 9\mu s_{\max}^{p-1}$. Then, B is $(\gamma n, \varepsilon)$ - ℓ_p -RIP, i.e., for every γn -sparse $x \in \mathbb{R}^n$,*

$$t^{\frac{1}{p}}(1 - \varepsilon) \|x\|_p \leq \|Bx\|_p \leq t^{\frac{1}{p}}(1 + \varepsilon) \|x\|_p.$$

The proof of Theorem 4 relies on the following technical lemma, which we prove in Section 6.3.

Lemma 6.1. *Let $A \in \{0, 1, -1\}^{m \times n}$ such that G_A is a t -left-regular, (γ, μ) -unique expander. Let s_{\max} denote the maximum degree of a right vertex of G_A . Then for any $p \geq 1$, $\delta_1 > 0$, $\delta_2 \in (0, 1)$, and γn -sparse $x \in \mathbb{R}^n$,*

$$\|Ax\|_p^p \geq \left(\frac{t(1 - \mu)}{(1 + \delta_1)^{p-1}} - \frac{\mu t}{\delta_1^{p-1}} (s_{\max} - 1)^{p-1} \right) \|x\|_p^p \quad (7)$$

$$\|Ax\|_p^p \leq \left(\frac{t}{(1 - \delta_2)^{p-1}} + \frac{\mu t}{\delta_2^{p-1}} (s_{\max} - 1)^{p-1} \right) \|x\|_p^p. \quad (8)$$

Note that for $p \geq 2$, Eq. (7) is trivial, as the right hand side is always ≤ 0 .

We now proceed with the proofs as follows. Below, we prove Theorem 4 from Lemma 6.1. We then prove Corollary 5 in Section 6.1, Corollary 6 in Section 6.2, and finally Lemma 6.1 in Section 6.3.

Proof of Theorem 4. Suppose without loss of generality that $\|x\|_p = 1$. Take $\delta_1 = \delta_2 = \frac{\varepsilon}{3}$. By Eq. (7),

$$\|Ax\|_p^p \geq \frac{t(1-\mu)}{(1+\delta_1)^{p-1}} - \frac{\mu t}{\delta_1^{p-1}}(s_{\max}-1)^{p-1} .$$

Our lower bound on ε implies, in particular, that $\mu \leq \frac{\varepsilon}{3}$. Thus, by our choice of δ_1 ,

$$\frac{t(1-\mu)}{(1+\delta_1)^{p-1}} \geq \frac{1-\frac{\varepsilon}{3}}{(1+\frac{\varepsilon}{3})^{p-1}} \cdot t \geq \frac{1-\frac{\varepsilon}{3}}{1+\frac{\varepsilon}{3}} \cdot t \geq \left(1 - \frac{2\varepsilon}{3}\right) t .$$

We also have

$$\frac{\mu t}{\delta_1^{p-1}}(s_{\max}-1)^{p-1} \leq \frac{\varepsilon^2 t}{9\delta_1} = \frac{\varepsilon t}{3} .$$

We conclude that

$$\|Ax\|_p^p \geq (1-\varepsilon)t .$$

We similarly obtain an upper bound on $\|Ax\|_p^p$. By Eq. (8),

$$\|Ax\|_p^p \leq \frac{t}{(1-\delta_2)^{p-1}} + \frac{\mu t}{\delta_2^{p-1}}(s_{\max}-1)^{p-1} .$$

Now,

$$\frac{t}{(1-\delta_2)^{p-1}} \leq \frac{t}{1-\frac{\varepsilon}{3}} \leq \left(1 + \frac{2\varepsilon}{3}\right) t$$

since $\varepsilon < \frac{1}{2}$. Also,

$$\frac{\mu t}{\delta_2^{p-1}}(s_{\max}-1)^{p-1} = \frac{\mu t}{\delta_1^{p-1}}(s_{\max}-1)^{p-1} \leq \frac{\varepsilon t}{3} ,$$

and so,

$$\|Ax\|_p^p \leq (1+\varepsilon)t .$$

Finally, we observe that $(1+\varepsilon)^{1/p} \leq 1+\varepsilon$ and $(1-\varepsilon)^{1/p} \geq 1-\varepsilon$, which finishes the proof. \square

6.1 Proof of Corollary 5: ℓ_p -RIP and ℓ_p -spread of sparse random matrices

We now prove Corollary 5, which we restate below.

Corollary 5 (Good ℓ_p -RIP and ℓ_p -spread of A). Fix $p \in [1, 2)$, $0 < \varepsilon < \frac{1}{2}$, and suppose that $s \geq \left(\frac{18}{\alpha\varepsilon^2}\right)^{\frac{1}{2-p}}$.

Then, with high probability over A , the matrix A is $(\Omega(\gamma n), \varepsilon)$ - ℓ_p -RIP for $\gamma = \frac{\alpha^2}{t^4}$: for every $\Omega(\gamma n)$ -sparse $x \in \mathbb{R}^n$, it holds that

$$t^{\frac{1}{p}}(1-\varepsilon)\|x\|_p \leq \|Ax\|_p \leq t^{\frac{1}{p}}(1+\varepsilon)\|x\|_p .$$

In particular, the subspace $\ker(A)$ is $(\Omega(\gamma n), \Omega\left(\gamma^{1-\frac{1}{p}}\right))$ - ℓ_p -spread and $\Delta_p(\ker(A)) \leq O\left(1/\gamma^{2-\frac{2}{p}}\right)$.

Proof. Fix ε , and let s satisfy $s \geq \left(\frac{18}{\alpha\varepsilon^2}\right)^{\frac{1}{2-p}}$. Note that this implies $\varepsilon^2 \geq \frac{18}{t} \cdot s^{p-1}$, as $t = \alpha s$.

By Proposition 4.2, with high probability it holds that G_A is a (γ, μ) -unique expander with $\gamma = \Omega(\alpha^2/t^4)$ and $\mu = \frac{2}{t}$. Hence, we have $\varepsilon^2 \geq 9\mu s^{p-1} = \frac{18}{t} \cdot s^{p-1}$, and $s_{\max} = s$ as G_A is s -right-regular. By Theorem 4, we thus have that A is $(\gamma n, \varepsilon)$ - ℓ_p -RIP. Applying Proposition 3.8, we conclude that $\ker(A)$ is $(\gamma n, \varepsilon')$ - ℓ_p -spread for $\varepsilon' = \frac{1-\varepsilon}{2+\varepsilon(1+(\frac{2}{\gamma})^{1-\frac{1}{p}})}$. As $\varepsilon < \frac{1}{2}$, it follows that $\varepsilon' = \Omega(\gamma^{1-\frac{1}{p}}) = \Omega((\frac{\alpha^2}{t^4})^{1-\frac{1}{p}})$.

Then, applying Proposition 3.11, we conclude that $\Delta_p(\ker(A)) \leq \frac{1}{\gamma \varepsilon'} = O(1/\gamma^{2-\frac{1}{p}}) = O((\frac{t^4}{\alpha^2})^{2-\frac{1}{p}})$. \square

6.2 Proof of Corollary 6: explicit construction of ℓ_p -RIP matrices

In this subsection, we prove Corollary 6, which we restate below.

Corollary 6 (Explicit construction of ℓ_p -RIP matrices). *Let $0 < \varepsilon < \frac{1}{2}$, $\alpha \in (0, 1)$, and let $n \in \mathbb{N}$ be sufficiently large. For some universal constant $1 < p_0 < 2$, there exists a deterministic algorithm which, given $p \in [1, p_0)$, ε , α and n , outputs in time $\text{poly}(n/\delta) + 2^{O(1/\delta)}$ a matrix $B \in \{0, 1\}^{m \times n}$, for some $m \leq \alpha n$, such that B is $(\gamma n, \varepsilon)$ - ℓ_p -RIP, for some $\delta, \gamma = \text{poly}(\varepsilon, \alpha)^{\frac{1}{p_0-p}}$. In particular, $\ker(B)$ is $(\gamma n, \gamma^{1-\frac{1}{p}})$ - ℓ_p -spread and $\Delta_p(\ker(B)) \leq 1/\gamma^{2-\frac{2}{p}}$.*

Proof. We prove Corollary 6 by combining Theorem 4 with the explicit constructions of vertex expanders due to [CRVW02]. We note that these expanders do not necessarily have bounded right degree, which is necessary to use Theorem 4. Because of this, we first give a simple preprocessing algorithm to convert an expander to a new graph with similar expansion and bounded right degree.

Lemma 6.2. *Let $G = (V_L, V_R, E)$ be a bipartite t -left regular (γ, μ) -unique expander where $n := |V_L| \geq |V_R|$. Then, there exists a bipartite t -left regular (γ, μ) -unique expander $G' = (V'_L, V'_R, E')$, with $V'_L = V_L$ and $|V'_R| \leq 3|V_R|$, such that every vertex in V'_R has degree at most $\frac{tn}{|V'_R|}$. Furthermore, G' can be computed from G in $\text{poly}(|V_L|)$ time.*

Proof. Write $\beta = \frac{|V_R|}{|V_L|} \leq 1$. We modify G to create the new graph G' via the following algorithm: While the graph contains a right vertex r with $\deg(r) > \frac{t}{\beta}$, create a new vertex r' . Let D be some arbitrary subset of $N(r)$ such that $|D| = \lfloor \frac{t}{\beta} \rfloor$. Remove all edges between r and D , and instead create edges between r' and every vertex in D .

Clearly, when this process terminates, the maximum right degree of the graph is at most $\frac{t}{\beta}$, and the left-degrees remain as before. Also, observe that the neighbor set of any set of left vertices can only increase in size, so our modification to the graph does not hurt its unique expansion.

It is left to show that we did not add too many right vertices. Indeed, the number of new right vertices in the graph is equal to the number of iterations of our algorithm. Let J_i denote the number of edges in graph that touch a right vertex with degree larger than $\frac{t}{\beta}$, after the i -th iteration. Clearly, $J_0 \leq |E| = nt$, and $J_{i+1} \leq J_i - \lfloor \frac{t}{\beta} \rfloor$. Hence, the number of iterations is at most $J_0 / \lfloor \frac{t}{\beta} \rfloor \leq J_0 \cdot \frac{2\beta}{t} \leq 2|V_R|$. \square

We now turn to the proof of Corollary 6. Let c be the constant in Theorem 3.4, and let $p \geq 1$ satisfy $p < p_0 := 1 + \frac{1}{c}$. Let $\varepsilon \in (0, \frac{1}{2})$, and let $\alpha \in (0, 1)$. Let $\beta = \frac{\alpha}{3}$. We shall assume without loss of generality that β is an inverse power of 2, as otherwise we can simply decrease α until this holds, and this will only lose a factor of 2 in α .

Let $\mu = (\alpha^{(1+c)(p-1)} \varepsilon^2 C)^{\frac{1}{1-c(p-1)}}$, where C is a universal constant. Note that $1 - c(p-1) > 0$ as $p < 1 + \frac{1}{c}$.

Now, let n be sufficiently large. We shall assume that n is a power of 2 without loss of generality, as otherwise we can simply increase n until this holds. By Theorem 3.4, we can construct in $\text{poly}(n)$ -time (as β, μ are constants) a bipartite graph $G = (V_L, V_R, E)$ with $|V_L| = n$, $|V_R| = \beta n$, such that G is a t -left-regular $(\Omega(\frac{\mu\beta n}{t}), \mu)$ -unique expander and $t = O\left(\left(\frac{1}{\beta\mu}\right)^c\right)$. Applying Lemma 6.2, we thus construct in $\text{poly}(n)$ -time a bipartite graph $G' = (V'_L, V'_R, E')$ with $|V'_L| = n$, $|V'_R| = 3\beta n = \alpha n$, and G' is a t -left-regular $(\Omega(\frac{\mu\alpha n}{t}), \mu)$ -unique expander with max right degree $s_{\max} \leq \frac{tn}{\beta n} = \frac{3t}{\alpha}$.

We now observe that $\varepsilon^2 \geq 9\mu s_{\max}^{p-1}$. Indeed, this is because

$$9\mu s_{\max}^{p-1} = 9\mu \cdot \left(\frac{3t}{\alpha}\right)^{p-1} \leq O(1) \cdot \mu \cdot \left(\frac{1}{\alpha(\alpha\mu)^c}\right)^{p-1} = O(1) \cdot \frac{1}{\alpha^{(1+c)(p-1)}} \cdot \mu^{1-c(p-1)} \leq \varepsilon^2,$$

by our choice of μ . Thus, by Theorem 4, we conclude that the adjacency matrix $B \in \{0, 1\}^{m \times n}$ of G' , defined by $B_{r,u} = 1$ if $(u, r) \in E'$ and 0 otherwise, is $(\Omega(\gamma n), \varepsilon)$ - ℓ_p -RIP, where $\gamma = \frac{\mu\alpha}{t}$.

By Proposition 3.8, we thus have that $\ker(B)$ is $(\gamma n, \varepsilon')$ - ℓ_p -spread, where $\varepsilon' = \Omega(\gamma^{1-\frac{1}{p}})$, and by Proposition 3.11, we have $\Delta_p(\ker(B)) \leq O(\frac{1}{\varepsilon'\gamma}) = O(1/\gamma^{2-\frac{2}{p}})$.

We finish the proof by simply observing that $\frac{1}{1-c(p-1)} = \frac{p_0-1}{p_0-p} = O(\frac{1}{p_0-p})$, and so $\gamma = \text{poly}(\mu, \alpha) = \text{poly}(\varepsilon, \alpha)^{\frac{1}{p_0-p}}$. \square

6.3 Proof of Lemma 6.1: ℓ_p -RIP from unique expansion

Let $x \in \mathbb{R}^n$ be γn -sparse, and let $S = \text{supp}(x) \subseteq V_L$. Let $E(S, N(S)) := \{\{v, r\} \in E : v \in S, r \in N(S)\}$. The following claim asserts the existence of a certain many-to-one matching between S and $N(S)$.

Claim 6.3. There exists a set of edges $M \subseteq E(S, N(S))$ with the following properties:

1. Every $r \in N(S)$ touches exactly one edge in M .
2. Every $v \in S$ touches at least $t(1 - \mu)$ edges in M .

Proof. We construct M via an iterative algorithm. The algorithm records a set of *processed* left vertices $P \subseteq S \subseteq V_L$. It also maintains the edge set M . The algorithm is initialized with $P := \emptyset$ and $M := \emptyset$.

There are $|S|$ iterations. On each iteration, the algorithm picks a vertex $v \in S \setminus P$ such that the set $U_v := N(v) \setminus (N(S \setminus P))$ of unique neighbors of v has size at least $t(1 - \mu)$. Note that by unique expansion of the set $S \setminus P$, there is always such a vertex. For each $r \in U_v$, the algorithm adds the edge $\{v, r\}$ to M . The vertex v is added to P .

We turn to analyzing this algorithm. It is straightforward to observe that when the algorithm terminates we have $P = S$, and every vertex $v \in S$ touches at least $t(1 - \mu)$ edges of M . We claim that every vertex $r \in N(S)$ touches exactly one edge of M . Indeed, let $h = |N(r) \cap S|$ denote the number of neighbors of r in S . Let u_1, \dots, u_h be these neighbors, ordered so that u_1 is the first neighbor of r added to P , u_2 is the second one, etc. We observe that for $i < h$, the edge $\{u_i, r\}$ cannot be in M . Indeed, when u_i is added to P , u_h is not in P , and thus r is not a unique neighbor of u_i . Next, we observe that when u_h is added to P , the edge $\{u_h, r\}$ is added to M . This is because all $u_i \in P$

for all $i < h$, and so at this iteration of the algorithm r has only one neighbor not in P , namely u_h . This finishes the proof. \square

We are now ready to establish the lower and upper bounds on $\|Ax\|_p$ claimed in Eqs. (7) and (8). We begin with Eq. (7).

Lower bound on $\|Ax\|_p$. Let $\kappa_1 = \frac{1}{\delta_1} > 0$. Let M be the set of edges obtained from Claim 6.3. For each $r \in N(S)$, let v_r denote the unique vertex in S such that $\{v_r, r\} \in M$, and let $W_r = (N(r) \cap S) \setminus \{v_r\}$ denote the rest of the neighbors of r in S . Define

$$a_r = |(Ax)_r|^p + \kappa_1^{p-1} (s_{\max} - 1)^{p-1} \sum_{u \in W_r} |x_u|^p .$$

Now note that

$$\begin{aligned} \sum_{r \in N(S)} a_r &= \|Ax\|_p^p + \kappa_1^{p-1} (s_{\max} - 1)^{p-1} \sum_{r \in N(S)} \sum_{u \in W_r} |x_u|^p \\ &\leq \|Ax\|_p^p + \kappa_1^{p-1} (s_{\max} - 1)^{p-1} \mu t \sum_{u \in S} |x_u|^p \\ &= \|Ax\|_p^p + \kappa_1^{p-1} (s_{\max} - 1)^{p-1} \mu t \|x\|_p^p , \end{aligned} \tag{9}$$

where the inequality is due to each $u \in S$ belonging to at most μt sets W_r ($r \in N(S)$). Indeed, the number of such sets for a given u is $t - |\{r \in N(S) \mid u = v_r\}| \leq t - (1 - \mu)t = \mu t$.

Next, for any fixed $r \in N(S)$, we claim that

$$a_r \geq \frac{1}{\left(1 + \frac{1}{\kappa_1}\right)^{p-1}} |x_{v_r}|^p . \tag{10}$$

Denote $z = \sum_{u \in W_r} |x_u|$. Note that $|W_r| = |(N(r) \cap S) \setminus \{v_r\}| \leq \deg(r) - 1 \leq s_{\max} - 1$, so Hölder's inequality (Lemma 3.5) yields

$$a_r = |(Ax)_r|^p + \kappa_1^{p-1} (s_{\max} - 1)^{p-1} \sum_{u \in W_r} |x_u|^p \geq |(Ax)_r|^p + \kappa_1^{p-1} z^p .$$

To deduce Eq. (10) from the above, we distinguish two cases:

- If $z \geq |x_{v_r}|$, we can bound

$$a_r \geq \kappa_1^{p-1} z^p \geq \kappa_1^{p-1} |x_{v_r}|^p \geq \frac{1}{\left(1 + \frac{1}{\kappa_1}\right)^{p-1}} |x_{v_r}|^p .$$

- If $z \leq |x_{v_r}|$, we let $z = \beta |x_{v_r}|$ for some $\beta \in [0, 1]$. We then have $|(Ax)_r|^p \geq (|x_{v_r}| - z)^p \geq (1 - \beta)^p |x_{v_r}|^p$, and conclude that

$$a_r \geq |(Ax)_r|^p + \kappa_1^{p-1} z^p \geq (1 - \beta)^p |x_{v_r}|^p + \kappa_1^{p-1} \beta^p |x_{v_r}|^p \geq \frac{1}{\left(1 + \frac{1}{\kappa_1}\right)^{p-1}} |x_{v_r}|^p ,$$

where we use that $\min_{\beta \in [0,1]} (1 - \beta)^p + \kappa_1^{p-1} \beta^p = \left(1 + \frac{1}{\kappa_1}\right)^{1-p}$ for $\kappa_1 > 0$.

Now, Eq. (10) yields

$$\begin{aligned} \sum_{r \in N(S)} a_r &\geq \frac{1}{\left(1 + \frac{1}{\kappa_1}\right)^{p-1}} \sum_{v \in S} (|x_v|^p \cdot |\{r \in N(S) \mid v = v_r\}|) \\ &\geq \frac{1}{\left(1 + \frac{1}{\kappa_1}\right)^{p-1}} \sum_{v \in S} |x_v|^p \cdot (1 - \mu)t = \frac{1}{\left(1 + \frac{1}{\kappa_1}\right)^{p-1}} (1 - \mu)t \|x\|_p^p . \end{aligned}$$

Eq. (7) follows from the above and Eq. (9), as $\kappa_1 = \frac{1}{\delta_1}$. □

Upper bound on $\|Ax\|_p$. We now prove Eq. (8) in a similar way. Let $\kappa_2 = \frac{1}{\delta_2} > 1$, and let

$$b_r = |(Ax)_r|^p - \kappa_2^{p-1} (s_{\max} - 1)^{p-1} \cdot \sum_{u \in W_r} |x_u|^p .$$

Note that

$$\begin{aligned} \sum_{r \in N(S)} b_r &\geq \|Ax\|_p^p - \kappa_2^{p-1} (s_{\max} - 1)^{p-1} \sum_{r \in N(S)} \sum_{u \in W_r} |x_u|^p \\ &\geq \|Ax\|_p^p - \mu t \kappa_2^{p-1} (s_{\max} - 1)^{p-1} \sum_{u \in S} |x_u|^p \\ &= \|Ax\|_p^p - \mu t \kappa_2^{p-1} (s_{\max} - 1)^{p-1} \|x\|_p^p . \end{aligned} \tag{11}$$

We next seek an upper bound on b_r . Fix some $r \in N(S)$, and write $\sum_{u \in W_r} |x_u| = \beta |x_{v_r}|$ for some $\beta \geq 0$. Note that

$$|(Ax)_r|^p \leq |x_{v_r}|^p + \sum_{u \in W_r} |x_u|^p \leq (1 + \beta)^p |x_{v_r}|^p .$$

Also, by Hölder's inequality (Lemma 3.5),

$$\sum_{u \in W_r} |x_u|^p \geq |W_r|^{1-p} \left(\sum_{u \in W_r} |x_u| \right)^p = |W_r|^{1-p} \cdot (\beta |x_{v_r}|)^p \geq (s_{\max} - 1)^{1-p} \cdot (\beta |x_{v_r}|)^p .$$

It follows that

$$b_r \leq \left((1 + \beta)^p - \kappa_2^{p-1} \cdot \beta^p \right) |x_{v_r}|^p \leq \frac{1}{\left(1 - \frac{1}{\kappa_2}\right)^{p-1}} |x_{v_r}|^p ,$$

since $\sup_{\beta \geq 0} \left\{ (1 + \beta)^p - \kappa_2^{p-1} \cdot \beta^p \right\} = \frac{1}{\left(1 - \frac{1}{\kappa_2}\right)^{p-1}}$ for $\kappa_2 > 1$.

By Eq. (11) and the above,

$$\|Ax\|_p^p \leq \sum_{r \in N(S)} b_r + \kappa_2^{p-1} (s_{\max} - 1)^{p-1} \mu t \|x\|_p^p \leq \left(\frac{t}{\left(1 - \frac{1}{\kappa_2}\right)^{p-1}} + \kappa_2^{p-1} (s_{\max} - 1)^{p-1} \mu t \right) \|x\|_p^p ,$$

since each $u \in S$ can appear as v_r for some $r \in N(S)$ at most t times. As $\kappa_2 = \frac{1}{\delta_2}$, this yields Eq. (8). \square

7 Singular values of random biregular matrices

In this section, we prove Theorem 2, which is restated below.

Theorem 2 (Singular value bound). *With high probability, the set of singular values $\sigma(A)$ of A satisfy*

$$\sigma(A) \subseteq \left[\sqrt{s-1} - (1+o(1)) \cdot \sqrt{t-1}, \sqrt{s-1} + (1+o(1)) \cdot \sqrt{t-1} \right].$$

Moreover, the above bound holds even without our (otherwise global) assumption that $s \leq n^c$ for some absolute constant $c < 1$.

For the entire section, we will let $A \in \mathcal{M}_{m,n,s,t}$, and let $G := G_A = (V_L, V_R, E)$ and $\text{sign} = \text{sign}_A$. For convenience, let us define $a \pm b := [a-b, a+b]$.

We begin our proof of Theorem 2 with the following simple claim.

Claim 7.1. Let $M = AA^\top - s \cdot I$. Suppose that

$$\text{Spec}(M) \subseteq t - 2 \pm (2 + \varepsilon) \sqrt{(s-1)(t-1)} \tag{12}$$

for some $\varepsilon = o(1)$. Then,

$$\sigma(A) \subseteq \sqrt{s-1} \pm (1 + \varepsilon') \sqrt{t-1} \tag{13}$$

for some $\varepsilon' = o(1)$.

Proof. We immediately have that

$$\text{Spec}(AA^\top) \subseteq (s-1) + (t-1) \pm (2 + \varepsilon) \sqrt{(s-1)(t-1)}.$$

The singular values of A are obtained by taking the square root of the eigenvalues of AA^\top . Hence, Eq. (13) holds whenever ε' satisfies

$$2\varepsilon' - (2\varepsilon' + \varepsilon'^2) \sqrt{\frac{t-1}{s-1}} \geq \varepsilon.$$

Because $\frac{t-1}{s-1} \leq \frac{t}{s} = \alpha$ is a constant < 1 , we can take $\varepsilon' \leq O_\alpha(\varepsilon)$. \square

Thus, we focus on showing that Eq. (12) holds with high probability when A is sampled uniformly from $\mathcal{M}_{m,n,s,t}$. To do this, we use the following definitions and theorem from [MOP20b].

Definition 7.2 (Nomadic walk matrix). A *nomadic pair* is a 2-tuple (e, e') of (undirected) edges in E , where $e = \{u, v\}$ and $e' = \{v, w\}$ with $u \in V_R, v \in V_L$, and $w \in V_R$ with $w \neq u$. Let

$$V_{\text{Nomadic}(G)} = \{(e_1, e_2) \mid (e_1, e_2) \text{ is a nomadic pair in } G\}.$$

The *nomadic walk matrix* of A is the matrix $B \in \mathbb{R}^{V_{\text{Nomadic}(G)} \times V_{\text{Nomadic}(G)}}$, such that $B[(e_1, e_2), (e_3, e_4)] =$

$\text{sign}(e_3) \cdot \text{sign}(e_4)$ if (e_1, e_2, e_3, e_4) forms a non-backtracking¹⁷ walk of length 4 in G . Otherwise, $B[(e_1, e_2), (e_3, e_4)] = 0$.

Theorem 7.3 (A modified Ihara-Bass formula [MOP20b, Theorem 3.1]). *Let $A \in \mathcal{M}_{m,n,s,t}$ and let $M = AA^\top - s \cdot I$. Let $L(z) = I - z \cdot M + z \cdot (t-2)I + z^2 \cdot (s-1)(t-1)I$. Then,*

$$(1-z)^{m \cdot \frac{s(t-1)}{t}-1} (1+(t-1)z)^{m \cdot \frac{s}{t}-1} \det(L(z)) = \det(I - Bz) , \quad (14)$$

where B is the nomadic walk matrix of A .

We note that the above theorem is stated in [MOP20b] in more generality: for simplicity we only state the version specific to our application.

We use Theorem 7.3 to connect the spectrum of M with the spectrum of the nomadic walk matrix B .

Claim 7.4. Let $M = AA^\top - s \cdot I$ and let B denote the nomadic walk matrix of A . Suppose that

$$\rho(B) \leq (1 + \varepsilon) \sqrt{(s-1)(t-1)} \quad (15)$$

for some $\varepsilon \leq \frac{1}{2}$. Then

$$\text{Spec}(M) \subseteq t - 2 \pm (2 + 4\varepsilon^2) \sqrt{(s-1)(t-1)} .$$

In particular, if $\varepsilon = o_{n \rightarrow \infty}(1)$, then Eq. (12) holds.

Proof. Let λ be an eigenvalue of M so that there is a nonzero vector x with $Mx = \lambda x$. Let $z \in \mathbb{R}$ be such that

$$1 + z(t-2-\lambda) + z^2(s-1)(t-1) = 0. \quad (16)$$

Note that $z \neq 0$, and so $\lambda = \frac{(1+z(t-2)+z^2(s-1)(t-1))}{z}$. Then, rearranging, we have that $L(z)x = 0$, so that $\det(L(z)) = 0$, and therefore $\det(I - zB) = 0$ due to Eq. (14). It then follows that $\frac{1}{z}$ is an eigenvalue of B .

Now, suppose that there exists an eigenvalue λ of M such that either $\lambda = t - 2 - (2 + \delta)\sqrt{(s-1)(t-1)}$ or $t - 2 + (2 + \delta)\sqrt{(s-1)(t-1)}$ for some $\delta \geq 0$. In either case, by choosing z to be the solution to Eq. (16) that minimizes $|z|$, we see that $z = \pm \frac{1 + \frac{\delta}{2} - \frac{1}{2}\sqrt{2\delta + \delta^2}}{\sqrt{(s-1)(t-1)}}$. Hence, B has an eigenvalue μ with

$$\rho(B) \geq |\mu| \geq \frac{\sqrt{(s-1)(t-1)}}{1 + \frac{\delta}{2} - \frac{1}{2}\sqrt{2\delta + \delta^2}} .$$

If $\delta \geq 1$, then this implies that $\rho(B) > \frac{3}{2} \cdot \sqrt{(s-1)(t-1)}$, which contradicts Eq. (15), as $\varepsilon \leq \frac{1}{2}$. If $\delta < 1$, then we observe that $\frac{1}{1 + \frac{\delta}{2} - \frac{1}{2}\sqrt{2\delta + \delta^2}} \geq 1 + \frac{1}{2}\sqrt{\delta}$. Hence, $1 + \varepsilon \geq \frac{\rho(B)}{\sqrt{(s-1)(t-1)}} \geq 1 + \frac{1}{2}\sqrt{\delta}$, which implies that $\delta \leq 4\varepsilon^2$. \square

In light of Claims 7.1 and 7.4, in order to finish the proof of Theorem 2, it suffices to prove the following lemma.

¹⁷A non-backtracking walk of length ℓ is a sequence of vertices v_0, \dots, v_ℓ where $(v_{i-1}, v_i) \in E$ and $v_{i-1} \neq v_i$ for all $i \in [\ell]$.

Lemma 7.5. *Let A be uniformly sampled from $\mathcal{M}_{m,n,s,t}$, and let B be the respective nomadic walk matrix. Then, $\rho(B) \leq (1 + \varepsilon)\sqrt{(s-1)(t-1)}$ with high probability, for some $\varepsilon \leq o(1)$.*

The rest of this section is devoted to the proof of Lemma 7.5. This requires some definitions.

Definition 7.6 (Hikes of various kinds). A 2ℓ -hike \mathcal{H} in G is a closed walk of length exactly 4ℓ , starting in V_R , which is non-backtracking except possibly between the 2ℓ -th and $(2\ell + 1)$ -th steps. We let $\text{sign}(\mathcal{H})$ denote the product of $\text{sign}(e)$ over all edges in \mathcal{H} , with multiplicity. We say that a hike is *even* if all edges are traversed with even multiplicity, and *singleton-free* if no edge is traversed exactly once. Finally, we say that a hike with edges $(e_1, e_2, \dots, e_{4\ell})$ is *special* if $e_1 = e_{4\ell}$, $e_2 = e_{4\ell-1}$, $e_{2\ell} = e_{2\ell+1}$, and $e_{2\ell-1} = e_{2\ell+2}$, namely, the last two steps are the reverse of the first two, and the $(2\ell + 1)$ -th and $(2\ell + 2)$ -th steps are the reverse of the two steps preceding them. Note that an odd-numbered step is always from V_R to V_L , and vice-versa for an even numbered step.

Proof of Lemma 7.5. Fix $\ell \in \mathbb{N}$ to be determined later. Let $T := \text{tr}(B^\ell \cdot (B^\top)^\ell)$, and note that

$$\rho(B)^{2\ell} \leq \rho(B^\ell)^2 \leq \|B^\ell\|_2^2 = \|B^\ell (B^\top)^\ell\|_2 \leq T . \quad (17)$$

We turn to bounding T . First, note that

$$\begin{aligned} (B^\ell) [(f_1, f_2), (g_1, g_2)] &= \sum_{\substack{(e_1, e_2), (e_3, e_4), \dots, (e_{2\ell+1}, e_{2\ell+2}) \in V_{\text{Nomadic}(G)} \\ (e_1, e_2) = (f_1, f_2), (e_{2\ell+1}, e_{2\ell+2}) = (g_1, g_2)}} \prod_{i=1}^{\ell} B[(e_{2\ell-1}, e_{2\ell}), (e_{2\ell+1}, e_{2\ell+2})] \\ &= \sum_{\substack{(e_1, e_2, \dots, e_{2\ell+2}) \\ \text{is a non-backtracking walk in } G \text{ and} \\ e_1 = f_1, e_2 = f_2, e_{2\ell+1} = g_1, e_{2\ell+2} = g_2}} \prod_{i=3}^{2\ell+2} \text{sign}(e_i) , \end{aligned}$$

for $(f_1, f_2), (g_1, g_2) \in V_{\text{Nomadic}(G)}$. Similarly,

$$\begin{aligned} ((B^\top)^\ell) [(g_1, g_2), (f_1, f_2)] &= \sum_{\substack{(e_1, e_2, \dots, e_{2\ell+2}) \\ \text{is a non-backtracking walk in } G \text{ and} \\ e_1 = g_2, e_2 = g_1, e_{2\ell+1} = f_2, e_{2\ell+2} = f_1}} \prod_{i=1}^{2\ell} \text{sign}(e_i) . \end{aligned}$$

Hence,

$$\begin{aligned} T &= \sum_{(f_1, f_2), (g_1, g_2) \in V_{\text{Nomadic}(G)}} (B^\ell) [(f_1, f_2), (g_1, g_2)] ((B^\top)^\ell) [(g_1, g_2), (f_1, f_2)] \\ &= \sum_{(f_1, f_2), (g_1, g_2) \in V_{\text{Nomadic}(G)}} \sum_{\substack{(e_1, e_2, \dots, e_{2\ell+2}) \\ \text{and} \\ (e'_1, e'_2, \dots, e'_{2\ell+2}) \\ \text{are non-backtracking walks in } G \text{ and} \\ e_1 = e'_{2\ell+2} = f_1, e_2 = e'_{2\ell+1} = f_2, \\ e_{2\ell+1} = e'_2 = g_1, e_{2\ell+2} = e'_1 = g_2}} \prod_{i=3}^{2\ell+2} \text{sign}(e_i) \prod_{i=1}^{2\ell} \text{sign}(e'_i) \end{aligned}$$

$$= \sum_{\substack{(e_1, e_2, \dots, e_{2\ell+2}) \\ \text{and} \\ (e'_1, e'_2, \dots, e'_{2\ell+2}) \\ \text{are non-backtracking walks in } G \text{ and} \\ e_1 = e'_{2\ell+2}, e_2 = e'_{2\ell+1}, \\ e_{2\ell+1} = e'_2, e_{2\ell+2} = e'_1}} \prod_{i=3}^{2\ell+2} \text{sign}(e_i) \prod_{i=1}^{2\ell} \text{sign}(e'_i) .$$

Observe that the each sequence $(e_1, e_2, \dots, e_{2\ell+2}, e'_1, e'_2, e'_3, \dots, e'_{2\ell+2})$ in the above sum is a special $(2\ell + 2)$ -hike \mathcal{H} in G . Moreover,

$$\text{sign}(\mathcal{H}) = \text{sign}(e_1)\text{sign}(e_2)\text{sign}(e'_{2\ell+1})\text{sign}(e'_{2\ell+2}) \left(\prod_{i=3}^{2\ell+2} \text{sign}(e_i) \prod_{i=1}^{2\ell} \text{sign}(e'_i) \right) = \prod_{i=3}^{2\ell+2} \text{sign}(e_i) \prod_{i=1}^{2\ell} \text{sign}(e'_i) ,$$

as $e_1 = e'_{2\ell+2}$ and $e_2 = e'_{2\ell+1}$. Hence,

$$T = \sum_{\mathcal{H} \text{ is a special } 2(\ell + 1)\text{-hike in } G} \text{sign}(\mathcal{H}) .$$

We proceed by conditioning on G . When G is fixed, each $\text{sign}(e)$ for $e \in E$ is drawn independently from $\{1, -1\}$. Thus, we see that $\mathbb{E}_{\text{sign}}[\text{sign}(\mathcal{H}) \mid G] = 1$ if every edge in \mathcal{H} appears with even multiplicity, and otherwise the expectation is 0. So, we have

$$\mathbb{E}_{\text{sign}}[T \mid G] = \# \text{ of even special } 2(\ell + 1)\text{-hikes in } G .$$

We can upper bound the latter by $(st)^2 \cdot |\mathcal{E}_{2(\ell-1)}(G)|$, where $\mathcal{E}_{2(\ell-1)}(G)$ is the set of even $2(\ell - 1)$ -hikes in G . Indeed, this is because every even special $2(\ell + 1)$ -hike can be formed from an even $2(\ell - 1)$ hike by attaching 2 steps at the beginning and their reverse at the end (at most st choices in total) and by attaching 2 steps and their reverse in the middle (at most st choices again). Let $\varepsilon = o_{n \rightarrow \infty}(1)$ be a function to be chosen later, and let $a = (1 + \varepsilon) \sqrt{(s-1)(t-1)}$. By Eq. (17) and Markov's inequality,

$$\begin{aligned} \Pr_{\text{sign}} [\rho(B) \geq a \mid G] &\leq \Pr_{\text{sign}} [T \geq a^{2\ell} \mid G] \\ &\leq \frac{\mathbb{E}_{\text{sign}} [T \mid G]}{a^{2\ell}} \leq \frac{(st)^2 \cdot |\mathcal{E}_{2(\ell-1)}(G)|}{a^{2\ell}} . \end{aligned}$$

Thus, to prove the lemma, it suffices to show that, for some ℓ and ε of our choice,

$$(st)^2 \cdot |\mathcal{E}_{2(\ell-1)}(G)| \leq o(a^{2\ell})$$

with high probability over the choice of G . Therefore, the following lemma implies Lemma 7.5.

Lemma 7.7. *Let $\ell = \lfloor \log_2^2 m \rfloor$. Then with high probability over the choice of G , the number of even 2ℓ hikes in G is at most $m(s-1)^\ell(t-1)^\ell(1+o(1))^\ell$.*

We prove Lemma 7.7 in the next two subsections. We break the proof into two cases, according to whether the sparsity s of G is $O(\log^c n)$ for some constant c , or $\omega(\log^c n)$ for every constant c . This then finishes the proof of Lemma 7.5, and thus also Theorem 2. \square

7.1 Counting hikes when $s \leq \text{polylog}(n)$

We prove Lemma 7.7 when $s \leq \log^c n$ for some $c > 0$ by showing the following lemma.

Lemma 7.8. *Fix $\ell \in \mathbb{N}$. Suppose that G is bicycle-free at radius r for some $r \geq 20 \ln \ell$. Then the number of singleton-free 2ℓ -hikes in G is at most*

$$O(\ell^4 m)[(s-1)(t-1)]^\ell (2r\ell(s-1)(t-1))^{O(\frac{\log \ell}{r}) \cdot \ell}.$$

Proof of Lemma 7.7 for $s \leq \text{polylog}(n)$ from Lemma 7.8. Let $c > 0$ so that $s \leq \log^c n$. By Proposition 4.1, it holds with high probability that G is bicycle-free at radius r for $r = c' \log n / \log s$ for some absolute constant c' . Assume that this event holds. As $\ell \leq \log_2^2 m \leq \log_2^2 n$, it follows that $20 \ln \ell = O(\log \log n)$ and $r = \Omega(\log n / \log \log n)$, so that $r \geq 20 \ln \ell$ for n sufficiently large. Hence, by Lemma 7.8, we conclude that the number of singleton-free 2ℓ -hikes in G , and in particular the number of even 2ℓ -hikes in G , is at most

$$O(\ell^4 m)[(s-1)(t-1)]^\ell (2r\ell(s-1)(t-1))^{O(\frac{\log \ell}{r}) \cdot \ell}.$$

We clearly have that $O(\ell^4) = (1 + o(1))^\ell$. We have shown that $\log \ell / r = O((\log'' \log n)^2 / \log n)$, and we have $2r\ell(s-1)(t-1) \leq \log^{c''} n$ for some constant c'' . We thus have that

$$(2r\ell(s-1)(t-1))^{O(\frac{\log \ell}{r})} = \exp(\text{poly}(\log \log n) / \log n) = 1 + o(1),$$

which shows that the number of even 2ℓ -hikes is at most

$$m(1 + o(1))^\ell [(s-1)(t-1)]^\ell,$$

as required. □

Proof of Lemma 7.8. We now turn to proving Lemma 7.8. Recall that since G is bicycle-free at radius r , the subgraph $B_G(v, r)$ of G contains at most one cycle for every $v \in V_L$. We will count the number of even 2ℓ -hikes via an encoding argument.

Let us first consider a fixed even 2ℓ -hike \mathcal{H} . For this hike, we let $G_{\mathcal{H}}$ denote the subgraph of G consisting of the edges traversed by \mathcal{H} . For intuition, one should think of the graph $G_{\mathcal{H}}$ as being “discovered” by the decoding algorithm as it follows the hike \mathcal{H} .

We classify each step of \mathcal{H} as either *fresh*, *stale*, or *boundary*, as follows:

- If a step of \mathcal{H} traverses a new edge and it steps into a previously unexplored vertex, it is *fresh*.
- If a step of \mathcal{H} traverses a new edge but it steps into a previously explored vertex, then it is *boundary*.
- If a step of \mathcal{H} traverses a previously used edge, then it is *stale*.

We note that a singleton-free 2ℓ -hike must have at most 2ℓ fresh steps, as every edge must be traversed at least twice. Next, we observe that the number of boundary steps is exactly $|E(G_{\mathcal{H}})| + 1 - |V(G_{\mathcal{H}})|$. As $|V(G_{\mathcal{H}})| \leq 4\ell$, we have $r \geq 10 \ln(4\ell)$. By [MOP20a, Theorem 2.13], we thus have that the number

of boundary steps is at most

$$|E(G_{\mathcal{H}})| + 1 - |V(G_{\mathcal{H}})| \leq \frac{1 + \ln|V(G_{\mathcal{H}})|}{r} \cdot |V(G_{\mathcal{H}})| \leq O\left(\frac{\log \ell}{r}\right) \cdot \ell .$$

We group the stale steps into “stretches” of contiguous stale steps. First, we can group the stale steps into maximal contiguous blocks of stale steps. Note that we have at most $O\left(\frac{\log \ell}{r}\right) \cdot \ell$ blocks, as each block of stale steps must follow a boundary step or the midpoint of the hike. If the midpoint of the hike is contained in one block, we split the block into two. This makes the sequence of stale steps in each block non-backtracking. Next, we can divide the blocks so that they each have size at most r . As there are trivially at most 4ℓ stale steps, this can only increase the number of blocks by at most $\frac{4\ell}{r} = O\left(\frac{\log \ell}{r}\right) \cdot \ell$ (as the “worst case” is when all blocks have size exactly $r + 1$). We have thus argued that we can divide the stale steps into $O\left(\frac{\log \ell}{r}\right) \cdot \ell$ blocks of contiguous stale steps, each of size at most r , and all steps in each block are non-backtracking.

We now describe the encoder and decoder simultaneously. For any hike \mathcal{H} , we let the *shape* of \mathcal{H} be the sequence $\sigma \in \{\mathbb{F}, \mathbb{B}, \mathbb{S}\}^{4\ell}$ that indicates which steps are “fresh”, “boundary”, and “stale”. The first part of the encoding is the following sequence of numbers:

- $c_{\mathbb{F}}$, the number of times \mathbb{F} appears in σ ,
- $c_{\mathbb{B}}$, the number of times \mathbb{B} appears in σ ,
- $c_{\mathbb{S}}$, the number of “stale stretches”. That is, the number of blocks of contiguous stale steps where each block has size at most r and the midpoint of the hike does not occur inside the block.

There are naively at most $O(\ell^3)$ choices for these numbers.

Next, we let σ' denote the shape σ , only we compress a stale stretch into one symbol \mathbb{s} . We call σ' the *compressed shape* of \mathcal{H} . The next part of the encoding is σ' . Note that there are naively $O(\ell) \cdot \ell^{O\left(\frac{\log \ell}{r}\right) \cdot \ell}$ choices of σ' (once $c_{\mathbb{F}}$, $c_{\mathbb{B}}$ and $c_{\mathbb{S}}$ are chosen), as we can specify its length ($O(\ell)$ choices), and then the locations of the \mathbb{B} and \mathbb{s} symbols (ℓ choices for each symbol, and there are $O\left(\frac{\log \ell}{r}\right) \cdot \ell$ symbols).

Now, the encoder specifies the start vertex of \mathcal{H} (m choices). For each \mathbb{F} and \mathbb{B} step, the encoder specifies which neighbor¹⁸ the hike moves to ($(s - 1)$ choices if we start at a right vertex, and $(t - 1)$ if we start at a left vertex¹⁹). A stale stretch is encoded as follows. Let u be the start vertex of the stretch, and let v be the end vertex of the stretch. Let $K = \mathbb{B}_G(u, r)$, which is bicycle-free. The stale stretch is a path of length $\leq r$ from u to v in K , so it can be specified by: (1) specifying the end vertex v (at most 4ℓ choices, and is specified by giving the first step²⁰ at which v is reached in the hike), (2) specifying the number of the times the (unique) cycle in K is traversed by the path (which is an integer $\leq \frac{r}{4}$), and (3) specifying which direction²¹ the cycle is traversed in (at most 2 choices). For each stale stretch, the encoder specifies the above information.

¹⁸We fix some arbitrary order for the neighbors of each vertex.

¹⁹Note that there are s choices for the first step of the hike, but we will count this as $s - 1$ and instead pay an extra factor of $\frac{s}{s-1} \leq 2$ in the final bound.

²⁰Note that since the step is *stale*, the hike must have reached v at some earlier point.

²¹The cycle can only be traversed in one direction, as K has only one cycle and each stale stretch is non-backtracking.

Above, we have argued that the information given by the encoder is enough to reconstruct the hike \mathcal{H} uniquely. This implicitly defines the decoder. It therefore remains to upper bound the total number of strings that can be outputted by the encoder.

Let \mathcal{H} be a singleton-free 2ℓ -hike with shape σ , and let c_F, c_B, c_S be the counts defined earlier (which are fixed by the shape σ). We observe that the number of singleton-free 2ℓ -hikes with shape σ is at most

$$2m(s-1)^{c_F^{(L)}+c_B^{(L)}} \cdot (t-1)^{c_F^{(R)}+c_B^{(R)}} \cdot (2r\ell)^{c_S},$$

where $c_F^{(L)}, c_B^{(L)}$ are the number of fresh/boundary steps taken that start at a left vertex, and similarly for $c_F^{(R)}, c_B^{(R)}$. Note that this is well-defined, as $c_F^{(L)}, c_B^{(L)}, c_F^{(R)}, c_B^{(R)}$ are determined by the shape σ .

We next observe that every $c_F^{(R)}$ step is followed by either a $c_F^{(L)}$ or $c_B^{(L)}$ step, so $c_F^{(R)} \leq c_F^{(L)} + c_B^{(L)}$, and similarly $c_F^{(L)} \leq c_F^{(R)} + c_B^{(R)}$. Thus, we have

$$\begin{aligned} |c_F^{(R)} - c_F^{(L)}| &\leq c_B^{(L)} + c_B^{(R)} = c_B \\ \implies |c_F^{(R)} - \frac{c_F}{2}| &\leq \frac{c_B}{2} \text{ and } |c_F^{(L)} - \frac{c_F}{2}| \leq \frac{c_B}{2}. \end{aligned}$$

Hence, the number of singleton-free 2ℓ -hikes with shape σ is at most

$$2m(s-1)^{\frac{1}{2}c_F + \frac{3}{2}c_B} \cdot (t-1)^{\frac{1}{2}c_F + \frac{3}{2}c_B} \cdot (2r\ell)^{c_S}. \quad (18)$$

As the above bound is determined completely by the compressed shape σ' , we see that it upper bounds the number of singleton-free 2ℓ -hikes with compressed shape σ' .

Recall that the number of valid compressed shapes σ' is at most $O(\ell^4) \cdot \ell^{O(\frac{\log \ell}{r}) \cdot \ell}$. Because Eq. (18) is an increasing function in c_F, c_B and c_S , we can upper bound the total number of singleton-free 2ℓ -hikes by (# of choices of σ') \cdot (maximum value of Eq. (18)). We thus have that the number of singleton-free 2ℓ -hikes is at most

$$\begin{aligned} O(\ell^4) \ell^{O(\frac{\log \ell}{r}) \cdot \ell} \cdot 2m(s-1)^{\ell + O(\frac{\log \ell}{r}) \cdot \ell} \cdot (t-1)^{\ell + O(\frac{\log \ell}{r}) \cdot \ell} \cdot (2r\ell)^{O(\frac{\log \ell}{r}) \cdot \ell} \\ = O(m\ell^4) [(s-1)(t-1)]^\ell \cdot (2r\ell(s-1)(t-1))^{O(\frac{\log \ell}{r}) \cdot \ell}, \end{aligned}$$

which finishes the proof of Lemma 7.8. \square

7.2 Counting hikes when $s = \omega(\text{polylog}(n))$

Lemma 7.9. *Let $s = \omega(\log^c n)$ for every constant c , let $\ell = \lfloor \log_2^2 m \rfloor$, and let G be an arbitrary (t, s) -biregular graph. Then the number of even 2ℓ -hikes in G is at most*

$$m \cdot (s-1)^\ell \cdot (t-1)^\ell \cdot \left(1 + \frac{\ell^{O(1)}}{t}\right).$$

Note that as $\ell = O(\log^2 m) = O(\log^2 n)$ and $t = \Omega(s) = \omega(\log^c n)$ for every constant c , it follows that $\frac{\ell^{O(1)}}{t} = o(1)$, so Lemma 7.9 implies Lemma 7.7 when $s = \omega(\text{polylog}(n))$.

Proof of Lemma 7.9. We begin by introducing some terminology. Let \mathcal{H} be an even 2ℓ -hike in G ,

and let $G_{\mathcal{H}}$ be the subgraph of G induced by the edges in \mathcal{H} . As in the case for $s \leq \text{polylog}(n)$, one should think of the graph $G_{\mathcal{H}}$ as being “discovered” as we traverse the hike \mathcal{H} .

For an edge e , we let a_e denote the multiplicity of the edge e in \mathcal{H} , i.e., the number of times that it is traversed (in either direction). We say that e is high multiplicity if $a_e > 2$.

A step in \mathcal{H} is a traversal of an edge in \mathcal{H} . That is, a step is an edge in \mathcal{H} , but we count different occurrences of an edge in \mathcal{H} as distinct steps. For a step in \mathcal{H} , we say that the endpoint of the step is the vertex where the hike “ends up” after this step.

We say that a step in the hike is high multiplicity (labeled as \mathfrak{H}) if it uses a high multiplicity edge. We say that a non- \mathfrak{H} step in the hike is fresh (labeled as \mathfrak{F}) if it traverses a new edge for the first time and the endpoint is a previously unvisited vertex. If the endpoint is instead a vertex that has been previously visited, we call it boundary (labeled as \mathfrak{B}). We note that any non-high multiplicity edge is traversed exactly twice in \mathcal{H} . We say that a non- \mathfrak{H} step is forced (labeled as \mathfrak{R}) if at that point in the hike, the current vertex has exactly one non-high multiplicity neighboring edge that has not yet been traversed twice. All other steps in the hike are unforced (labeled as \mathfrak{U}).

For an edge e in \mathcal{H} with $a_e = 2$, we call the edge “fresh” if when e is traversed by \mathcal{H} for the first time, the step is fresh, and otherwise we call e “boundary”.

Remark 7.10. We note that the definitions above differ slightly from those in Section 7.1. For instance, here we separate out the \mathfrak{H} steps into their own class, and only split non- \mathfrak{H} steps into \mathfrak{F} and \mathfrak{B} steps.

For a hike \mathcal{H} , we define the following quantities:

- $E(\mathcal{H})$, the number of distinct edges in \mathcal{H} ,
- $n(\mathcal{H})$, the number of distinct left vertices in \mathcal{H} ,
- $m(\mathcal{H})$, the number of distinct right vertices in \mathcal{H} ,
- $h(\mathcal{H})$, the number of distinct high multiplicity edges in \mathcal{H} ,
- $f(\mathcal{H})$, the number of distinct fresh edges in \mathcal{H} ,
- $b(\mathcal{H})$, the number of distinct boundary edges in \mathcal{H} ,
- $u(\mathcal{H})$, the number of unforced steps in \mathcal{H} .
- $\tau(\mathcal{H})$, the type of \mathcal{H} , which the tuple $(h(\mathcal{H}), f(\mathcal{H}), b(\mathcal{H}), u(\mathcal{H}))$, along with a string $\{\mathfrak{H}, \mathfrak{F}, \mathfrak{B}, \mathfrak{R}, \mathfrak{U}\}^{4\ell}$ that specifies the type of each step in \mathcal{H} .

Finally, we say that a type τ is *valid* if there exists a hike \mathcal{H} of type τ .

We first prove some basic facts about the above quantities.

Claim 7.11. For all \mathcal{H} , $n(\mathcal{H}) \leq \ell$ and $m(\mathcal{H}) \leq \ell + 1$.

Proof. Consider a vertex v traversed in the hike \mathcal{H} that is not the midpoint. We claim that v must appear at least twice in the hike. (The start and end of the hike, which are the same vertex, are each counted separately.) Indeed, this is because any edge $e = (u, v)$ that is first traversed as $u \rightarrow v$ must be traversed at least twice, and cannot be traversed via the path $u \rightarrow v \rightarrow u$ because \mathcal{H} is non-backtracking. We note that since \mathcal{H} has length 4ℓ , the midpoint of \mathcal{H} is a right vertex, and \mathcal{H} traverses 2ℓ left vertices and $2\ell + 1$ right vertices (counting multiplicities). Since each vertex counted in $n(\mathcal{H})$ appears at least twice, it follows that $2n(\mathcal{H}) \leq 2\ell$. Since each vertex except for one (the

midpoint) counted in $m(\mathcal{H})$ appears at least twice, it follows that $1 + 2(m(\mathcal{H}) - 1) \leq 2\ell + 1$, and so $m(\mathcal{H}) \leq \ell + 1$. \square

Claim 7.12. $f(\mathcal{H}) + h(\mathcal{H}) + b(\mathcal{H}) = E(\mathcal{H})$, $h(\mathcal{H}) \leq 2\ell - E(\mathcal{H})$, and $u(\mathcal{H}) \leq 4\ell - 2f(\mathcal{H})$.

Proof. We observe that each distinct edge is counted exactly once in $f(\mathcal{H})$, $h(\mathcal{H})$, $b(\mathcal{H})$, so the first equality holds.

For the first inequality, we observe that $4h(\mathcal{H}) \leq \sum_{e \in \mathcal{H}: a_e > 2} a_e = 4\ell - \sum_{e \in \mathcal{H}: a_e = 2} a_e = 4\ell - 2(f(\mathcal{H}) + b(\mathcal{H}))$, where we use the fact that a_e is always even. This implies that $h(\mathcal{H}) \leq 2\ell - (f(\mathcal{H}) + b(\mathcal{H}) + (\mathcal{H})) = 2\ell - E(\mathcal{H})$.

To show the final inequality, we first make the following definition. We call a “return step” a step that is either \cup or \mathfrak{R} , and we call a non-high multiplicity edge “available” if it has been traversed exactly once.

Let v be any vertex in \mathcal{H} that is not the start vertex. We keep track of a counter c_v , that counts the number of available edges at v every time \mathcal{H} is at v . When the hike first arrives at v , it does so via a \mathfrak{F} or \mathfrak{H} step, and so v has at most one available edge. So, c_v starts at either 0 or 1. Every time the hike leaves v via a return step, it decreases the number of available edges by 1. So, c_v decreases by 1 in this case, and moreover the step must be \mathfrak{R} if $c_v = 1$ prior to the return step. Every time the hike leaves v via a non-return step that is also a non-high multiplicity edge, it adds 1 to the number of available edges adjacent to v , and every time the hike returns to v it will remove an available edge, unless it returns to v via a high multiplicity or boundary edge. This means that the c_v increases by at most 1 (resp. 2) every time \mathcal{H} returns to v using a high multiplicity edge (resp. boundary edge), and cannot increase otherwise.

By the above, it follows that the number of \cup steps from v is at most $-1 + \#$ of times c_v increases. We note that $\#$ of times c_v increases is $\leq \#$ of times a high multiplicity edge enters $v + 2 \cdot \#$ of times a boundary edge enters v . Summing over all vertices, we see that $u(\mathcal{H}) \leq \sum_{e: a_e > 2} a_e + 2b(\mathcal{H})$. Since $\sum_{e: a_e > 2} a_e = 4\ell - 2(f(\mathcal{H}) + b(\mathcal{H}))$, we are done. \square

We are now ready to bound the number of even 2ℓ -hikes \mathcal{H} . Fix ordering on the vertices of G , so that if we have a vertex u , then the neighbors of u are numbered uniquely from $1, \dots, s-1$ (if u is a left vertex) or $1, \dots, t-1$ (if u is a right vertex).

First, we consider the case when $E(\mathcal{H}) = 2\ell$ and $b(\mathcal{H}) = 0$. This will be the dominant term. By Claim 7.12, it follows that $h(\mathcal{H}) = 0$, so \mathcal{H} has no high multiplicity edges. Since $b(\mathcal{H}) = 0$, we see that $f(\mathcal{H}) = E(\mathcal{H}) = 2\ell$. Now, we consider the sequence of steps made by \mathcal{H} . The first step is \mathfrak{F} , and a \mathfrak{F} step must always be followed by a \mathfrak{H} , \mathfrak{F} , or \mathfrak{B} step, except at the midpoint of the hike. Because $h(\mathcal{H}) = b(\mathcal{H}) = 0$, this means that all steps must be \mathfrak{F} until the midpoint. But then this means that there can be no more \mathfrak{F} steps, as the number of \mathfrak{F} steps is at most 2ℓ , and so all the remaining steps must be \mathfrak{R} . We can count the number of such hikes by (1) picking the start vertex in V_R (m choices), and (2) for every fresh step, picking the neighbor of the current vertex to move to. We see that there are exactly ℓ fresh steps starting from a right vertex and ℓ fresh steps starting from a left vertex, so the total number of such hikes is at most

$$m(s-1)^\ell(t-1)^\ell. \quad (19)$$

Now, we assume that $E(\mathcal{H}) < 2\ell$ or $b(\mathcal{H}) > 0$. Fix a type τ with either $E < 2\ell$ or $b > 0$. We bound the number of hikes \mathcal{H} where $\tau(\mathcal{H}) = \tau$ via an encoding argument.

The encoding of \mathcal{H} is as follows.

- For each high multiplicity edge (u, v) (where the edge is first traversed as $u \rightarrow v$ in \mathcal{H}), specify which neighbor v of u is the endpoint of the edge.
- Specify the start vertex.
- For each \mathfrak{H} step, specify which of the $h(\mathcal{H})$ edges is the high multiplicity edge being used at this step.
- For each fresh step \mathfrak{F} , specify which neighbor of the current vertex is the endpoint of the edge.
- For each boundary step \mathfrak{B} , we specify the first location in the hike of the vertex that is the endpoint of this edge.
- For each unforced step \mathfrak{U} , specify the first location in the hike of the vertex that is the endpoint of this edge.
- For each forced step \mathfrak{R} , we do nothing.

We now specify the decoder. That is, we show how to uniquely construct the hike \mathcal{H} from the above data. Indeed, suppose we have reconstructed the hike \mathcal{H} correctly for the first i steps, ending at vertex u . (Note that the base case is trivial, as we are given the start vertex and so we can reconstruct \mathcal{H} after 0 steps.) Then, if the $(i + 1)$ -th step is \mathfrak{H} , we know which high multiplicity edge will be traversed, and so we know which neighbor of u to move to.²² If the step is \mathfrak{F} , then we also know which neighbor to move to. If the step is $\mathfrak{B}/\mathfrak{U}$, then we know that we move to the vertex that we were at in step $j < i$, so we can correctly reconstruct this step. Finally, if the step is \mathfrak{R} , then there is only one possible edge that can be traversed (and moreover, we can easily find out which edge this is, by simply keeping track of the number of times each “seen” edge has been traversed), so we can also reconstruct this step. Thus, one can uniquely reconstruct the hike \mathcal{H} from the encoding.

We now count the number of hikes \mathcal{H} with $\tau(\mathcal{H}) = \tau$. We observe that the total number of vertices specified (other than the start vertex) is $f + h$. When choosing to move to a right vertex, there are $(t - 1)$ choices (except for the start vertex, where there are m choices), and when choosing to move to a left vertex there are $(s - 1)$ choices. So, we see that there are at most $m \cdot (s - 1)^\ell \cdot (t - 1)^{f+h-\ell}$ possibilities, using the fact that $n(\mathcal{H}) \leq \ell$ always. This accounts for the choice of vertices from the \mathfrak{H} and \mathfrak{F} steps.

For each \mathfrak{H} step, we specify which of the h edges we use. As $h \leq 4\ell$ naively, this is at most 4ℓ choices per \mathfrak{H} step. Each \mathfrak{B} step and \mathfrak{U} step has naively at most 4ℓ choices, as we traverse at most 4ℓ distinct vertices. So, the total number of even 2ℓ -hikes \mathcal{H} of type τ is therefore at most

$$m \cdot (s - 1)^\ell \cdot (t - 1)^{f+h-\ell} \cdot (4\ell)^{h+b+u} .$$

Next, we observe that for a fixed f, h, b, u , there are at most $(4\ell)(4\ell)^{2(h+b+u)}$ valid types τ (namely, types τ corresponding to at least one \mathcal{H}). Indeed, we only need to specify the string of $\{\mathfrak{H}, \mathfrak{F}, \mathfrak{B}, \mathfrak{R}, \mathfrak{U}\}^{4\ell}$, which can be done as follows. We choose the locations of the \mathfrak{H} , \mathfrak{B} , and \mathfrak{U} steps, of which there are naively $(4\ell)^{h+b+u}$ choices. This gives us a partially filled in string, with the locations of the \mathfrak{F} and \mathfrak{R}

²²This technically holds only for the first traversal of the high multiplicity edge, as later traversals may be in the other direction. However, after the first traversal we have constructed the edge (u, v) , so this can be done trivially on later steps.

steps still undetermined. However, we observe that we cannot have a \mathbb{R} step after a \mathbb{F} step (except at the midpoint), and so we can determine these steps by specifying the number of \mathbb{R} steps in each “unfilled gap” in the string, where we split the gap containing the midpoint (if it exists) into two gaps. There are at most $h + b + u + 2$ such gaps, as there are $h + b + u + 1$ unfilled gaps, and we split the gap with the midpoint into two. So, we have at most $(4\ell)^{h+b+u+2}$ choices here.

Thus, for a fixed f, h, b, u , by Section 7.2 there are at most

$$m \cdot (s-1)^\ell \cdot (t-1)^{f+h-\ell} \cdot (4\ell)^{h+b+u} \cdot (4\ell)(4\ell)^{2(h+b+u)}$$

even 2ℓ -hikes with $f(\mathcal{H}) = f$, $h(\mathcal{H}) = h$, $b(\mathcal{H}) = b$ and $u(\mathcal{H}) = u$.

By Claim 7.12, the total number of \mathcal{H} with $E(\mathcal{H}) = E$ and $b(\mathcal{H}) = b$ is therefore at most

$$m \cdot (s-1)^\ell \cdot (t-1)^{E-b-\ell} \cdot (4\ell)^{15(2\ell-E)+9b} \cdot (4\ell)^3,$$

where we multiply by $(4\ell)^2$ to account for the choices of h and u . Now, to get the total number of hikes \mathcal{H} with either $E(\mathcal{H}) < 2\ell$ or $b(\mathcal{H}) > 0$, we sum over all possible choices of E and b . We have $b \leq E$, $E \leq 2\ell$, and that either $E < 2\ell$ or $b > 0$. Using the fact that $t = \omega(\log^c m)$ for all constants c and that $\ell = \log^2 m$, we conclude that the total number of such hikes is at most

$$m \cdot (s-1)^\ell \cdot (t-1)^\ell \cdot \frac{\ell^{O(1)}}{t}.$$

Thus, by Eq. (19), the total number of even 2ℓ hikes is $m \cdot (s-1)^\ell \cdot (t-1)^\ell \cdot (1 + \frac{\ell^{O(1)}}{t})$, as required. \square

A A weak positive bound on ℓ_2 -spread

We prove Theorem 3, which we recall below.

Theorem 3 (Converse to Theorem 1). *Assume that $t \geq 9$. Then, with high probability over A , the space $\ker(A)$ is $(\Omega(\alpha^2 n/t^4), \alpha^{O(\log n/\log t)})$ - ℓ_2 -spread.*

We use the following lemma whose proof we defer to the end of this section.

Lemma A.1. *Let $A \in \mathcal{M}_{m,n,s,t}$ such that G_A is a (γ, μ) -unique expander, where $0 < \mu \leq \frac{2}{9}$, $0 < \gamma \leq 2\mu$ and $\gamma n \geq \left(\frac{1}{\mu}\right)^c$. Then, for every γn -sparse $x \in \mathbb{R}^n$, there holds*

$$\|Ax\|_2 \geq c_1 \cdot \sqrt{t} \cdot \left(\frac{\sqrt{t}}{\|A\|_2} \right)^{\frac{c_2 \cdot \log(\gamma n)}{\log \frac{1}{\mu}}} \cdot \|x\|_2. \quad (20)$$

Remark A.2. For every $A \in \mathcal{M}_{m,n,s,t}$ we have $\|A\|_2 \geq \|Ae_1\|_2 = \sqrt{t}$. Hence, the term $\frac{\sqrt{t}}{\|A\|_2}$ in the right-hand side of Eq. (20) is at most 1.

Proof of Theorem 3. By Proposition 4.2 and Theorem 2, it holds with high probability, G_A is a (γ, μ) -unique expander with $\gamma = \Omega(\alpha^2/t^4)$, $\mu = 2/t$, and also that $\|A\|_2 \leq 2\sqrt{s}$. Assume that these events hold.

Let $y \in \mathbb{R}^n$ be (γ, ε) -compressible for some $\varepsilon > 0$ to be determined later, and let $x \in \mathbb{R}^n$ be γn -sparse with $\|x - y\|_2 \leq \varepsilon$. Then,

$$\|Ay\|_2 \geq \|Ax\|_2 - \|A(x - y)\|_2 \geq \|Ax\|_2 - \varepsilon \|A\|_2 \geq \|Ax\|_2 - 2\varepsilon\sqrt{s}. \quad (21)$$

Note that for n large enough, γ and μ satisfy the hypothesis of Lemma A.1. In particular, $\mu \leq \frac{2}{9}$ by our assumption that $t \geq 9$. Hence, Eq. (20) applies to x . It follows that

$$\|Ax\|_2 \geq c_1 \sqrt{t} \cdot \left(\frac{\sqrt{t}}{2\sqrt{s}} \right)^{\frac{c_2 \log \frac{a^2 \mu}{t^4}}{\log t}}.$$

Denote the right-hand side of the above by a , so Eq. (21) yields

$$\|Ay\|_2 \geq a - 2\varepsilon\sqrt{s}.$$

Taking $\varepsilon = \frac{a}{4\sqrt{s}}$, we then have that $\|Ay\|_2 > 0$, and, in particular, $y \notin \ker(A)$. Thus, $\ker(A)$ is (γ, ε) - ℓ_2 -spread. The proposition follows since $\varepsilon \geq \alpha^{O\left(\frac{\log n}{\log t}\right)}$. \square

Proof of Lemma A.1. Write $G = G_A = (V_L, V_R, E)$. We need the following claim.

Claim A.3. Fix $k, b \in \mathbb{N}$ such that $kb \leq \gamma n$. Let $S_1, \dots, S_b \subseteq V_L$ be disjoint sets, each of size k . Then, there exist sets T_1, \dots, T_b , each of size $\geq (1 - \mu b)tk$, such that for each $1 \leq i \leq b$, every $r \in T_i$ has exactly one neighbor in S_i and no neighbors in any of the sets S_j ($j \neq i$).

Proof. Let $S = \sqcup_{i=1}^b S_i$, and let $T_i = U(S) \cap N(S_i)$. Note that T_1, \dots, T_b are pairwise disjoint. As $|S| \leq \gamma n$, the unique expansion of G yields $\sum_{i=1}^b |T_i| = \left| \sqcup_{i=1}^b T_i \right| = |U(S)| \geq t(1 - \mu)|S| = t(1 - \mu)kb$. Hence,

$$|T_i| \geq t(1 - \mu)kb - \sum_{j \neq i} |T_j| \geq t(1 - \mu)kb - tk(b - 1) = (1 - \mu b)tk. \quad \square$$

Let $b = \left\lceil \frac{1}{2\mu} \right\rceil \geq 2$. Let $M \in \mathbb{N}$ satisfy $b^{M-1} < \gamma n \leq b^M$. We partition the interval $[b^M]$ into consecutive intervals R_0, R_1, \dots, R_M where $R_\ell = \{b^{\ell-1} + 1, \dots, b^\ell\}$ for $0 \leq \ell \leq M$.

Let x be a γn -sparse vector with $\|x\|_2 = 1$. Without loss of generality, assume that $|x_1| \geq |x_2| \geq \dots \geq |x_n|$. In particular, $\text{supp}(x) \subseteq [\lfloor \gamma n \rfloor] \subseteq [b^M] = \cup_{\ell=0}^M R_\ell \subseteq [n]$. The last inclusion is due to $b^M = b^{M-1}b < \gamma nb \leq \frac{\gamma n}{2\mu} \leq n$, where the last inequality uses our assumption that $\gamma \leq 2\mu$.

For $0 \leq \ell \leq M$, let $z_\ell = x_{R_\ell}$, i.e., the vector that is equal to x on the set R_ℓ , and is 0 otherwise. Let $\beta = \frac{t}{32\|A\|_2^2}$. Note that $\beta \leq \frac{1}{2}$, due to Remark A.2. Consequently,

$$\sum_{\ell=0}^M \|z_\ell\|_2^2 = \|x\|_2^2 = 1 \geq \frac{\beta}{1 - \beta} > \sum_{\ell=0}^M \beta^{\ell+1}.$$

Hence, there must exist $0 \leq \ell \leq M$ such that $\|z_\ell\|_2 \geq \beta^{\ell+1}$. Let ℓ^* denote the largest ℓ for which this occurs.

Let $u = \sum_{\ell=1}^{\ell^*} z_\ell$ and $v = \sum_{\ell=\ell^*+1}^M z_\ell$, so that $x = u + v$. Because the z_ℓ 's have disjoint support, we have

$$\|v\|_2^2 = \sum_{\ell=\ell^*+1}^M \|z_\ell\|_2^2 \leq \sum_{\ell=\ell^*+1}^{\infty} \beta^{\ell+1} \leq 2\beta^{\ell^*+2}, \quad (22)$$

where we used that $\beta \leq \frac{1}{2}$.

We claim that

$$\|Au\|_2^2 \geq \frac{t}{2}\beta^{\ell^*+1}. \quad (23)$$

Let $S = \bigcup_{\ell=0}^{\ell^*} R_\ell \supseteq \text{supp}(u)$. We consider two cases. First, if $\ell^* = 0$, then $|S| = 1$, so that $\|Au\|_2^2 = t\|u\|_2^2 \geq \frac{t}{2}\beta^2 = \frac{t}{2}\beta^{2(\ell^*+1)}$, implying Eq. (23).

Next, suppose that $\ell^* \geq 1$. Note that $|S| = \left| \bigcup_{\ell=0}^{\ell^*} R_\ell \right| = b^{\ell^*}$. Partition S into b consecutive intervals S_1, \dots, S_b , each of size $k := \frac{|S|}{b} = b^{\ell^*-1}$, defined by $S_\ell = \{(\ell-1)k+1, \dots, \ell k\}$.

Note that S_2, \dots, S_b partition R_{ℓ^*} , and $S_1 = \bigcup_{\ell=0}^{\ell^*-1} R_\ell$. By Claim A.3, there exist sets T_1, \dots, T_b , each of size $\geq (1-\mu b)tk \geq \frac{tk}{2}$, such that each $r \in T_i$ is in $U(S_i)$ and is not in $N(S_j)$, for all $j \neq i$. Let $\eta_i = \min_{j \in S_i} |u_j|$. Then, $|(Au)_r| \geq \eta_i$ for each $r \in T_i$. Moreover, we must also have $\eta_i^2 \geq \|u_{S_{i+1}}\|_2^2 / k$, where $u_{S_{i+1}}$ denotes the restriction of u to the set S_{i+1} , as there are k entries in S_{i+1} , each with absolute value $\leq \eta_i$. Eq. (23) follows since

$$\|Au\|_2^2 \geq \sum_{i=1}^b \sum_{r \in T_i} |(Au)_r|^2 \geq \sum_{i=1}^b \frac{tk}{2} \eta_i^2 \geq \frac{tk}{2} \sum_{i=2}^b \frac{\|u_{S_i}\|_2^2}{k} = \frac{t}{2} \|u_{\bigcup_{i=2}^b S_i}\|_2^2 = \frac{t}{2} \|z_{\ell^*}\|_2^2 \geq \frac{t}{2} \beta^{\ell^*+1}.$$

Eqs. (22) and (23) now yield

$$\begin{aligned} \|Ax\|_2 &= \|A(u+v)\|_2 \geq \|Au\|_2 - \|A\|_2 \|v\|_2 \geq \frac{\sqrt{t}}{\sqrt{2}} \cdot \beta^{\frac{\ell^*+1}{2}} - 2\|A\|_2 \cdot \beta^{\frac{\ell^*+2}{2}} \\ &= \beta^{\frac{\ell^*+1}{2}} \left(\frac{\sqrt{t}}{\sqrt{2}} - 2\|A\|_2 \sqrt{\beta} \right) \geq \beta^{\frac{\ell^*+1}{2}} \cdot \frac{\sqrt{t}}{2\sqrt{2}} \geq \beta^{\frac{M+1}{2}} \cdot \frac{\sqrt{t}}{2\sqrt{2}}, \end{aligned}$$

and the lemma follows from the definitions of β and M , and from our assumptions $\mu \leq \frac{2}{9}$ and $\gamma n \geq \left(\frac{1}{\mu}\right)^c$. \square

B General relations between spread, distortion and restricted isometry properties

We prove Propositions 3.7, 3.8 and 3.11 in Appendices B.1 to B.3, respectively.

B.1 ℓ_p -spread implies ℓ_q -spread for $q < p$

We prove Proposition 3.7, which we restate below.

Proposition 3.7 (ℓ_p -spread implies ℓ_q -spread). *Suppose that $X \subseteq \mathbb{R}^n$ is $(2k, \varepsilon)$ - ℓ_p -spread. Then for every $1 \leq q < p$, X is (k, ε_q) - ℓ_q -spread for $\varepsilon_q = \varepsilon^2 \left(\frac{k}{n}\right)^{\frac{1}{q}}$.*

In particular, if X is $(\Omega(n), \Omega(1))$ - ℓ_p -spread, then X is also $(\Omega(n), \Omega(1))$ - ℓ_q -spread for every $1 \leq q < p$.

Proof. Let $x \in X$, and let $S \subseteq [n]$ with $|S| = k$ denote the k smallest coordinates (in absolute value) in x . Let $X_{\bar{S}}$ denote the subspace obtained by projecting X to the coordinates *not* in S . We first show the following claim.

Claim B.1. $X_{\bar{S}}$ is (k, ε) - ℓ_p -spread.

Proof of Claim B.1. Let $y_{\bar{S}} \in X_{\bar{S}}$ be obtained by projecting $y \in X$ to the coordinates not in S , and let $S' \subseteq \bar{S}$ be a subset of size k . The vector z with $\text{supp}(z) = S'$ that minimizes $\|y_{\bar{S}} - z\|_p$ is obtained by setting $z = y_{S'}$. Letting $w = y_{S \cup S'}$, we see that $y_{\bar{S}} - z = y - w$, as both are equal to $y_{\overline{S \cup S'}}$. As w is $2k$ -sparse,

$$\|y_{\bar{S}} - z\|_p = \|y - w\|_p \geq \varepsilon \|y\|_p \geq \varepsilon \|y_{\bar{S}}\|_p ,$$

where the first inequality uses that $y \in X$ and X is $(2k, \varepsilon)$ - ℓ_p -spread. As S' was arbitrary, this proves that $X_{\bar{S}}$ is (k, ε) - ℓ_p -spread. \square

We now use Claim B.1 to finish the proof. Let $y = x_S$. We have that

$$\|x_{\bar{S}}\|_p = \|x - y\|_p \geq \varepsilon \|x\|_p \geq \frac{\varepsilon}{n^{\frac{1}{q} - \frac{1}{p}}} \|x\|_q ,$$

by Hölder's inequality (Lemma 3.5) and using that X is $(2k, \varepsilon)$ - ℓ_p -spread. We also have that

$$\|x_{\bar{S}}\|_p \leq \|x_{\bar{S}}\|_q \cdot \frac{\Delta_{q,p}(X_{\bar{S}})}{n^{\frac{1}{q} - \frac{1}{p}}} .$$

As $X_{\bar{S}}$ is (k, ε) - ℓ_p -spread, we have $\Delta_{q,p}(X_{\bar{S}}) \leq \frac{1}{\varepsilon} \cdot \left(\frac{n}{k}\right)^{\frac{1}{q}}$ by Proposition 3.11. Hence,

$$\|x_{\bar{S}}\|_q \cdot \frac{1}{\varepsilon} \cdot \left(\frac{n}{k}\right)^{\frac{1}{q}} \geq \varepsilon \|x\|_q .$$

As $\|x_{\bar{S}}\|_q = \|x - y\|_q$ and S was arbitrary, we are done. \square

B.2 ℓ_p -RIP implies ℓ_p -spread

We prove Proposition 3.8, which we restate below. Our proof is based on an argument of [KT07].

Proposition 3.8 (ℓ_p -RIP implies ℓ_p -spread). *Let $p \in [1, \infty]$, and let $A \in \mathbb{R}^{m \times n}$ be a (k, ε) - ℓ_p -RIP matrix. Then, $\ker(A)$ is (k, ε') - ℓ_p -spread for $\varepsilon' = \frac{1-\varepsilon}{2+\varepsilon(1+(\frac{2n}{k})^{1-\frac{1}{p}})}$.*

Proof. Let A be a (k, ε) - ℓ_p -RIP matrix. Without loss of generality, we will assume that the “normalization factor” K is 1. We first upper bound $\|A\|_p$. Let $x \in \mathbb{R}^n$ be arbitrary. Partition $[n]$ arbitrarily into $b = \lceil \frac{n}{k} \rceil$ sets S_1, \dots, S_b , each of size $\leq k$, and for $i \in [b]$, let $z_i := x_{S_i}$. That is, z_i is x restricted to

the coordinates in S_i . By definition, the z_i 's are each k -sparse and satisfy $\sum_{i=1}^b z_i = x$. We thus have

$$\|Ax\|_p \leq \sum_{i=1}^b \|Az_i\|_p \leq \sum_{i=1}^b (1 + \varepsilon) \|z_i\|_p \leq (1 + \varepsilon) b^{1-\frac{1}{p}} \left(\sum_{i=1}^b \|z_i\|_p^p \right)^{1/p} = (1 + \varepsilon) b^{1-\frac{1}{p}} \|x\|_p ,$$

where the second inequality uses that A is (k, ε) - ℓ_p -RIP and the third inequality is by Lemma 3.5.

We thus have that $\|A\|_p \leq (1 + \varepsilon) b^{1-\frac{1}{p}}$.

Now, let $x \in \mathbb{R}^n$ with $\|x\|_p = 1$ be a (k, ε') -compressible vector, where $\varepsilon' < \frac{1-\varepsilon}{2+\varepsilon(1+b^{1-\frac{1}{p}})}$. We show that $\|Ax\|_p > 0$, so that $x \notin \ker(A)$, and thus $\ker(A)$ is (k, ε') -spread. Let $y \in \mathbb{R}^n$ be a k -sparse vector such that $\|x - y\|_p \leq \varepsilon'$. Note that $\|y\|_p \geq \|x\|_p - \|x - y\|_p \geq 1 - \varepsilon'$. We have that

$$\|Ax\|_p \geq \|Ay\|_p - \|A\|_p \|x - y\|_p \geq (1 - \varepsilon)(1 - \varepsilon') - (1 + \varepsilon) b^{1-\frac{1}{p}} \varepsilon' > 0 ,$$

by choice of ε' .

Finally, we note that $b^{1-\frac{1}{p}} < \left(\frac{2n}{k}\right)^{1-\frac{1}{p}}$, and so we can take $\varepsilon' = \frac{1-\varepsilon}{2+\varepsilon(1+(\frac{2n}{k})^{1-\frac{1}{p}})}$. \square

B.3 Equivalence between ℓ_p -spread and (ℓ_q, ℓ_p) -distortion

We prove Proposition 3.11, restated below. Our proof generalizes an argument in [GLR10].

Proposition 3.11 (Compressibility and distortion). *The following holds for all $1 \leq q < p$, $k \in \mathbb{N}$ and $x \in \mathbb{R}^n$.*

1. Let $\varepsilon > 0$. If x is (k, ε) - ℓ_p -compressible then $\Delta_{q,p}(x) \geq \frac{1}{\left(\frac{k}{n}\right)^{\frac{1}{q}-\frac{1}{p}} + \varepsilon}$.
2. The vector x is $\left(k, \frac{\left(\frac{n}{k}\right)^{\frac{1}{q}}}{\Delta_{q,p}(x)}\right)$ - ℓ_p -compressible.

In particular, if a subspace X is (k, ε) - ℓ_p -spread, then $\Delta_{q,p}(X) \leq \frac{1}{\varepsilon} \left(\frac{n}{k}\right)^{\frac{1}{q}}$ for all $1 \leq q < p$.

Proof. Clearly, it suffices to prove both statements under the assumption that $\|x\|_p = 1$. Let $a = \frac{1}{q} - \frac{1}{p}$.

1. Let $y \in \mathbb{R}^n$ be k -sparse with $\|x - y\|_p \leq \varepsilon$. Clearly, we may take y so that that y and $x - y$ have disjoint supports. In particular, this implies that $\|y\|_p \leq \|x\|_p = 1$. Thus,

$$\frac{1}{\Delta_{q,p}(x)} = \|x\|_q \cdot n^{-a} \leq (\|y\|_q + \|x - y\|_q) \cdot n^{-a} \leq (\|y\|_p \cdot k^a + \|x - y\|_p \cdot n^a) \cdot n^{-a} \leq \left(\frac{k}{n}\right)^a + \varepsilon,$$

where the penultimate inequality follows from Lemma 3.5.

2. Let $S \subseteq [n]$ be a set consisting of the k largest coordinates of x in absolute value (with arbitrary tie-breaking). Namely, $|S| = k$ and $|x_j| \geq |x_i|$ for all $j \in S, i \in [n] \setminus S$. Define $y \in \mathbb{R}^n$ by

$$y_i = \begin{cases} x_i & \text{if } i \in S \\ 0 & \text{if } i \in [n] \setminus S . \end{cases}$$

It now suffices to show that $\|x - y\|_p \leq \left(\frac{n}{k}\right)^{1/q} \frac{1}{\Delta_{p,q}(v)}$. Let $\sigma = \min\{|x_i| \mid i \in S\}$. Note that

$$\frac{1}{\Delta_{p,q}(x)} = \frac{\|x\|_q}{n^a} \geq \frac{\|y\|_q}{n^a} \geq \frac{k^{\frac{1}{q}} \sigma}{n^a}, \quad \text{so that} \quad \sigma \leq \frac{n^a}{\Delta_{p,q}(x) \cdot k^{\frac{1}{q}}}.$$

The claim follows since

$$\|x - y\|_p \leq \|x - y\|_\infty \cdot n^{\frac{1}{p}} \leq \sigma \cdot n^{\frac{1}{p}} \leq \left(\frac{n}{k}\right)^{\frac{1}{q}} \cdot \frac{1}{\Delta_{p,q}(x)}.$$

Finally, suppose that X is (k, ε) - ℓ_p -spread. Let $x \in X$ be any vector, and note that x is (k, ε) - ℓ_p -spread.

We also have that x is $\left(k, \frac{\left(\frac{n}{k}\right)^{\frac{1}{q}}}{\Delta_{q,p}(x)}\right)$ - ℓ_p -compressible, which implies that $\varepsilon \leq \frac{\left(\frac{n}{k}\right)^{\frac{1}{q}}}{\Delta_{q,p}(x)}$. Rearranging and taking the sup over $x \in X$, we conclude that $\Delta_{q,p}(X) \leq \frac{1}{\varepsilon} \left(\frac{n}{k}\right)^{\frac{1}{q}}$. \square

Acknowledgements

We thank Sidhanth Mohanty for helpful discussions about the works of [BDH18, Bor19, BC19, MOP20a, MOP20b, OW20], and Yuval Peled for helpful discussions about convergence theorems for graph spectra.

We thank Amir Shpilka for bringing the work of [Kar11] to our attention, and Ioana Dumitriu for bringing the works of [BDH18, Zhu20] to our attention.

References

- [AGR15] Zeyuan Allen-Zhu, Rati Gelashvili, and Ilya P. Razenshteyn. Restricted isometry property for general p -norms. In *31st International Symposium on Computational Geometry, SoCG 2015, June 22-25, 2015, Eindhoven, The Netherlands*, volume 34 of *LIPICs*, pages 451–460. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2015.
- [BBH⁺12] Boaz Barak, Fernando G. S. L. Brandão, Aram Wettroth Harrow, Jonathan A. Kelner, David Steurer, and Yuan Zhou. Hypercontractivity, sum-of-squares proofs, and their applications. In *Proceedings of the 44th Symposium on Theory of Computing Conference, STOC 2012, New York, NY, USA, May 19 - 22, 2012*, pages 307–326. ACM, 2012.
- [BC19] Charles Bordenave and Benoît Collins. Eigenvalues of random lifts and polynomials of random permutation matrices. *Annals of Mathematics*, 190(3):811–875, 2019.
- [BDDW07] Richard Baraniuk, Mark Davenport, Ronald DeVore, and Michael Wakin. A simple proof of the Restricted Isometry property for random matrices. Available at <https://www.math.tamu.edu/~rdevore/publications/131.pdf>, 2007.
- [BDH18] Gerandy Brito, Ioana Dumitriu, and Kameron Decker Harris. Spectral gap in random bipartite biregular graphs and applications. *arXiv preprint arXiv:1804.07808*, 2018.

- [BGI⁺08] R. Berinde, A. C. Gilbert, P. Indyk, H. Karloff, and M. J. Strauss. Combining geometry and combinatorics: A unified approach to sparse signal recovery. In *2008 46th Annual Allerton Conference on Communication, Control, and Computing*, pages 798–805, 2008.
- [Bor19] Charles Bordenave. A new proof of friedman’s second eigenvalue theorem and its extension to random lifts. In *Annales scientifiques de l’Ecole normale supérieure*, 2019.
- [BR17] Anirban Basak and Mark Rudelson. Invertibility of sparse non-hermitian matrices. *Advances in Mathematics*, 310:426–483, 2017.
- [BV16] Afonso S Bandeira and Ramon Van Handel. Sharp nonasymptotic bounds on the norm of random matrices with independent entries. *Annals of Probability*, 44(4):2479–2506, 2016.
- [CRT06] Emmanuel J. Candès, Justin K. Romberg, and Terence Tao. Stable signal recovery from incomplete and inaccurate measurements. *Comm. Pure Appl. Math.*, 59:1207–1223, 2006.
- [CRVW02] Michael R. Capalbo, Omer Reingold, Salil P. Vadhan, and Avi Wigderson. Randomness conductors and constant-degree lossless expanders. In *Proceedings on 34th Annual ACM Symposium on Theory of Computing, May 19-21, 2002, Montréal, Québec, Canada*, pages 659–668. ACM, 2002.
- [CT05] Emmanuel J. Candès and Terence Tao. Decoding by linear programming. *IEEE Trans. Inf. Theory*, 51(12):4203–4215, 2005.
- [Don06] David L. Donoho. Compressed sensing. *IEEE Trans. Inf. Theory*, 52(4):1289–1306, 2006.
- [FLM77] T. Figiel, J. Lindenstrauss, and V. D. Milman. The dimension of almost spherical sections of convex bodies. *Acta Math.*, 139(1-2):53–94, 1977.
- [Gal63] R. G. Gallager. *Low-Density Parity-Check Codes*. MIT Press, 1963.
- [GG84] A. Garnaev and E. D. Gluskin. The widths of Euclidean balls. *Doklady An. SSSR.*, 277:1048—1052, 1984.
- [GLR10] Venkatesan Guruswami, James R. Lee, and Alexander A. Razborov. Almost Euclidean subspaces of ℓ_1^N via expander codes. *Combinatorica*, 30(1):47–68, 2010.
- [Glu83] E. D. Gluskin. Norms of random matrices and diameters of finite-dimensional sets. *Mat. Sb. (N.S.)*, 120(162)(2):180–189, 286, 1983.
- [GLW08] V. Guruswami, J. Lee, and A. Wigderson. Euclidean sections of with sublinear randomness and error-correction over the reals. In *12th International Workshop on Randomization and Combinatorial Optimization: Algorithms and Techniques (RANDOM)*, pages 444–454, 2008.
- [Ind06] Piotr Indyk. Stable distributions, pseudorandom generators, embeddings, and data stream computation. *Journal of the ACM*, 53(3):307–323, 2006.
- [Ind07] Piotr Indyk. Uncertainty principles, extractors, and explicit embeddings of L_1 into L_2 . In *Proceedings of the 39th Annual ACM Symposium on the Theory of Computing*, pages 615–620, 2007.

- [IS10] Piotr Indyk and Stanislaw J. Szarek. Almost-euclidean subspaces of ℓ_1^n via tensor products: A simple approach to randomness reduction. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, 13th International Workshop, APPROX 2010, and 14th International Workshop, RANDOM 2010, Barcelona, Spain, September 1-3, 2010. Proceedings*, volume 6302 of *Lecture Notes in Computer Science*, pages 632–641. Springer, 2010.
- [Kar11] Zohar Shay Karnin. Deterministic construction of a high dimensional l_p section in l_1^n for any $p < 2$. In Lance Fortnow and Salil P. Vadhan, editors, *Proceedings of the 43rd ACM Symposium on Theory of Computing, STOC 2011, San Jose, CA, USA, 6-8 June 2011*, pages 645–654. ACM, 2011.
- [Kas77] B. S. Kashin. The widths of certain finite-dimensional sets and classes of smooth functions. *Izv. Akad. Nauk SSSR Ser. Mat.*, 41(2):334–351, 478, 1977.
- [KT07] Boris S Kashin and Vladimir N Temlyakov. A remark on compressed sensing. *Mathematical notes*, 82(5):748–755, 2007.
- [MOP20a] Sidhanth Mohanty, Ryan O’Donnell, and Pedro Paredes. Explicit near-ramanujan graphs of every degree. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing, STOC 2020, Chicago, IL, USA, June 22-26, 2020*, pages 510–523. ACM, 2020.
- [MOP20b] Sidhanth Mohanty, Ryan O’Donnell, and Pedro Paredes. The SDP value for random two-eigenvalue csps. In *37th International Symposium on Theoretical Aspects of Computer Science, STACS 2020, March 10-13, 2020, Montpellier, France*, volume 154 of *LIPICs*, pages 50:1–50:45. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020.
- [MRR⁺20] Jonathan Mosheiff, Nicolas Resch, Noga Ron-Zewi, Shashwat Silas, and Mary Wootters. LDPC codes achieve list decoding capacity. In *Proceedings of the 61st IEEE Annual Symposium on Foundations of Computer Science*, pages 458–469, 2020.
- [OW20] Ryan O’Donnell and Xinyu Wu. Explicit near-fully X-ramanujan graphs. In *61st IEEE Annual Symposium on Foundations of Computer Science, FOCS 2020, Durham, NC, USA, November 16-19, 2020*, pages 1045–1056. IEEE, 2020.
- [RU08] Tom Richardson and Ruediger Urbanke. *Modern Coding Theory*. Cambridge University Press, USA, 2008.
- [SS96] Michael Sipser and Daniel A. Spielman. Expander codes. *IEEE Trans. Inform. Theory*, 42(6, part 1):1710–1722, 1996. Codes and complexity.
- [Tan81] Robert M. Tanner. A recursive approach to low complexity codes. *IEEE Transactions on Information Theory*, 27(5):533–547, 1981.
- [Tao12] Terence Tao. *Topics in random matrix theory*, volume 132. American Mathematical Soc., 2012.
- [Vad12] Salil P. Vadhan. Pseudorandomness. *Foundations and Trends® in Theoretical Computer Science*, 7(1–3):1–336, 2012.

[Zhu20] Yizhe Zhu. On the second eigenvalue of random bipartite biregular graphs. *arXiv preprint arXiv:2005.08103*, 2020.