Received: 14 September 2022 | Revised: 24 March 2023 | Accepted: 20 April 2023

DOI: 10.1112/plms.12538

RESEARCH ARTICLE

Proceedings of the London Mathematical Society

Transversals in quasirandom latin squares

Sean Eberhard[1] | **Freddie Manners**[2] | **Rudi Mrazović**[3]

[1]Mathematical Sciences Research Centre, Queen's University Belfast, Belfast, UK

[2]Department of Mathematics, UC San Diego, La Jolla, California, USA

[3]Faculty of Science, Department of Mathematics, University of Zagreb, Zagreb, Croatia

Correspondence
Rudi Mrazović, Faculty of Science, Department of Mathematics, University of Zagreb, Zagreb, Croatia.
Email: Rudi.Mrazovic@math.hr

Funding information
Horizon 2020, Grant/Award Number: 803711; Royal Society; Croatian Science Foundation, Grant/Award Number: UIP-2017-05-4129; Sloan Fellowship

Abstract

A transversal in an $n \times n$ latin square is a collection of n entries not repeating any row, column, or symbol. Kwan showed that almost every $n \times n$ latin square has $\big((1+o(1))n/e^2\big)^n$ transversals as $n \to \infty$. Using a loose variant of the circle method we sharpen this to $(e^{-1/2}+o(1))n!^2/n^n$. Our method works for all latin squares satisfying a certain quasirandomness condition, which includes both random latin squares with high probability as well as multiplication tables of quasirandom groups.

MSC 2020
05B15 (primary)

1 | INTRODUCTION

A transversal in an $n \times n$ latin square is a set of n entries such that no two of them come from the same row or column or contain the same symbol.

Although there are examples of latin squares with no transversals (e.g., the multiplication table of $\mathbf{Z}/n\mathbf{Z}$ for n even), it is widely believed that these are rare. For example, a famous conjecture of Ryser claims that an $n \times n$ latin square contains a transversal whenever n is odd. In the same direction, Kwan [6] proved that a uniformly random $n \times n$ latin square has a transversal with probability $1-o(1)$. Moreover, he showed that, with probability $1-o(1)$, the number of transversals is $\big((1-o(1))n/e^2\big)^n$.

In this paper, we improve Kwan's result by finding the precise asymptotic of the number of transversals in a uniformly random latin square.

Theorem 1.1. *Let $\mathbb{L}$ be a uniformly random $n \times n$ latin square. Then $\mathbb{L}$ has $\big(e^{-1/2}+o(1)\big)n!^2/n^n$ transversals with probability $1-o(1)$ as $n \to \infty$.*

More generally, we find a (deterministic) *quasirandomness condition* for latin squares which is sufficient to guarantee this same asymptotic number of transversals.

Theorem 1.2. *There is a constant $\rho > 0$ such that the following holds. Let L be an $n \times n$ latin square which is $\mathcal{A}$-quasirandom with parameter ρ. Then L has $\left(e^{-1/2} + o(1)\right)n!^2/n^n$ transversals.*

The precise definition of "$\mathcal{A}$-quasirandom" is in terms of the spectral gap of some operator associated to L: see Definition 7.1. Despite the language, it is not actually obvious that a uniformly random $n \times n$ is quasirandom with high probability as $n \to \infty$, and hence that Theorem 1.2 implies Theorem 1.1. Indeed, it is incredibly delicate to prove *any* statistical properties of a uniform random latin square, for a number of reasons: the exact asymptotic count of $n \times n$ latin squares is not known; the latin square property is too rigid to make local changes; and no efficient way of sampling uniform random latin squares is known.

However, using a recent result of Kwan, Sah, Sawhney, and Simkin [7] we are indeed able to establish that a random latin square is $\mathcal{A}$-quasirandom with parameter $o(1)$, with high probability, and we can thus prove Theorem 1.1 as a consequence of Theorem 1.2.

Theorem 1.3. *Let L be a uniformly random $n \times n$ latin square. Then L is $\mathcal{A}$-quasirandom with parameter $o(1)$, with probability $1 - o(1)$ as $n \to \infty$.*

Somewhat opposite to random latin squares are latin squares that are multiplication tables of finite groups. In [3], we proved that as long as the underlying group satisfies a necessary condition to have at least one transversal, we have the count as in Theorem 1.2 with an extra factor equal to the size of the group's abelianization. For some groups, this result is implied by Theorem 1.2 and the following (easy) result.

Theorem 1.4. *Let G be a group and let L_G be the multiplication table of G. Then L_G is $\mathcal{A}$-quasirandom with parameter $1/D$, where D is the minimal dimension of a nontrivial linear representation of G.*

This shows that the $\mathcal{A}$-quasirandomness condition when restricted to group multiplication tables coincides with the usual notion of quasirandomness for groups due to Gowers [5]. Thus, together Theorems 1.2 and 1.4 recover the main result of [3] for sufficiently quasirandom groups.

There appears to be no single universal definition of a "quasirandom latin square," in the same way that there is no single definition of a "quasirandom set of integers". Instead there are various possible qualitatively inequivalent definitions, some more natural than others, and the correct choice depends on the problem at hand. For this reason we prefer to talk about a *quasirandomness condition* than about a "definition of quasirandomness," and we do not claim that the condition in Definition 7.1 is necessarily the correct one for other problems. In particular it is not directly related to the notion introduced in [1, 4], as that depends on some additional structure (namely, an ordering on the set of symbols) to which our condition is oblivious. See Section 7 for further remarks.

2 | OUTLINE

Our approach is analytical rather than combinatorial. Let X, Y, Z be n-element sets of rows, columns and symbols. We identify an $n \times n$ latin square L with a subset of $X \times Y \times Z$ satisfying the

latin square property, that is, every pair from $X \times Y$, $Y \times Z$ and $Z \times X$ is in exactly one triple from L. We let $L^2(X), L^2(Y), L^2(Z)$ denote the spaces of complex-valued functions on X, Y, Z (equipped with the standard hermitian inner product). The *latin square tensor* $\Lambda = \Lambda_{\mathsf{L}}$ is defined by

$$\Lambda(f, g, h) := \mathbf{E}_{(x,y,z)\in\mathsf{L}} f(x)g(y)h(z)$$

for $f \in L^2(X), g \in L^2(Y), h \in L^2(Z)$.

We stress that the latin square tensor Λ_{L} depends on L, but we will always just write Λ for brevity. We use the same notation for powers of L, in the following sense. If L_1 and L_2 are latin squares then $\mathsf{L}_1 \times \mathsf{L}_2$ is also a latin square, where $\mathsf{L}_1 \times \mathsf{L}_2$ is simply the cartesian product of L_1 and L_2 as subsets of $X_1 \times Y_1 \times Z_1$ and $X_2 \times Y_2 \times Z_2$. Accordingly the powers L^m are latin squares of order n^m for all $m \geqslant 0$, and if $f \in L^2(X^m), g \in L^2(Y^m), h \in L^2(Z^m)$ then we write

$$\Lambda(f, g, h) := \mathbf{E}_{(x,y,z)\in\mathsf{L}^m} f(x)g(y)h(z).$$

Of particular interest is the latin square L^n. We write S (or sometimes S_X to emphasize the domain, and similarly S_Y and S_Z) for the subset $S \subseteq X^n$ of all bijections $[n] \to X$. Then one can check that the number of transversals in L is

$$\Lambda(1_S, 1_S, 1_S)\frac{n^{2n}}{n!}.$$

Our goal is therefore to show that

$$\Lambda(1_S, 1_S, 1_S) = \big(e^{-1/2} + o(1)\big)\Big(\frac{n!}{n^n}\Big)^3, \tag{2.1}$$

provided that L satisfies an appropriate quasirandomness condition.

We approach (2.1) principally by studying how 1_S deviates from its density $n!/n^n$. We do this as follows. For any set $A \subseteq [m]$, we may identify $L^2(X^A)$ with the subspace of $L^2(X^m)$ consisting of functions $X^m \to \mathbf{C}$ that factor as $X^m \to X^A \to \mathbf{C}$; that is, functions $f(x_1, \dots, x_m)$ that only depend on $(x_i : i \in A)$. These spaces are nested: if $A \subseteq B$ then $L^2(X^A) \subseteq L^2(X^B)$. We write Q_A for the orthogonal projection $L^2(X^m) \to L^2(X^A)$ and P_A for the orthogonal projection

$$P_A : L^2(X^m) \to L^2(X^A) \cap \bigcap_{B\subsetneq A} L^2(X^B)^\perp. \tag{2.2}$$

Here $L^2(X^B)^\perp$ is the space of functions $f(x_1, \dots, x_m)$ orthogonal to functions depending only on $(x_i : i \in B)$, that is, such that $\mathbf{E}_{x_i : i\notin B} f(x_1, \dots, x_m) = 0$ for any choice of $(x_i : i \in B)$. Therefore, the space on the right-hand side of (2.2) is the space of functions depending only on $(x_i : i \in A)$ and such that $\mathbf{E}_{x_i} f(x_1, \dots, x_m) = 0$ for any $i \in A$.

The operators P_A, Q_A are related via inclusion–exclusion rules:

$$Q_A = \sum_{B\subseteq A} P_B,$$

$$P_A = \sum_{B\subseteq A} (-1)^{|A\setminus B|} Q_B.$$

Hence, we have a kind of "Fourier expansion"

$$f = \sum_{A \subseteq [m]} P_A f,$$

for any function $f \in L^2(X^m)$ (which is only truly a Fourier expansion if $n = 2$ and X^m is identified with $\mathbf{F}_2^m$). Applying this to $f = 1_S \in L^2(X^n)$,

$$1_S = \sum_{A \subseteq [n]} P_A 1_S.$$

By the discussion above, $P_A 1_S$ can be thought of as the component of 1_S that depends exactly on $(x_i : i \in A)$ (and is orthogonal to all functions depending only on variables in a strict subset of A). For example, $P_\emptyset 1_S$ is equal to the density $n!/n^n$.

The relevance of the P_A projections is that any latin square tensor $\Lambda_\square$ is diagonal with respect to this decomposition: that is,

$$\Lambda(1_S, 1_S, 1_S) = \sum_{A \subseteq [n]} \Lambda(P_A 1_S, P_A 1_S, P_A 1_S). \tag{2.3}$$

This is a consequence of the following lemma.

Lemma 2.1. *Let $f \in L^2(X^n), g \in L^2(Y^n), h \in L^2(Z^n)$ and $A, B, C \subseteq [n]$. Then*

$$\Lambda(P_A f, P_B g, P_C h) = 0$$

unless $A = B = C$.

Proof. Assume it is not the case that $A = B = C$. By symmetry we may assume $A \not\subseteq B$, say $i \in A \setminus B$. We may also assume $P_A f = f$, $P_B g = g$, $P_C h = h$, by replacing f, g, h with their images under P_A, P_B, P_C, respectively. Now consider

$$\Lambda(f, g, h) = \mathbf{E}_{(x,y,z) \in \square^n} f(x) g(y) h(z).$$

In particular, consider the average over the variables $(x_i, y_i, z_i) \in \square$. As $i \notin B$, there is no dependence on y_i, so it is equivalent by the latin square property to average over all $(x_i, z_i) \in X \times Z$. As $\mathbf{E}_{x_i} f(x_1, \dots, x_m) = 0$, it follows that $\Lambda(f, g, h) = 0$. □

We now divide up the sum (2.3) according to the size m of A.

2.1 | Major arcs

The terms in this decomposition where A is very sparse (of size up to $cn^{1/2}$) form the *major arcs*.

Theorem 2.2. *There is a constant $c > 0$ such that for $\log n < m \leqslant cn^{1/2}$,*

$$\sum_{\substack{A \subseteq [n] \\ |A| \leqslant m}} \Lambda(P_A 1_S,\ P_A 1_S,\ P_A 1_S) = \left(\frac{n!}{n^n}\right)^3 e^{-1/2}\big(1 + O(m^2/n)\big).$$

The proof is a mostly mechanical adaptation of [3, section 4], which did not use group theory in an essential way.

2.2 | Sparse minor arcs

The next range, the *sparse minor arcs*, concerns A of size up to cn for some small absolute constant c.

Theorem 2.3. *There is a constant $c > 0$ such that for $1 \leqslant m \leqslant cn$,*

$$\sum_{\substack{A \subseteq [n] \\ |A|=m}} \Lambda(|P_A 1_S|, |P_A 1_S|, |P_A 1_S|) \leqslant \left(\frac{n!}{n^n}\right)^3 O(1)^{m+n/m} (m/n)^{m/8}.$$

Note $|\Lambda(P_A 1_S, P_A 1_S, P_A 1_S)| \leqslant \Lambda(|P_A 1_S|, |P_A 1_S|, |P_A 1_S|)$ by the triangle inequality. The point is we have an exponential-in-m gain over the main term provided

$$m \log(n/m) > C'(m + n/m),$$

for some large enough $C' > 0$. This would be satisfied as long as

$$C(n/\log n)^{1/2} < m < \epsilon n, \tag{2.4}$$

for some large enough $C > 0$ and small enough $\epsilon > 0$.

We prove Theorem 2.3 by exhibiting a majorant for $|P_A 1_S|$ and then using generating function methods.

2.3 | Dense minor arcs

Finally, we have the dense range, where $m \geqslant cn$. Here we use quasirandomness. To be precise we define a certain Markov chain on $X \times Y$, with adjacency operator $\mathcal{A}$, and we consider L to be $\mathcal{A}$-quasirandom with parameter ρ if $\mathcal{A}$ has a spectral gap at least $1 - \rho$, that is, if the spectral radius of $\mathcal{A} - \mathcal{U}$ is at most ρ, where $\mathcal{U}$ is the projection to constants (the uniform distribution). See Definition 7.1.

Theorem 2.4. *For every $\epsilon > 0$ there is $\rho > 0$ such that if L is $\mathcal{A}$-quasirandom with parameter ρ, then*

$$\sum_{\substack{A \subseteq [n] \\ |A| \geqslant \epsilon n}} |\Lambda(P_A 1_S, P_A 1_S, P_A 1_S)| \leqslant \left(\frac{n!}{n^n}\right)^3 10^{-n}.$$

2.4 | Quasirandomness

It remains (for Theorems 1.1 and 1.4) to demonstrate that the latin squares in scope are quasirandom in this sense. If L is the multiplication table of a group G we compute the entire spectrum of

$\mathcal{A}$ and find $\rho = 1/D$ where D is the minimal dimension of a nontrivial representation of G, which shows that our notion of quasirandomness is equivalent to the usual one due to Gowers [5] in the case of groups. For genuinely random latin squares we use recent work of Kwan, Sah, Sawhney, and Simkin [7] to show that $\operatorname{tr} \mathcal{A}^6 = 1 + o(1)$ with high probability, and this implies that $\rho = o(1)$.

2.5 | Proof of Theorem 1.2

Putting Theorems 2.2 to 2.4 together, it is straightforward to deduce Theorem 1.2.

Proof of Theorem 1.2. Let C and ϵ be as in (2.4) and $M := C(n/\log n)^{1/2}$. Theorems 2.2, 2.3, and 2.4 give us that for some $c > 0$

$$\begin{aligned}\Lambda(1_S, 1_S, 1_S) = \left(e^{-1/2} + O(M^2/n)\right)\left(\frac{n!}{n^n}\right)^3 \\ + \sum_{M < m \leqslant \epsilon n} O(e^{-cm})\left(\frac{n!}{n^n}\right)^3 \\ + 10^{-n}\left(\frac{n!}{n^n}\right)^3,\end{aligned}$$

as long as L is $\mathcal{A}$-quasirandom with parameter ρ for small enough ρ (depending on ϵ). The choice of M implies (2.1) and hence Theorem 1.2. □

2.6 | Layout of the paper

To prove Theorems 2.2 to 2.4, we need some background material on partition systems (Section 3.1) and on the primitive "Fourier analysis" of coordinate projections Q_A, P_A discussed above (Section 4). This builds on similar material from [3].

Then, Sections 5 to 7 give the proofs of the three key theorems above. Finally, Section 8 proves the quasirandomness properties from Subsection 2.4.

3 | PARTITIONS AND PARTITION SYSTEMS

3.1 | Partitions

Most of our language relating to the partition lattice is standard.

(1) If A is a set, Π_A is the set of all partitions of A. If $A = [m]$, we will conserve brackets by writing simply Π_m.
(2) $\Pi_A^{(k)}$ is the set of partitions all of whose cells have size at most k.
(3) If $A \subseteq B$, then any partition of A is identified with a partition of B by adding singletons $\big\{\{b\} : b \in B \setminus A\big\}$. With this convention, $\Pi_A \subseteq \Pi_B$.
(4) The *support* $\operatorname{supp} \pi$ of a partition $\pi \in \Pi_A$ is the union of the nonsingleton cells of π. It is the smallest set $B \subseteq A$ such that $\pi \in \Pi_B$.

(5) Π'_A is the set of $\pi \in \Pi_A$ with $\operatorname{supp} \pi = A$.
(6) If $\pi, \pi' \in \Pi_A$, $\pi \leqslant \pi'$ means that π is a *refinement* of π' (i.e., every cell of π' is a union of cells of π). Synonymously, π' is a *coarsening* of π.
(7) The *meet* $\pi \wedge \pi'$ is the coarsest partition refining both π and π'; the *join* $\pi \vee \pi'$ is the finest partition coarsening both π and π'.
(8) The partition $\{\{a\} : a \in A\}$ is the *discrete partition*; the partition $\{A\}$ is the *indiscrete partition*.
(9) The *rank* of a partition $\pi \in \Pi_A$ is $\operatorname{rank}(\pi) = |A| - |\pi|$; equivalently it is the greatest r such that there are partitions $\pi_0 < \pi_1 < \cdots < \pi_r = \pi$. (Note that $\operatorname{rank}(\pi)$ is meaningful without specifying A, unlike $|\pi|$; that is, it is invariant under adding or removing singletons.)
(10) The *Möbius function* μ at $\pi \in \Pi_A$ is given by $\mu(\pi) = (-1)^{\operatorname{rank}(\pi)} \prod_{p \in \pi} (|p| - 1)!$.
(11) A function $f : A \to X$ is *π-measurable* if f is constant on the cells of π. A subset $S \subseteq A$ is called π-measurable if 1_S is π-measurable.
(12) If $\pi \in \Pi_A$, $c_\pi \in L^2(X^A)$ is the indicator of π-measurability (i.e., $c_\pi(f)$ is 1 if f is π-measurable, and 0 otherwise).

The *exponential formula* for partitions states

$$\sum_{m \geqslant 0} \frac{1}{m!} \sum_{\pi \in \Pi_m} \prod_{p \in \pi} x_{|p|} = \exp\left(\sum_{k \geqslant 1} \frac{1}{k!} x_k \right). \tag{3.1}$$

Here $x_1, x_2, \ldots$ are formal variables. We will apply (3.1) several times in Section 6.

3.2 | Partition systems

In Sections 5 and 6, it will be essential to have good bounds on the quantity $\Lambda(c_{\pi_1}, c_{\pi_2}, c_{\pi_3})$ for $A \subseteq [n]$ and various choices $\pi_1, \pi_2, \pi_3 \in \Pi_A$. This motivates the following definitions.

(1) A *partition triple* on a set A is a triple $\mathfrak{P} = (\pi_1, \pi_2, \pi_3) \in \Pi_A^3$.
(2) We call $\mathfrak{P}$ a *partition system* if $\operatorname{supp} \pi_1 = \operatorname{supp} \pi_2 = \operatorname{supp} \pi_3$.
(3) The *support* of $\mathfrak{P}$ is $\operatorname{supp} \mathfrak{P} = \operatorname{supp} \pi_1 \cup \operatorname{supp} \pi_2 \cup \operatorname{supp} \pi_3$.

Definition 3.1 (Combinatorial rank). Let $\mathfrak{P} = (\pi_1, \pi_2, \pi_3) \in \Pi_A^3$ be a partition triple. We write $S \subseteq \mathfrak{P}$ to mean that $S \subseteq \pi_1 \sqcup \pi_2 \sqcup \pi_3$, that is, S is a collection of cells labeled 1, 2, or 3. A subset $S \subseteq \mathfrak{P}$ is *closed* (with respect to $\mathfrak{P}$) if whenever $p_i \in \pi_i$ for $i = 1, 2, 3$ and $p_1 \cap p_2 \cap p_3 \neq \emptyset$, if two of p_1, p_2, p_3 are in S then so is the third. The *closure* $\langle S \rangle$ of S is the intersection of all closed sets containing S. The *combinatorial rank* of $\mathfrak{P} = (\pi_1, \pi_2, \pi_3)$ is defined as

$$\operatorname{crank}(\mathfrak{P}) = 2|A| - \min\{|S| : S \subseteq \mathfrak{P}, \langle S \rangle = \mathfrak{P}\}.$$

The motivation for combinatorial rank is the following bound.

Lemma 3.2. *For a set A, partitions $\pi_1, \pi_2, \pi_3 \in \Pi_A$, and latin square $\mathsf{L} \subseteq X \times Y \times Z$,*

$$0 \leqslant \Lambda(c_{\pi_1}, c_{\pi_2}, c_{\pi_3}) \leqslant n^{-\operatorname{crank}(\pi_1, \pi_2, \pi_3)}.$$

The idea of the proof is the same as for the related result [3, Lemma 4.6].

Proof. The Λ value is, by definition, $n^{-2|A|}$ times the number of triples of functions

$$f_1 : A \to X, \qquad f_2 : A \to Y, \qquad f_3 : A \to Z$$

such that f_i is π_i-measurable for $i = 1, 2, 3$ and such that $(f_1(a), f_2(a), f_3(a)) \in \mathsf{L}$ for all $a \in A$. Note we can think of f_i as a function on the cells of π_i, as it is π_i-measurable.

We claim that, given $S \subseteq \mathfrak{P}$ with $\langle S \rangle = \mathfrak{P}$, the triple (f_1, f_2, f_3) is determined by the values of f_i on cells in S. Hence, the number of such triples is at most $n^{|S|}$, giving the result.

Indeed, suppose f_1', f_2', f_3' is another triple of measurable functions with the same restriction to S. Let $W \subseteq \mathfrak{P}$ be the set of all cells $p_i \in \pi_i$ such that $f_i|_{p_i} = f_i'|_{p_i}$. By hypothesis $W \supseteq S$. If $p_i \in \pi_i$ for $i = 1, 2, 3$, $a \in p_1 \cap p_2 \cap p_3$, and two of p_1, p_2, p_3 are in W, then so is the third, as the triples $(f_1(a), f_2(a), f_3(a)), (f_1'(a), f_2'(a), f_3'(a)) \in \mathsf{L}$ agree at two coordinates and so are equal by the latin square property. Hence, W is a closed set, so $W \supseteq \langle S \rangle = \mathfrak{P}$ and $f_i = f_i'$, as required. □

This reduces the problem of bounding $\Lambda(c_{\pi_1}, c_{\pi_2}, c_{\pi_3})$ from above to the problem of bounding $\operatorname{crank}(\pi_1, \pi_2, \pi_3)$ from below. In [3], we did this using two slightly weaker notions of rank, called *triple rank* and *lower rank*, defined, respectively, as

$$\operatorname{trank}(\mathfrak{P}) = \max_{\sigma \in S_3} \big(\operatorname{rank}(\pi_{\sigma(1)}) + \operatorname{rank}(\pi_{\sigma(2)} \vee \pi_{\sigma(3)})\big)$$

$$\operatorname{lrank}(\mathfrak{P}) = \big(\operatorname{rank}(\pi_1) + \operatorname{rank}(\pi_2) + \operatorname{rank}(\pi_3) + \operatorname{rank}(\pi_1 \vee \pi_2 \vee \pi_3)\big)/2.$$

Lemma 3.3. $\operatorname{crank}(\mathfrak{P}) \geqslant \operatorname{trank}(\mathfrak{P}) \geqslant \operatorname{lrank}(\mathfrak{P})$.

Proof. For the first inequality, let $S \subseteq \mathfrak{P}$ contain all of π_1 and one cell of π_2 from each cell of $\pi_2 \vee \pi_3$. Then $|S| = |\pi_1| + |\pi_2 \vee \pi_3|$ and $\langle S \rangle = \mathfrak{P}$, so $\operatorname{crank}(\mathfrak{P}) \geqslant \operatorname{rank}(\pi_1) + \operatorname{rank}(\pi_2 \vee \pi_3)$, and equally for other permutations of 1, 2, 3. The second inequality was proved in [3, Lemma 4.8], and in any case will not be used in this paper. □

For continuity with [3], we define the *complexity* of a partition system $\mathfrak{P}$ to be

$$\operatorname{cx}(\mathfrak{P}) = \operatorname{trank}(\mathfrak{P}) - |\operatorname{supp} \mathfrak{P}|.$$

The complexity of a partition system is nonnegative, and it is zero if and only if $\mathfrak{P} = (\pi, \pi, \pi)$ for some *matching* π, that is, a partition of $A = \operatorname{supp} \pi$ into $|A|/2$ pairs.

3.3 | Combinatorial rank of matching systems

In this subsection, we compute $\operatorname{crank}(\pi_1, \pi_2, \pi_3)$ for all $(\pi_1, \pi_2, \pi_3) \in \Pi_A^{(2)}$, that is, partition triples such that all cells of π_1, π_2, π_3 have size at most 2. Where it applies, this is a significant improvement on what Lemma 3.3 gives us.

Lemma 3.4. *Let $\pi_1, \pi_2, \pi_3 \in \Pi_A^{(2)}$. Suppose there are precisely k cells $p \in \pi_1 \vee \pi_2 \vee \pi_3$ such that $\pi_i|_p$ has full support (i.e., is a matching) for each $i \in [3]$. Then*

$$\operatorname{crank}(\pi_1, \pi_2, \pi_3) = \operatorname{rank}(\pi_1) + \operatorname{rank}(\pi_2) + \operatorname{rank}(\pi_3) - k.$$

Proof. As all terms are additive across cells of $\pi_1 \vee \pi_2 \vee \pi_3$, we may assume $\pi_1 \vee \pi_2 \vee \pi_3$ is indiscrete. In particular, $k \in \{0, 1\}$, and $k = 0$ if and only if one of π_1, π_2, π_3 has a singleton.

Case $k = 1$: In this case π_1, π_2, π_3 are matchings, so

$$\operatorname{rank}(\pi_1) = \operatorname{rank}(\pi_2) = \operatorname{rank}(\pi_3) = |A|/2,$$

and we must show

$$\operatorname{crank}(\pi_1, \pi_2, \pi_3) = 3|A|/2 - 1.$$

Let $\mathcal{G}$ be the multigraph whose vertex set is A and with edges given by the cells of π_1, π_2, π_3 (which are all 2-cells). Clearly $\mathcal{G}$ is 3-regular, with $|A|$ vertices and $3|A|/2$ edges. As $\pi_1 \vee \pi_2 \vee \pi_3$ is indiscrete, $\mathcal{G}$ is connected.

According to Definition 3.1, we want to infect as few edges as possible in such a way that, if two infected edges incident at a vertex always spread infection to the third edge, then infection spreads to all edges. Note that for this to happen, it is necessary and sufficient to infect at least one edge in each cycle, as the edges that are uninfected at the end of the process form a subgraph with no vertex of degree 1. Hence, equivalently, we want to delete as few edges as possible to get a forest.

As $\mathcal{G}$ has $3|A|/2$ edges and any forest has at most $|A| - 1$ edges, we must delete at least $|A|/2 + 1$ edges. Conversely, given any connected 3-regular multigraph, we can delete edges until we have a (simple) tree. Hence, the minimal number of generators is precisely $|A|/2 + 1$, so $\operatorname{crank}(\pi_1, \pi_2, \pi_3) = 2|A| - (|A|/2 + 1) = 3|A|/2 - 1$, as claimed.

Case $k = 0$: In this case, at least one of π_1, π_2, π_3 has a singleton, and we must show that

$$\operatorname{crank}(\pi_1, \pi_2, \pi_3) = \operatorname{rank}(\pi_1) + \operatorname{rank}(\pi_2) + \operatorname{rank}(\pi_3).$$

We define a graph $\mathcal{G}$ as in the previous case but additionally for every singleton $\{v\} \in \pi_1 \sqcup \pi_2 \sqcup \pi_3$ we add an edge $\{v, *\}$, where $*$ is a special additional vertex at which infection does not spread. As $\pi_1 \vee \pi_2 \vee \pi_3$ is indiscrete, $\mathcal{G} \setminus *$ is connected. As there is at least one singleton, $\mathcal{G}$ is connected. Again we want to delete as few edges as possible to get a forest. The number of vertices in $\mathcal{G}$ is $|A| + 1$ and the number of edges is $|\pi_1| + |\pi_2| + |\pi_3|$, so the number of edges we must delete is precisely $|\pi_1| + |\pi_2| + |\pi_3| - |A|$. Hence,

$$\operatorname{crank}(\pi_1, \pi_2, \pi_3) = 3|A| - |\pi_1| - |\pi_2| - |\pi_3| = \operatorname{rank}(\pi_1) + \operatorname{rank}(\pi_2) + \operatorname{rank}(\pi_3),$$

as claimed. □

4 | THE "FOURIER" EXPANSION OF 1_S

Recall from Section 2 that Q_A denotes the orthogonal projection $L^2(X^m) \to L^2(X^A)$ and P_A denotes the orthogonal projection

$$P_A : L^2(X^m) \to L^2(X^A) \cap \bigcap_{B \subsetneq A} L^2(X^B)^\perp,$$

and these operators are related via inclusion–exclusion rules:

$$Q_A = \sum_{B \subseteq A} P_B,$$
$$P_A = \sum_{B \subseteq A} (-1)^{|A \setminus B|} Q_B. \tag{4.1}$$

In this section, we study the terms in the expansion

$$1_S = \sum_{A \subseteq [n]} P_A 1_S.$$

To express some of the results, it is convenient to use the linear map $U : \mathbf{C}[z] \to \mathbf{C}$ defined by

$$U(z^k) = \begin{cases} n^k/(n)_k & : k \leqslant n, \\ 0 & : k > n. \end{cases}$$

Here $(n)_k = n(n-1)\cdots(n-k+1)$.

4.1 | Formulae for $P_A 1_S$

Let $S_A \subseteq X^A$ denote the set of injections $A \to X$. Thus, if $|A| = m$, $|S_A| = (n)_m$.

Lemma 4.1. *If $A \subseteq [n]$ and $|A| = m$,*

$$Q_A 1_S = \frac{n!}{n^n} \frac{n^m}{(n)_m} 1_{S_A}.$$

Proof. A function $f : A \to X$ can be extended to a bijection $[n] \to X$ in $(n-m)!$ ways if f is injective and 0 ways otherwise, and by definition $Q_A 1_S(f)$ is the number of such extensions normalized by $1/n^{n-m}$. □

Lemma 4.2 [3, Lemma 4.3].

$$1_{S_A} = \sum_{\pi \in \Pi_A} \mu(\pi) c_\pi.$$

Lemma 4.3. *If $A \subseteq [n]$ and $|A| = m$, then*

$$P_A 1_S = \frac{n!}{n^n} \frac{n^m}{(n)_m} \sum_{\pi \in \Pi'_A} \mu(\pi) P_A c_\pi.$$

Proof. Combining the previous two lemmas,

$$P_A 1_S = P_A Q_A 1_S = \frac{n!}{n^n} \frac{n^m}{(n)_m} \sum_{\pi \in \Pi_A} \mu(\pi) P_A c_\pi.$$

As $c_\pi \in L^2(X^{\operatorname{supp}\pi})$, only the terms with $\operatorname{supp}\pi = A$ survive. □

We can use U to give another formula for $P_A 1_S$. If $x \in X^A$ (i.e., $x : A \to X$), the *kernel* $\ker x \in \Pi_A$ of x is the level set partition

$$\ker x = \left\{x^{-1}(t) : t \in X, x^{-1}(t) \neq \emptyset\right\}.$$

Note that

$$c_\pi(x) = 1 \iff x \text{ is } \pi\text{-measurable} \iff \pi \leqslant \ker x.$$

Lemma 4.4. *Let $A \subseteq [n]$, $|A| = m$. For $x \in X^n$, let $\pi = \ker(x|_A) \in \Pi_A$. Then*

$$P_A 1_S(x) = (-1)^{\operatorname{rank}(\pi)} \frac{n!}{n^n} U \prod_{p \in \pi} (|p|z - 1).$$

Proof. From (4.1) and Lemma 4.1, we have

$$P_A 1_S = \frac{n!}{n^n} \sum_{B \subseteq A} (-1)^{|A \setminus B|} \frac{n^{|B|}}{(n)_{|B|}} 1_{S_B}.$$

Now, the sets B such that $x|_B$ is injective are precisely those which intersect each cell of π in at most one point. Hence,

$$\begin{aligned} P_A 1_S(x) &= \frac{n!}{n^n} U \sum_{B \subseteq A} (-1)^{|A \setminus B|} z^{|B|} 1_{S_B}(x) \\ &= \frac{n!}{n^n} (-1)^{|A| - |\pi|} U \prod_{p \in \pi} (|p|z - 1). \end{aligned}$$

□

4.2 | Sparseval

The word *sparseval* is our playful term for the computation of $\|P_A f\|_2^2$ for any $A \subseteq [n]$. This is possible by inclusion–exclusion and orthogonality: as

$$\|Q_A f\|_2^2 = \sum_{B \subseteq A} \|P_B f\|_2^2,$$

it follows that

$$\|P_A f\|_2^2 = \sum_{B \subseteq A} (-1)^{|A \setminus B|} \|Q_B f\|_2^2. \tag{4.2}$$

Lemma 4.5. *If $A \subseteq [n]$ and $|A| = m$,*

$$\|P_A 1_S\|_2^2 = \left(\frac{n!}{n^n}\right)^2 U\left((z-1)^m\right).$$

Proof. Note that $\|1_{S_B}\|_2^2 = (n)_{|B|}/n^{|B|}$ for every $B \subseteq A$. Hence, from (4.2) and Lemma 4.1,

$$\|P_A 1_S\|_2^2 \Big(\frac{n!}{n^n}\Big)^{-2} = \sum_{B\subseteq A} (-1)^{|A\setminus B|} \frac{n^{|B|}}{(n)_{|B|}} = U\big((z-1)^m\big).$$

□

Proposition 4.6. *Assume $0 \leqslant m \leqslant n$ and let $t = m/n$. Then*

$$0 \leqslant U\big((z-1)^m\big) \ll \binom{n}{m}^{-1} e^{s(t)n},$$

where

$$s(t) = t^{1/2} - t\log t^{1/2} - (1-t)\log(1+t^{1/2}).$$

In particular,

$$U\big((z-1)^m\big) \leqslant e^{O(m)}(m/n)^{m/2}.$$

Sketch. The inequality $U\big((z-1)^m\big) \geqslant 0$ follows from the previous lemma. For the main claim, by expanding we have

$$U\big((z-1)^m\big) = \frac{1}{n!}\sum_{k=0}^{m} \binom{m}{k} (-1)^{m-k} n^k (n-k)!,$$

and this can be identified as $\binom{n}{m}^{-1}$ times the coefficient of X^m in $e^{nX}/(1+X)^{n-m+1}$. The stated bound follows by taking a contour integral (chosen in the spirit of the saddle-point method) to extract the coefficient. For details, see [2, bound for the sum in (5.4)]. Extra care is needed for t near 1, but we omit the details because we will not use the claim for $t > 1/2$. The second bound follows by Stirling's formula. □

The following corollary will not be used but is included for interest.

Corollary 4.7. *The sign of $P_A 1_S(x)$ is $(-1)^{\operatorname{rank}(\ker x|_A)}$.*

Proof. Let $\pi = \ker(x|_A)$. By Lemma 4.4, it suffices to prove that

$$U\prod_{p\in\pi}(|p|z-1) > 0.$$

There are nonnegative integers $r_\omega \geqslant 0$ such that

$$\prod_{p\in\pi}(|p|z-1) = \prod_{p\in\pi}\big(|p|(z-1)+(|p|-1)\big) = \sum_{\omega\subseteq\pi} r_\omega (z-1)^{|\omega|}.$$

Hence, the claim follows from $U\big((z-1)^m\big) \geqslant 0$. □

5 | MAJOR ARCS

The goal in this section is to prove Theorem 2.2. Define

$$\mathfrak{S}_m = \sum_{2k \leqslant m} \frac{(-1)^k}{2^k k!},$$

$$M_m = \sum_{\substack{A \subseteq [n] \\ |A| \leqslant m}} \Lambda(P_A 1_S, P_A 1_S, P_A 1_S).$$

Our aim is to prove that, for $m \leqslant cn^{1/2}$,

$$M_m = \big(\mathfrak{S}_m + O(m^2/n)\big)\Big(\frac{n!}{n^n}\Big)^3. \tag{5.1}$$

In particular, this implies Theorem 2.2.

5.1 | The quantities γ and γ_0

From Lemma 4.3, it is clear that to estimate $\Lambda(P_A 1_S, P_A 1_S, P_A 1_S)$ it suffices to estimate $\Lambda(P_A c_{\pi_1}, P_A c_{\pi_2}, P_A c_{\pi_3})$ for every partition system $\mathfrak{P} = (\pi_1, \pi_2, \pi_3)$ with support A and aggregate the results with the appropriate weighting. For continuity with [3, section 4], we define the normalized quantities

$$\gamma_0(\mathfrak{P}) = n^{\operatorname{trank}(\mathfrak{P})} \Lambda(c_{\pi_1}, c_{\pi_2}, c_{\pi_3})$$

and

$$\gamma(\mathfrak{P}) = n^{\operatorname{trank}(\mathfrak{P})} \Lambda(P_A c_{\pi_1}, P_A c_{\pi_2}, P_A c_{\pi_3})$$

for any partition triple $\mathfrak{P} = (\pi_1, \pi_2, \pi_3)$. Note that

$$0 \leqslant \gamma_0(\mathfrak{P}) \leqslant 1$$

by Lemmas 3.2 and 3.3. As $c_\pi \in L^2(X^{\operatorname{supp}\pi})$, $\gamma(\mathfrak{P}) = 0$ unless $\mathfrak{P}$ is a partition system.

Lemma 5.1. *Let $\mathfrak{P}$ be a partition system with support $\operatorname{supp}\mathfrak{P} = A$ of size m, and suppose m' points of A are contained in cells $\pi_1 \vee \pi_2 \vee \pi_3$ of size at least 3. Then*

$$|\gamma(\mathfrak{P})| \leqslant 2^{m'}.$$

Sketch. The idea is that

$$\begin{aligned}\Lambda(P_A c_{\pi_1}, P_A c_{\pi_2}, P_A c_{\pi_3}) &= \Lambda(c_{\pi_1}, c_{\pi_2}, P_A c_{\pi_3}) \\ &= \sum_{B \subseteq A} (-1)^{|A \setminus B|} \Lambda(c_{\pi_1}, c_{\pi_2}, Q_B c_{\pi_3}) \\ &= \sum_{B \subseteq A} (-1)^{|A \setminus B|} \Lambda(Q_B c_{\pi_1}, Q_B c_{\pi_2}, Q_B c_{\pi_3}),\end{aligned}$$

and

$$Q_B c_\pi = n^{-\operatorname{rank}(\pi)+\operatorname{rank}(\pi|_B)} c_{\pi|_B},$$

where

$$\pi|_B = \{p \cap B : p \in \pi, p \cap B \neq \emptyset\}.$$

Let $\mathfrak{P}|_B = (\pi_1|_B, \pi_2|_B, \pi_3|_B)$. Then, normalizing,

$$\gamma(\mathfrak{P}) = \sum_{B \subseteq A} (-1)^{|A\setminus B|} \gamma_0(\mathfrak{P}|_B) n^{-t(\mathfrak{P},B)},$$

where

$$t(\mathfrak{P}, B) = \operatorname{trank}(\mathfrak{P}|_B) - \operatorname{trank}(\mathfrak{P}) + \sum_{i=1}^{3} \big(\operatorname{rank}(\pi_i) - \operatorname{rank}(\pi_i|_B)\big).$$

In [3, section 4] we showed $t(\mathfrak{P}, B) \geqslant 0$. As $\gamma_0(\mathfrak{P}|_B) \in [0,1]$ for all B this shows $|\gamma(\mathfrak{P})| \leqslant 2^m$. The stronger bound with m' in place of m follows by separating off the doubleton cells of $\pi_1 \vee \pi_2 \vee \pi_3$. See [3, section 4] for details. □

5.2 | The $M_m(z)$ series

For a partition triple $\mathfrak{P} = (\pi_1, \pi_2, \pi_3)$ we use the shorthand

$$\mu(\mathfrak{P}) = \mu(\pi_1)\,\mu(\pi_2)\,\mu(\pi_3).$$

From Lemma 4.3 we have

$$M_m = \left(\frac{n!}{n^n}\right)^3 \sum_{|\operatorname{supp}\mathfrak{P}| \leqslant m} \left(\frac{n^{|\operatorname{supp}\mathfrak{P}|}}{(n)_{|\operatorname{supp}\mathfrak{P}|}}\right)^3 \mu(\mathfrak{P})\gamma(\mathfrak{P}) n^{-\operatorname{trank}(\mathfrak{P})},$$

where the sum is over all partition systems on $[n]$. For $z \in \mathbf{C}$ define

$$M_m(z) = \left(\frac{n!}{n^n}\right)^3 \sum_{|\operatorname{supp}\mathfrak{P}| \leqslant m} \left(\frac{n^{|\operatorname{supp}\mathfrak{P}|}}{(n)_{|\operatorname{supp}\mathfrak{P}|}}\right)^3 \mu(\mathfrak{P})\gamma(\mathfrak{P}) n^{-|\operatorname{supp}\mathfrak{P}|} z^{\operatorname{cx}(\mathfrak{P})}.$$

As we have mentioned, $\operatorname{cx}(\mathfrak{P}) \geqslant 0$ for any partition system $\mathfrak{P}$, so $M_m(z)$ is a polynomial such that $M_m = M_m(1/n)$. By bounding $M_m(z)$ and using some complex analysis we will show $M_m(1/n) \approx M_m(0)$, and then we will directly estimate $M_m(0)$.

Proposition 5.2. *There is a constant $c > 0$ such that, for $|z|^{1/2} \leqslant c/m$, we have*

$$|M_m(z)| \ll \left(\frac{n^m}{(n)_m}\right)^2 \left(\frac{n!}{n^n}\right)^3.$$

Proof. By the definition of $M_m(z)$, the triangle inequality, and Lemma 5.1, the quantity $|M_m(z)|/(\frac{n!}{n^n})^3$ is bounded by

$$\sum_{|A|\leqslant m} n^{-|A|}\left(\frac{n^{|A|}}{(n)_{|A|}}\right)^3 \sum_{\operatorname{supp}\mathfrak{P}=A} 2^{m'(\mathfrak{P})}|\mu(\mathfrak{P})|\,|z|^{\operatorname{cx}(\mathfrak{P})},$$

where $m'(\mathfrak{P})$ is the number of points of $\operatorname{supp}\mathfrak{P}$ contained in cells of $\pi_1\vee\pi_2\vee\pi_3$ of size at least 3. This exact sum was analyzed in [3, section 4.4], and we showed that it is $O(n^m/(n)_m)^2$ provided $|z|^{1/2}<c/m$. The proposition follows. □

Corollary 5.3. *There is a constant $c>0$ such that, for $m<cn^{1/2}$,*

$$|M_m - M_m(0)| \ll (m^2/n)\left(\frac{n!}{n^n}\right)^3.$$

Proof. By the residue theorem,

$$M_m(u) - M_m(0) = \frac{1}{2\pi i}\oint_{|z|=R}\frac{M_m(z)u}{(z-u)z}\,dz$$

as long as $|u|<R$. Hence,

$$\big|M_m(u) - M_m(0)\big| \leqslant \max_{|z|=R}|M_m(z)|\frac{|u|/R}{1-|u|/R}.$$

Take $u=1/n$ and $R=c^2/m^2$, where c is as in the previous proposition. Then as long as $1/n<c^2/m^2$, that is, $m<cn^{1/2}$, we get

$$|M_m - M_m(0)| \ll (1+n^{-1/2})^{2m}\left(\frac{n^m}{(n)_m}\right)^2\left(\frac{n!}{n^n}\right)^3\frac{m^2/n}{1-m^2/(c^2n)}.$$

Hence, as long as say $m<(c/2)n^{1/2}$ we get the claimed bound. □

5.3 | The constant term $M_m(0)$

By definition,

$$M_m(0) = \left(\frac{n!}{n^n}\right)^3 \sum_{\substack{|\operatorname{supp}\mathfrak{P}|\leqslant m\\ \operatorname{cx}(\mathfrak{P})=0}}\left(\frac{n^{|\operatorname{supp}\mathfrak{P}|}}{(n)_{|\operatorname{supp}\mathfrak{P}|}}\right)^3\mu(\mathfrak{P})\gamma(\mathfrak{P})n^{-|\operatorname{supp}\mathfrak{P}|}.$$

As remarked, $\operatorname{cx}(\mathfrak{P})=0$ if and only if $\mathfrak{P}=(\pi,\pi,\pi)$ for some matching π. In this case, if say $|\operatorname{supp}\mathfrak{P}|=2k$,

$$\frac{n^{|\operatorname{supp}\mathfrak{P}|}}{(n)_{|\operatorname{supp}\mathfrak{P}|}} = \frac{n^{2k}}{(n)_{2k}} = 1+O(k^2/n),$$

$$\mu(\mathfrak{P}) = \mu(\pi)^3 = (-1)^k,$$

$$\gamma(\mathfrak{P}) = (1-1/n)^k.$$

The last identity holds by a direct calculation analogous to [3, Lemma 4.10]. The number of matchings π in $[n]$ of support size $2k$ is

$$\frac{(n)_{2k}}{2^k k!} = \frac{n^{2k}}{2^k k!}\big(1 + O(k^2/n)\big).$$

Thus,

$$\begin{aligned} M_m(0) &= \left(\frac{n!}{n^n}\right)^3 \sum_{k=0}^{\lfloor m/2 \rfloor} \frac{n^{2k}}{2^k k!}(-1)^k n^{-2k}\big(1 + O(k^2/n)\big) \\ &= \left(\frac{n!}{n^n}\right)^3 \big(\mathfrak{S}_m + O(1/n)\big). \end{aligned}$$

By combining with Corollary 5.3, we have

$$M_m = \left(\frac{n!}{n^n}\right)^3 \big(\mathfrak{S}_m + O(m^2/n)\big)$$

provided $m < cn^{1/2}$. This finishes the proof of (5.1).

6 | SPARSE MINOR ARCS

To prove Theorem 2.3, we need a bound on $\Lambda(|P_A 1_S|, |P_A 1_S|, |P_A 1_S|)$ for larger $|A|$. Note that in any latin square $\mathsf{L}' \subseteq (X', Y', Z')$,

$$\begin{aligned} |\Lambda(f, g, h)| &= |\mathbf{E}_{(x,y,z)\in \mathsf{L}'} f(x)g(y)h(z)| \\ &\leqslant \mathbf{E}_{x\in X'} |f(x)| \Big|\mathbf{E}_{y,z\,:\,(x,y,z)\in \mathsf{L}'} g(y)h(z)\Big| \leqslant \|f\|_1 \|g\|_2 \|h\|_2 \end{aligned} \tag{6.1}$$

using the latin square property and Cauchy–Schwarz, and similarly permuting f, g, h. One approach to Theorem 2.3 might be to find upper bounds on $|P_A 1_S(x)|$, pointwise or in L^1, and simply apply (6.1). However, by itself this approach is too crude, even assuming optimal upper bounds.

Another idea is to seek a majorant for $|P_A 1_S|$ of the form

$$|P_A 1_S| \leqslant \sum_{\pi\in\Pi_A} t_\pi c_\pi \tag{6.2}$$

for some coefficients $t_\pi \geqslant 0$. Then

$$\Lambda(|P_A 1_S|, |P_A 1_S|, |P_A 1_S|) \leqslant \sum_{\pi_1,\pi_2,\pi_3\in\Pi_A} t_{\pi_1} t_{\pi_2} t_{\pi_3} \Lambda(c_{\pi_1}, c_{\pi_2}, c_{\pi_3})$$

and Lemma 3.2, together with generating function techniques, gives a way to control the right-hand side. This bound is particularly effective if $\pi_i \in \Pi_A^{(2)}$, given Lemma 3.4.

Again this approach does not succeed by itself. Our final argument works by decomposing $|P_A 1_S|$ into two pieces and combining the two techniques discussed above.

6.1 | A majorant for $|P_A 1_S|$

Throughout this section, let $C > 0$ be some large enough constant, $A \subseteq [n]$ and $|A| = m \leqslant n/C$. Additionally, we let

$$\delta := (Cm/n)^{1/2}.$$

Although we will always have this specific value of δ in mind, most of the results in this section only rely on $\delta \leqslant 1$. The next proposition gives a useful bound for $|P_A 1_S|$. For $\delta := (Cm/n)^{1/2}$, $r \geqslant 1$, and π a partition define

$$\sigma_r^{(\delta)} = \begin{cases} \delta & : r = 1, \\ r - 1 & : r > 1, \end{cases}$$
$$\sigma_\pi^{(\delta)} = \prod_{p \in \pi} \sigma_{|p|}^{(\delta)}.$$

Proposition 6.1. *We have*

$$|P_A 1_S(x)| \leqslant \frac{n!}{n^n} e^{\delta m} \sigma_{\ker x}^{(\delta)} \qquad (x \in X^A).$$

Proof. From Lemma 4.4,

$$P_A 1_S(x) = (-1)^{\operatorname{rank}(\pi)} \frac{n!}{n^n} U\phi,$$

where $\pi = \ker x$ and

$$\phi = \prod_{p \in \pi} (|p| z - 1) = \sum_{\omega \subseteq \pi} r_\omega (z-1)^{|\omega|},$$
$$r_\omega = \prod_{p \in \pi \setminus \omega} (|p| - 1) \prod_{p \in \omega} |p|.$$

From Proposition 4.6 and crude estimates (Stirling's formula), for $0 \leqslant d \leqslant m$,

$$U\big((z-1)^d\big) \leqslant (Cd/n)^{d/2} \leqslant (Cm/n)^{d/2} = \delta^d$$

provided C is large enough. Then

$$U\phi \leqslant \sum_{\omega \subseteq \pi} r_\omega \delta^{|\omega|}$$
$$= \prod_{p \in \pi} (|p| - 1 + |p|\delta)$$

$$= \sigma_\pi^{(\delta)} \prod_{p\in\pi\colon |p|>1} \left(1 + \frac{|p|}{|p|-1}\delta\right)$$
$$\leqslant \sigma_\pi^{(\delta)}(1+2\delta)^{m/2}$$
$$\leqslant \sigma_\pi^{(\delta)} e^{\delta m}$$

as required. □

In light of the proposition, to find majorants for $|P_A 1_S|$ of the form (6.2) it suffices to find analogous bounds for $\sigma_\pi^{(\delta)}$. Recall that $\Pi_A^{(k)}$ is the set of all $\pi \in \Pi_A$ having no part of size greater than k. Let $r_k(\pi)$ be the number of k-cells in π and let $r_{3+}(\pi) = \sum_{k\geqslant 3} r_k(\pi)$.

Lemma 6.2. *Let π be a partition.*

(1)
$$\sigma_\pi^{(\delta)} \leqslant \sum \left\{ \sigma_{\pi'}^{(\delta)} : \pi' \leqslant \pi, \pi' \in \Pi^{(3)} \right\}.$$

(2)
$$\sigma_\pi^{(\delta)} \leqslant \sum \left\{ \sigma_{\pi'}^{(\delta)} : \pi' \leqslant \pi, \pi' \in \Pi^{(4)}, r_{3+}(\pi') = r_{3+}(\pi) \right\}.$$

(3)
$$\sigma_\pi^{(\delta)} \leqslant \delta^{-r_{3+}(\pi)} \sum \left\{ \sigma_{\pi'}^{(\delta)} : \pi' \leqslant \pi, \pi' \in \Pi^{(2)} \right\}.$$

Proof. Consider the first inequality. Both sides are multiplicative across cells of π, so we may assume π is a single cell, say of size r. The inequality is trivial for $r \leqslant 3$ (as $\sigma_\pi^{(\delta)}$ is one of the summands on the right-hand side), so we may assume $r \geqslant 4$. Then it suffices to check

$$r - 1 \leqslant \sum_{2a+3b=r} \frac{r!}{2!^a a! 3!^b b!} 2^b.$$

This is a calculation for $r \leqslant 10$ (say) and an uninteresting exercise for $r > 10$.

Now consider the second inequality. This time, the right-hand side is not itself multiplicative over cells of π, but if we replace the condition $r_{3+}(\pi') = r_{3+}(\pi)$ by the stronger one

$$\forall p \in \pi,\ |p| \geqslant 3 :\ \text{there is exactly one } p' \in \pi' \text{ with } p' \subseteq p \text{ and } |p'| \geqslant 3$$

then it becomes so, and it suffices to prove the corresponding stronger inequality. Now we may again assume that π is an r-cell, and we may assume $r \geqslant 5$. Then we must check

$$r - 1 \leqslant \sum_{\substack{2a+3b+4c=r \\ b+c=1}} \frac{r!}{2!^a a! 3!^b b! 4!^c c!} 2^b 3^c.$$

Again we omit further details.

Now consider the third inequality. Again it suffices to consider the case of an r-cell, and we may assume $r \geqslant 3$. Then the assertion is

$$r - 1 \leqslant \delta^{-1} \sum_{a+2b=r} \frac{r!}{a!2!^b b!} \delta^a.$$

As $\delta \leqslant 1$, it suffices to check

$$r - 1 \leqslant \sum_{\substack{a+2b=r \\ a \leqslant 1}} \frac{r!}{a!2!^b b!},$$

which is again essentially a calculation. □

Lemma 6.3. *Let $x \in X^A$. Then*

$$|P_A 1_{S_A}(x)| \leqslant \frac{n!}{n^n} e^{\delta m} \sum_{\pi \in \Pi_A^{(3)}} \sigma_\pi^{(\delta)} c_\pi(x).$$

Proof. By Proposition 6.1,

$$|P_A 1_{S_A}(x)| \leqslant \frac{n!}{n^n} e^{\delta m} \sigma_{\ker x}^{(\delta)}.$$

By Lemma 6.2(1),

$$\sigma_{\ker x}^{(\delta)} \leqslant \sum \left\{ \sigma_\pi^{(\delta)} : \pi \leqslant \ker x, \pi \in \Pi_A^{(3)} \right\} = \sum_{\pi \in \Pi_A^{(3)}} \sigma_\pi^{(\delta)} c_\pi(x).$$

□

6.2 | A splitting of $|P_A 1_S|$

We can use the bound on $|P_A 1_S|$ given in the previous section to bound the L^1 norm of $P_A 1_S$, but the bound would not be strong enough for what we need. To go further, we break up $R := |P_A 1_S|$ into two parts, a part whose L^1 norm we can control better, and a part we can analyze separately. Fix $\epsilon \geqslant 0$ and let

$$\Pi^\sharp = \{\pi \in \Pi_A : r_{3+}(\pi) < \epsilon m\}.$$

Let $\Pi^\flat = \Pi_A \setminus \Pi^\sharp$. Define

$$R^\sharp(x) = 1_{\Pi^\sharp}(\ker x) R(x),$$

$$R^\flat(x) = 1_{\Pi^\flat}(\ker x) R(x).$$

Clearly, $R = R^\sharp + R^\flat$.

Lemma 6.4. *We have*

$$R^\flat \leqslant \frac{n!}{n^n} e^{\delta m} \sum_{\pi \in \Pi^\flat \cap \Pi^{(4)}} \sigma_\pi^{(\delta)} c_\pi,$$

$$R^\sharp \leqslant \frac{n!}{n^n} e^{\delta m} \sum_{\pi \in \Pi^\sharp \cap \Pi^{(4)}} \sigma_\pi^{(\delta)} c_\pi,$$

$$R^\sharp \leqslant \frac{n!}{n^n} e^{\delta m} \delta^{-\epsilon m} \sum_{\pi \in \Pi^{(2)}} \sigma_\pi^{(\delta)} c_\pi.$$

Proof. By Proposition 6.1,

$$R(x) \leqslant \frac{n!}{n^n} e^{\delta m} \sigma^{(\delta)}_{\ker x}.$$

Suppose $\ker x \in \Pi^\flat$. Then by Lemma 6.2(2),

$$\sigma_{\ker x} \leqslant \sum \{\sigma_\pi : \pi \leqslant \ker x, \pi \in \Pi^{(4)}, r_{3+}(\pi) \geqslant \epsilon m\}.$$

This proves the bound on $R^\flat$. The first bound on $R^\sharp$ is proved identically. The second is proved in the same way using instead Lemma 6.2(3). □

Corollary 6.5. *We have*

$$\|R^\flat\|_1 \ll \frac{n!}{n^n} e^{O(m)} (m/n)^{(1+\epsilon)m/2}.$$

Proof. Using the previous lemma, $\delta \leqslant 1$ and $\|c_\pi\|_1 = n^{-\operatorname{rank}(\pi)}$,

$$\|R^\flat\|_1 \leqslant \frac{n!}{n^n} e^m \sum_{\pi \in \Pi^\flat \cap \Pi^{(4)}} \sigma_\pi^{(\delta)} n^{-\operatorname{rank}(\pi)}.$$

Let

$$\alpha_r(x, w) = \sum_{\pi \in \Pi_r^{(4)}} \sigma_\pi^{(\delta)} x^{\operatorname{rank}(\pi)} w^{r_{3+}(\pi)}.$$

Then, for real $w \geqslant 1$,

$$\sum_{\pi \in \Pi^\flat \cap \Pi^{(4)}} \sigma_\pi^{(\delta)} n^{-\operatorname{rank}(\pi)} \leqslant w^{-\epsilon m} \alpha_m(1/n, w).$$

Using the exponential formula (3.1) with $x_k = \sigma_k^{(\delta)} x^{k-1} y^k$ for $k = 1, 2$, $x_k = \sigma_k^{(\delta)} w x^{k-1} y^k$ for $k = 3, 4$, and $x_k = 0$ for $k \geqslant 5$, we obtain

$$\sum_{r \geqslant 0} \frac{1}{r!} \alpha_r(x, w) y^r = \exp(\delta y + xy^2/2 + wx^2y^3/3 + wx^3y^4/8).$$

Hence, for real $y > 0$,

$$w^{-\epsilon m}\alpha_m(x,w) \leqslant \frac{m!}{w^{\epsilon m}y^m}\exp(\delta y + xy^2/2 + wx^2y^3/3 + wx^3y^4/8).$$

Putting $x = 1/n$, $y = (mn)^{1/2}$, and $w = (n/m)^{1/2}$, we get

$$w^{-\epsilon m}\alpha_m(1/n,w) \leqslant \frac{m!}{(n/m)^{\epsilon m/2}(mn)^{m/2}}e^{O(m)}.$$

This proves what we want. □

Corollary 6.6. *We have*

$$\Lambda(R,R,R) \leqslant \Lambda(R^\sharp,R^\sharp,R^\sharp) + \left(\frac{n!}{n^n}\right)^3 O(1)^m (m/n)^{(1+\epsilon/2)m}.$$

Proof. We have $\|R^\sharp\|_2 \leqslant \|R\|_2$ because $0 \leqslant R^\sharp \leqslant R$ pointwise. Hence, from (6.1),

$$\begin{aligned}\Lambda(R,R,R) &= \Lambda(R^\flat,R,R) + \Lambda(R^\sharp,R^\flat,R) + \Lambda(R^\sharp,R^\sharp,R^\flat) + \Lambda(R^\sharp,R^\sharp,R^\sharp)\\ &\leqslant \Lambda(R^\sharp,R^\sharp,R^\sharp) + 3\|R\|_2^2\|R^\flat\|_1.\end{aligned}$$

From sparseval (Lemma 4.5 and Proposition 4.6),

$$\|R\|_2^2 \ll \left(\frac{n!}{n^n}\right)^2 e^{O(m)}(m/n)^{m/2}.$$

Combining with Corollary 6.5 gives the bound. □

6.3 | The contribution from $R^\sharp$

Finally, we must bound $\Lambda(R^\sharp,R^\sharp,R^\sharp)$. From Lemma 6.4,

$$R^\sharp \leqslant \frac{n!}{n^n}e^{\delta m}\delta^{-\epsilon m}Q \leqslant \frac{n!}{n^n}e^{O(m)}(m/n)^{-\epsilon m/2}Q, \tag{6.3}$$

where

$$Q = \sum_{\pi\in\Pi^{(2)}}\sigma_\pi^{(\delta)}c_\pi.$$

Hence, it suffices to bound $\Lambda(Q,Q,Q)$. The key ingredient for this is the knowledge of the exact value of combinatorial rank for $\pi_1,\pi_2,\pi_3 \in \Pi^{(2)}$ (Lemma 3.4).

Lemma 6.7.

$$\Lambda(Q,Q,Q) \leqslant (m/n)^{3m/2}e^{O(m+n/m)}.$$

Proof. Let $\mathcal{M}_A \subseteq \Pi_A^{(2)}$ be the set of matchings (partitions all of whose cells have size 2). For $\pi_1, \pi_2, \pi_3 \in \Pi_A^{(2)}$, let $k(\pi_1, \pi_2, \pi_2)$ be the number of cells $p \in \pi_1 \vee \pi_2 \vee \pi_3$ such that $\pi_i|_p \in \mathcal{M}_p$ for each $i \in [3]$. Then, from Lemmas 3.2 and 3.4,

$$\Lambda(c_{\pi_1}, c_{\pi_2}, c_{\pi_3}) \leqslant n^{k(\pi_1,\pi_2,\pi_3)-\mathrm{rank}(\pi_1)-\mathrm{rank}(\pi_2)-\mathrm{rank}(\pi_3)}.$$

Hence,

$$\begin{aligned}\Lambda(Q,Q,Q) &= \sum_{\pi_1,\pi_2,\pi_3 \in \Pi_A^{(2)}} \sigma_{\pi_1}^{(\delta)}\sigma_{\pi_2}^{(\delta)}\sigma_{\pi_3}^{(\delta)}\Lambda(c_{\pi_1}, c_{\pi_2}, c_{\pi_3}) \\ &\leqslant \sum_{\pi_1,\pi_2,\pi_3 \in \Pi_A^{(2)}} \sigma_{\pi_1}^{(\delta)}\sigma_{\pi_2}^{(\delta)}\sigma_{\pi_3}^{(\delta)} n^{k(\pi_1,\pi_2,\pi_3)-\mathrm{rank}(\pi_1)-\mathrm{rank}(\pi_2)-\mathrm{rank}(\pi_3)} \\ &= \sum_{\pi \in \Pi_A} \prod_{p \in \pi} \sum_{\substack{\pi_1,\pi_2,\pi_3 \in \Pi_p^{(2)} \\ \pi_1 \vee \pi_2 \vee \pi_3 = \{p\}}} n^{k(\pi_1,\pi_2,\pi_3)} \prod_{i \in [3]} \sigma_{\pi_i}^{(\delta)} n^{-\mathrm{rank}(\pi_i)}.\end{aligned}$$

In the last sum above, as $\pi_1 \vee \pi_2 \vee \pi_3 = \{p\}$, $k(\pi_1, \pi_2, \pi_3)$ is 0 or 1 according to whether $\pi_1, \pi_2, \pi_3 \in \mathcal{M}_p$. Splitting the sum according to these cases,

$$\Lambda(Q,Q,Q) \leqslant \sum_{\pi \in \Pi_A} \prod_{p \in \pi} \left(\sum_{\substack{\pi_1,\pi_2,\pi_3 \in \Pi_p^{(2)} \\ \pi_1 \vee \pi_2 \vee \pi_3 = \{p\}}} \prod_{i \in [3]} \sigma_{\pi_i}^{(\delta)} n^{-\mathrm{rank}(\pi_i)} + \sum_{\substack{\pi_1,\pi_2,\pi_3 \in \mathcal{M}_p \\ \pi_1 \vee \pi_2 \vee \pi_3 = \{p\}}} n \prod_{i \in [3]} \sigma_{\pi_i}^{(\delta)} n^{-\mathrm{rank}(\pi_i)} \right).$$

In the second sum, we will ignore the constraint $\pi_1 \vee \pi_2 \vee \pi_3 = \{p\}$; in the first sum we will use only $\mathrm{rank}(\pi_1) + \mathrm{rank}(\pi_2) + \mathrm{rank}(\pi_3) \geqslant \mathrm{rank}(\pi_1 \vee \pi_2 \vee \pi_3) = |p| - 1$.

Fix parameters $w_r \geqslant 1$ for all $r \geqslant 1$. Define

$$\begin{aligned}\alpha_r(x) &= \sum_{\pi \in \Pi_r^{(2)}} \sigma_\pi^{(\delta)} x^{\mathrm{rank}(\pi)}, \\ \alpha'_r(x) &= \sum_{\pi \in \mathcal{M}_r} \sigma_\pi^{(\delta)} x^{\mathrm{rank}(\pi)} = |\mathcal{M}_r| x^{r/2}, \\ \beta_r(x) &= \sum_{\pi \in \Pi_r} \prod_{p \in \pi} \left(w_{|p|}^{-(|p|-1)} \alpha_{|p|}(w_{|p|}x)^3 + x^{-1}\alpha'_{|p|}(x)^3 \right).\end{aligned}$$

Then, by the discussion above,

$$\Lambda(Q,Q,Q) \leqslant \beta_m(1/n).$$

Three applications of the exponential formula (3.1) give

$$\sum_{r \geqslant 0} \frac{y^r}{r!} \alpha_r(x) = \exp(\delta y + xy^2/2), \tag{6.4}$$

$$\sum_{r\geqslant 0}\frac{y^r}{r!}\alpha_r'(x)=\exp(xy^2/2), \tag{6.5}$$

$$\sum_{r\geqslant 0}\frac{y^r}{r!}\beta_r(x)=\exp\left(\sum_{r\geqslant 1}\frac{y^r w_r^{-r+1}\alpha_r(w_r x)^3}{r!}+\sum_{r\geqslant 2\text{ even}}\frac{y^r x^{-1}\alpha_r'(x)^3}{r!}\right). \tag{6.6}$$

From (6.4), for real $y>0$,

$$\alpha_r(x)\leqslant\frac{r!}{y^r}\exp(\delta y+xy^2/2).$$

Replacing x with $w_r x$, putting $w_r=\delta^2/(xr)$ (we will ensure later that $w_r\geqslant 1$ for $1\leqslant r\leqslant m$) and $y=r/\delta$ gives

$$w_r^{-r+1}\alpha_r(w_r x)^3\leqslant e^{O(r)}r^r\delta^{r+2}x^{r-1}.$$

From (6.5) with $y=(r/x)^{1/2}$, we have

$$\alpha_r'(x)\leqslant\frac{r!}{y^r}\exp(xy^2/2)\asymp r^{1/2}(rx/e)^{r/2}$$

(alternatively, this follows directly from $\alpha_r'(x)=|\mathcal{M}_r|x^{r/2}$). Hence, from (6.6) for $x,y>0$,

$$\beta_m(x)\leqslant\frac{m!}{y^m}\exp b(x,y), \tag{6.7}$$

where b is the truncated sum

$$\begin{aligned}b(x,y)&=\sum_{r=1}^m\frac{y^r w_r^{-r+1}\alpha_r(w_r x)^3}{r!}+\sum_{r=2}^m\frac{y^r x^{-1}\alpha_r'(x)^3}{r!}\\&\ll\sum_{r=1}^m e^{O(r)}\delta^{r+2}x^{r-1}y^r+\sum_{r=2}^m r^{O(1)}(e^{-1/2}r^{1/2}x^{3/2}y)^r x^{-1}.\end{aligned}$$

Inserting $x=1/n$ and $\delta=(Cm/n)^{1/2}$,

$$b(1/n,y)\ll\sum_{r=1}^m O(m^{1/2}y/n^{3/2})^r m+\sum_{r=2}^m r^{O(1)}(e^{-1/2}r^{1/2}y/n^{3/2})^r n.$$

Note that $w_r=Cm/r$, and this is indeed at least 1 for $r\leqslant m$ because we may assume $C\geqslant 1$. Finally, let $y=cn^{3/2}/m^{1/2}$ for a sufficiently small constant $c>0$. Then

$$b(1/n,y)\ll m+n/m.$$

Hence, from (6.7),

$$\Lambda(Q,Q,Q)\leqslant\beta_m(1/n)\leqslant\frac{m!}{y^m}\exp b(1/n,y)\ll(m/n)^{3m/2}e^{O(m+n/m)},$$

as claimed. □

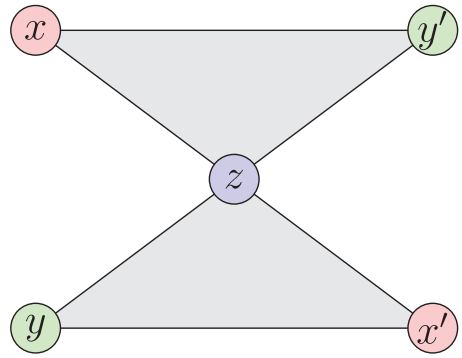

FIGURE 1 A transition $(x, y) \mapsto (x', y')$ in the Markov chain.

Putting the last few results together, we have the following theorem, which clearly implies Theorem 2.3.

Theorem 6.8. *We have*

$$\Lambda(|P_A 1_S|, |P_A 1_S|, |P_A 1_S||) \leqslant \left(\frac{n!}{n^n}\right)^3 (m/n)^{9m/8} e^{O(m+n/m)}.$$

Proof. From Corollary 6.6,

$$\Lambda(R, R, R) \leqslant \Lambda(R^\sharp, R^\sharp, R^\sharp) + \left(\frac{n!}{n^n}\right)^3 e^{O(m)} (m/n)^{(1+\epsilon/2)m}.$$

By (6.3) and the previous lemma, the main term is

$$\begin{aligned}\Lambda(R^\sharp, R^\sharp, R^\sharp) &\leqslant \left(\frac{n!}{n^n}\right)^3 e^{O(m)} (m/n)^{-3\epsilon m/2} \Lambda(Q, Q, Q) \\ &\leqslant \left(\frac{n!}{n^n}\right)^3 (m/n)^{(1-\epsilon)3m/2} e^{O(m+n/m)}.\end{aligned}$$

Set $\epsilon = 1/4$. □

7 | DENSE MINOR ARCS

Define a Markov chain on $X \times Y$ as follows. If the current state is (x, y), pick uniformly at random $z \in Z$. The next state is (x', y'), where x' and y' are the unique solutions to

$$(x, y', z), (x', y, z) \in \mathsf{L}$$

(see Figure 1). Let $\mathcal{A}$ be the transition operator for this Markov chain:

$$\mathcal{A}(f)(x, y) = \frac{1}{n} \sum_{(x,y',z),(x',y,z)\in\mathsf{L}} f(x', y').$$

The Markov chain is reversible with uniform stationary distribution, so $\mathcal{A}$ is self-adjoint and has the constant function on $X \times Y$ as a 1-eigenvector. Let $\mathcal{U}$ be the projection to constants:

$$\mathcal{U}(f)(x,y) = \frac{1}{n^2}\sum_{x',y'} f(x',y').$$

Definition 7.1. We say L is *$\mathcal{A}$-quasirandom* with parameter ρ if $\mathcal{A} - \mathcal{U}$ has spectral radius at most ρ.

In particular, $\rho < 1$ if and only if the Markov chain is connected, and in general ρ measures the rate of mixing.

Remark 7.2. For a finite set T, let $L^2(T)_0$ denote the subspace $\{f \in L^2(T) : \mathbf{E}f = 0\}$. Then equivalently, L is $\mathcal{A}$-quasirandom with parameter ρ if the restriction $\mathcal{A}|_{L^2(X\times Y)_0}$ has spectral radius at most ρ.

All our applications of quasirandomness go through the following lemma.

Lemma 7.3. *Assume L is $\mathcal{A}$-quasirandom with parameter ρ and*[†] *let $m \geqslant 1$. Then*

$$|\Lambda(f,g,h)| \leqslant \rho^{m/2}\|f\|_2\|g\|_2\|h\|_2 \tag{7.1}$$

for all $f \in L^2(X)_0^{\otimes m}$, $g \in L^2(Y)_0^{\otimes m}$, $h \in L^2(Z)_0^{\otimes m}$.

Remark 7.4. Identifying $L^2(X)^{\otimes m}$ with $L^2(X^m)$ in the usual way, $L^2(X)_0^{\otimes m}$ is identified with the subspace $\operatorname{im} P_{[m]} \subseteq L^2(X^m)$; see (2.2).

Proof of Lemma 7.3. By Cauchy–Schwarz,

$$\begin{aligned}
|\Lambda(f,g,h)| &= \left|\mathbf{E}_{(x,y,z)\in \mathsf{L}^m} f(x)g(y)h(z)\right| \\
&\leqslant \left(\mathbf{E}_z\left|\mathbf{E}_{x,y:(x,y,z)\in\mathsf{L}^m} f(x)g(y)\right|^2\right)^{1/2}\|h\|_2 \\
&= \left(\mathbf{E}_{z,x,y,x',y':(x,y,z),(x',y',z)\in\mathsf{L}^m} f(x)g(y)\bar{f}(x')\bar{g}(y')\right)^{1/2}\|h\|_2 \\
&= \langle \mathcal{A}^{\otimes m}(f\otimes\bar{g}), f\otimes\bar{g}\rangle^{1/2}\|h\|_2.
\end{aligned}$$

Note $\|f\otimes\bar{g}\|_2 = \|f\|_2\|g\|_2$, and that $f\otimes\overline{g} \in L^2(X\times Y)_0^{\otimes m}$. As $\mathcal{A}|_{L^2(X\times Y)_0}$ has spectral radius at most ρ, the tensor power $\mathcal{A}^{\otimes m}|_{L^2(X\times Y)_0^{\otimes m}}$ has spectral radius (and hence operator norm) at most ρ^m, so the last expression above is bounded by $\rho^{m/2}\|f\|_2\|g\|_2\|h\|_2$. □

Remark 7.5. As stated in the introduction, while Definition 7.1 has some nice properties (e.g., the spectral radius of $\mathcal{A}-\mathcal{U}$ can be computed efficiently), it is chosen for mainly practical rather than

[†] Note that the $m = 1$ case of (7.1) does not obviously imply the general case: the operator-type norm for trilinear forms does not behave well under taking tensor powers.

philosophical reasons, and there are similar but qualitatively inequivalent conditions that would work equally well.

One notable criticism of this definition is that latin squares associated to Steiner triple systems (i.e., where $X = Y = Z$ and L contains the diagonal $\{(x, x, x) : x \in X\}$ and is invariant under the S_3-action on triples) always fail to be $\mathcal{A}$-quasirandom with parameter $\rho < 1$ (as the diagonal $\{(x, x) : x \in X\}$ of $X \times X$ is a closed set for the Markov chain). On the other hand, a random Steiner triple system is far from having algebraic structure and presumably satisfies (7.1) for $\rho = o(1)$ with high probability as $n \to \infty$.

One point of view is that (7.1) itself is the more natural quasirandomness condition (but harder to verify), and Definition 7.1 is a convenient sufficient condition.

Proof of Theorem 2.4. Let $A \subseteq [n]$ and $|A| = m$. By Lemma 7.3 and Remark 7.2,

$$|\Lambda(P_A 1_S, P_A 1_S, P_A 1_S)| \leqslant \rho^{m/2} \|P_A 1_S\|_2^3 \leqslant \rho^{m/2} \|1_S\|_2^3 = \rho^{m/2} \Big(\frac{n!}{n^n}\Big)^3.$$

Hence, for $\rho \leqslant 1$,

$$\sum_{|A| \geqslant m} |\Lambda(P_A 1_S, P_A 1_S, P_A 1_S)| \leqslant 2^n \rho^{m/2} \Big(\frac{n!}{n^n}\Big)^3.$$

Taking $m = \epsilon n$ and ρ so that $2\rho^{\epsilon/2} \leqslant 1/10$, the result follows. □

8 | QUASIRANDOMNESS

In this section, we will verify that two natural classes of latin squares are $\mathcal{A}$-quasirandom with parameter $o(1)$:

- multiplication tables of quasirandom groups;
- uniformly random $n \times n$ latin squares, with high probability as $n \to \infty$.

In the case of a group, we can compute the whole spectrum of $\mathcal{A}$ using representation theory. In the case of a random latin square, we will use the bound

$$1 + \rho^6 \leqslant \operatorname{tr} \mathcal{A}^6$$

which holds because the spectrum of $\mathcal{A}$ is real and 6 is even. By interpreting $n^6 \operatorname{tr} \mathcal{A}^6$ as counting certain kinds of configuration in L (and using a recent result of [7]) we will show that $\operatorname{tr} \mathcal{A}^6 = 1 + o(1)$ with high probability, which implies that $\rho = o(1)$. (Using the same method one can show that $\operatorname{tr} \mathcal{A}^4 = 3 + o(1)$ with high probability, so 6 is the smallest even integer that we can use for this argument.)

8.1 | Quasirandom groups

The following proposition shows that our quasirandomness condition generalizes the definition of a quasirandom group (see [5]), implying Theorem 1.4.

Proposition 8.1. *Suppose* $\llcorner$ *is the multiplication table of a group* G*. Then the spectrum of* $\mathcal{A}$ *consists of* $d^3(d+1)/2$ *copies of* $1/d$ *and* $d^3(d-1)/2$ *copies of* $-1/d$ *for every* d*-dimensional irreducible representation of* G*, and* $n^2 - \sum_{\chi\in\mathrm{Irr}(G)} \chi(1)^4$ *zeros. In particular,* $\rho = 1/D$ *where* D *is the minimal dimension of a nontrivial representation of* G*.*

Proof. Here $X = Y = Z = G$ and $\llcorner = \{(x,y,z) \in G^3 : xy = z\}$, so $L^2(X\times Y) = L^2(G\times G)$ and $\mathcal{A}$ is the operator defined by

$$\mathcal{A}(f)(x,y) = \frac{1}{n}\sum_{z\in G} f(zy^{-1}, x^{-1}z).$$

By representation theory, $L^2(G)$ has an orthogonal basis consisting of the functions of the form $x \mapsto \langle \rho(x)e_i, e_j\rangle$, where $\rho : G \to U(V)$ is an irreducible unitary representation of G and $e_1, \dots, e_{\dim V}$ is an orthonormal basis of V.

It follows that $L^2(G\times G) \cong L^2(G)\otimes L^2(G)$ has an orthogonal basis consisting of functions of the form

$$f_{\rho,\rho',i,j,k,\ell}(x,y) = \left\langle \rho(x)e_i, e_j\right\rangle \left\langle e'_\ell, \rho'(y)e'_k\right\rangle,$$

where $\rho : G \to U(V)$ and $\rho' : G \to U(V')$ are two irreducible unitary representations of G and $1 \leqslant i,j \leqslant \dim V$, $1 \leqslant k,\ell \leqslant \dim V'$.

To find $\mathcal{A}(f_{\rho,\rho',i,j,k,\ell})$, we recall the Schur orthogonality relation for matrix coefficients, which states that for irreducible V, V' as above, $a, b \in V$ and $a', b' \in V'$,

$$\frac{1}{n}\sum_{z\in G}\langle \rho(z)a, b\rangle\langle b', \rho'(z)a'\rangle = \begin{cases} 0 & : (\rho,V) \ncong (\rho',V') \\ \frac{1}{\dim V}\langle a,a'\rangle\langle b',b\rangle & : (\rho,V) = (\rho',V'), \end{cases}$$

and thereby compute

$$\begin{aligned}\mathcal{A}(f_{\rho,\rho',i,j,k,\ell})(x,y) &= \frac{1}{n}\sum_{z\in G}\left\langle \rho(z)\rho(y^{-1})e_i, e_j\right\rangle\left\langle \rho(x)e'_\ell, \rho(z)e'_k\right\rangle \\ &= \begin{cases} 0 & : (\rho,V) \ncong (\rho',V') \\ \frac{1}{\dim V}\left\langle \rho(x)e_\ell, e_j\right\rangle\left\langle e_i, \rho(y)e_k\right\rangle & : (\rho,V) = (\rho',V') \end{cases} \\ &= \begin{cases} 0 & : (\rho,V) \ncong (\rho',V') \\ \frac{1}{\dim V} f_{\rho,\rho,\ell,j,k,i}(x,y) & : (\rho,V) = (\rho',V'). \end{cases}\end{aligned}$$

In the case $\rho \neq \rho'$, we get an eigenfunction with eigenvalue 0. When $\rho = \rho'$ and $i = \ell$ we get a $(1/\dim V)$-eigenfunction. Finally, when $\rho = \rho'$ and $i \neq \ell$, the functions

$$f_{\rho,\rho,i,j,k,\ell} \pm f_{\rho,\rho,\ell,j,k,i}$$

are eigenfunctions of $\mathcal{A}$ with eigenvalues $\pm 1/\dim V$, respectively.

Altogether we have $d^3 + d^3(d-1)/2 = d^3(d+1)/2$ copies of $1/d$ and $d^3(d-1)/2$ copies of $-1/d$, and the rest 0, as claimed. □

8.2 | Random latin squares

We will use a recent result of Kwan, Sah, Sawhney, and Simkin [7] on configuration counts in random latin squares. A *triple system* is a 3-uniform 3-partite hypergraph $\mathsf{H} \subseteq X_\mathsf{H} \times Y_\mathsf{H} \times Z_\mathsf{H}$ with vertex classes $X_\mathsf{H}, Y_\mathsf{H}, Z_\mathsf{H}$. The number of vertices is $v = |X_\mathsf{H}| + |Y_\mathsf{H}| + |Z_\mathsf{H}|$ and the number of triples (hyperedges) is $e = |\mathsf{H}|$. We say H is *latin* if every pair of vertices is in at most one triple. (A latin square of order n is then a latin triple system with three classes of n vertices and n^2 triples.)

Let H be a fixed triple system. A *copy* of H in a triple system L is a triple of injective maps

$$X_\mathsf{H} \to X_\mathsf{L}, \qquad Y_\mathsf{H} \to Y_\mathsf{L}, \qquad Z_\mathsf{H} \to Z_\mathsf{L}$$

which maps triples to triples. Let $N_\mathsf{H}(\mathsf{L})$ denote the number of copies of H in L.

Let B_n denote the random triple system $\mathsf{B}_n \subseteq [n] \times [n] \times [n]$ in which each possible triple is present independently with probability $1/n$. Note that $\mathbf{E}[N_\mathsf{H}(\mathsf{B}_n)] = (1 - o(1))n^{v-e}$ (when H is fixed and n is large). We say H is *α-stable* if $\alpha \geqslant v - e$ and

$$\mathbf{E}[N_\mathsf{H}(\mathsf{B}_n) \mid \mathsf{Q} \subseteq \mathsf{B}_n] - \mathbf{E}[N_\mathsf{H}(\mathsf{B}_n)] = o(n^\alpha)$$

for any latin triple system $\mathsf{Q} \subseteq [n] \times [n] \times [n]$ with at most $n(\log n)^3$ triples.

Theorem 8.2 [7, Theorem 7.2]. *Fix an α-stable latin triple system H with v vertices and e triples. Let L be a uniformly random latin square. Then*

$$N_\mathsf{H}(\mathsf{L}) \leqslant n^{v-e} + o(n^\alpha)$$

with high probability as $n \to \infty$.

To use this theorem effectively, we need a computable form of stability. Let H be a latin triple system. A subset of the vertices $S \subseteq X_\mathsf{H} \cup Y_\mathsf{H} \cup Z_\mathsf{H}$ is called *closed* if whenever two vertices of a triple of H is in S, so is the third. The *closure* $\langle S \rangle_\mathsf{H}$ of a subset S if the smallest closed set containing it. If $\mathsf{F} \subseteq \mathsf{H}$ let $X_\mathsf{F}, Y_\mathsf{F}, Z_\mathsf{F}$ denote the vertices incident with at least one member of F, and let $v(\mathsf{F}) = |X_\mathsf{F}| + |Y_\mathsf{F}| + |Z_\mathsf{F}|$ and $e(\mathsf{F}) = |\mathsf{F}|$. We say $\mathsf{F} \subseteq \mathsf{H}$ *generates* H if

$$\langle X_\mathsf{F} \cup Y_\mathsf{F} \cup Z_\mathsf{F} \rangle_\mathsf{H} = X_\mathsf{H} \cup Y_\mathsf{H} \cup Z_\mathsf{H}.$$

Let

$$d(\mathsf{H}) = \min\{e(\mathsf{F}) : \mathsf{F} \text{ generates } \mathsf{H}\}.$$

For example, if H_1 is the latin triple system shown in Figure 2, one generating set consists of both triples containing z_1, one triple containing z_3, and one triple containing z_5, and there is no smaller generating set, so $d(\mathsf{H}_1) = 4$.

Lemma 8.3. *Let H be a latin triple system with v vertices and e triples. Then H is α-stable provided $\alpha \geqslant v - e$ and*

$$\alpha > v - e + \max_{\emptyset \neq \mathsf{F} \subseteq \mathsf{H}} \big(d(\mathsf{F}) - v(\mathsf{F}) + e(\mathsf{F})\big).$$

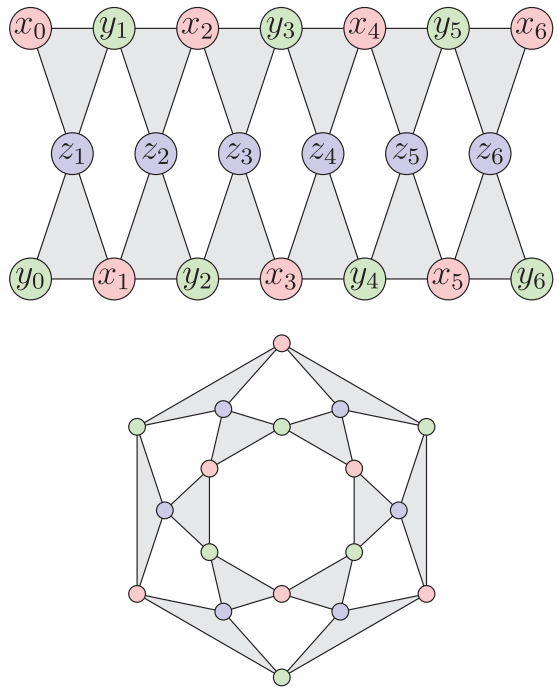

FIGURE 2 The chain $(x_0, y_0), \dots, (x_6, y_6)$ and the latin triple system H_1 defined by identifying x_0 with x_6 and y_0 with y_6.

Remark 8.4. A much simpler model problem is the following: given a fixed graph H and a random graph $G_{n,p}$, does G contain $n^{v(H)}p^{e(H)}(1 + o(1))$ copies of H (i.e., close to the expected number) with high probability? The answer might be no if H contains a subgraph H' with much greater density than H in some sense: indeed, if $n^{v(H')}p^{e(H')} = o(1)$ then with high probability $G(n, p)$ contains zero copies of H', and hence of H. However, this is essentially all that can go wrong. The condition for α-stability in the lemma captures a similar intuition.

Remark 8.5. Given a triple system $\mathsf{H} \subseteq X_{\mathsf{H}} \times Y_{\mathsf{H}} \times Z_{\mathsf{H}}$, one can construct a partition triple $\mathfrak{P} = (\pi_1, \pi_2, \pi_3) \in \Pi^3_{\mathsf{H}}$ in the sense of Subsection 3.1 (i.e., the ground set has size $e(\mathsf{H})$) where two triples $(x, y, z), (x', y', z') \in \mathsf{H}$ lie in the same cell of π_1 if and only if $x = x'$, and similarly for π_2 and $y = y'$, and π_3 and $z = z'$.

The construction can be reversed (up to the issue of repeated edges). In other words, triple systems and partition triples are more-or-less the same objects. Under this analogy, the notion of closure here coincides with that in Definition 3.1, and $\mathrm{crank}(\mathfrak{P}) = 2e(\mathsf{H}) - d(\mathsf{H})$.

Although using both languages is strictly speaking redundant, it is useful to keep the two notions separate, partly for minor technical reasons, but mainly because using partition systems follows our previous work in [2, 3] while using triple systems follows [7].

Proof of Lemma 8.3 [7, p. 15]. Let $\mathsf{Q} \subseteq [n]^3$ be a latin triple system with at most $n^{1+o(1)}$ triples. For a copy of H in B_n, say one of its triples is *forced* if it appears in Q. The difference

$$\mathbf{E}[N_{\mathsf{H}}(\mathsf{B}_n) \mid \mathsf{Q} \subseteq \mathsf{B}_n] - \mathbf{E}[N_{\mathsf{H}}(\mathsf{B}_n)] \tag{8.1}$$

arises from copies of H with at least one forced triple. Let $\mathsf{F} \subseteq \mathsf{H}$ be a nonempty subsystem and consider copies of H whose forced triples are precisely the images of those in F. Let $\mathsf{F}_0 \subseteq \mathsf{F}$ be a generating subsystem of size $d(\mathsf{F})$. Because Q satisfies the latin property, any copy of F in Q is determined by the image of F_0. Therefore, the number of copies of F in Q is at most $|\mathsf{Q}|^{|\mathsf{F}_0|}$. There are $v - v(\mathsf{F})$ vertices of H outside F, each with n possible images in $[n]^3$, and the image of each of the $e - e(\mathsf{F})$ triples outside F has probability $1/n$ (independently) of being present in B_n. Hence, the contribution to (8.1) from F is bounded by

$$|\mathsf{Q}|^{|\mathsf{F}_0|} n^{v-v(\mathsf{F})} (1/n)^{e-e(\mathsf{F})} = n^{v-e+d(\mathsf{F})-v(\mathsf{F})+e(\mathsf{F})+o(1)}.$$

This is $o(n^\alpha)$ provided the stated condition is satisfied. □

Now we can show that random latin squares are $\mathcal{A}$-quasirandom with parameter $o(1)$ with high probability (Theorem 1.3). This follows from the following proposition and the bound $1 + \rho^6 \leqslant \operatorname{tr} \mathcal{A}^6$.

Proposition 8.6. *For a uniformly random latin square* L,

$$\operatorname{tr} \mathcal{A}^6 = 1 + o(1)$$

with high probability as $n \to \infty$.

Proof (Computer-assisted). For $(x_0, y_0) \in X \times Y$, let (x_i, y_i) denote the iterates of (x_0, y_0) under the Markov chain defining $\mathcal{A}$. Then

$$\operatorname{tr} \mathcal{A}^6 = \sum_{x_0, y_0} \mathbf{P}\big((x_6, y_6) = (x_0, y_0)\big) = N/n^6,$$

where N is the number of configurations in L of the form shown in Figure 2 with $x_0 = x_6$ and $y_0 = y_6$. We do not assume the other vertices are distinct.

Let H_1 be the latin triple system depicted in Figure 2 and let $\mathsf{H}_2, \dots, \mathsf{H}_k$ (where k is bounded) be all the degenerations obtainable by identifying some (like-colored) vertices and identifying triangles as necessary to preserve the latin property.

Formally, we consider all triples of partitions (π_X, π_Y, π_Z) where $\pi_X \in \Pi_{X_{\mathsf{H}_1}}$, $\pi_Y \in \Pi_{Y_{\mathsf{H}_1}}$, $\pi_Z \in \Pi_{Z_{\mathsf{H}_1}}$ satisfying the following closure property: if (x, y, z) and (x', y', z') are two triples of H_1 and two of the pairs (x, x'), (y, y'), (z, z') are in the same cell of π_X, π_Y, π_Z, respectively, then so is the third. Number such triples of partitions $1, \dots, k$, where 1 corresponds to three copies of the discrete partition. Then H_i denotes the quotient hypergraph of H_1 with respect to partition i.

Let $N_i = N_{\mathsf{H}_i}(\mathsf{L})$. Then $N = N_1 + \cdots + N_k$. Let $v_i = v(\mathsf{H}_i)$ and $e_i = e(\mathsf{H}_i)$. Then $v_1 - e_1 = 18 - 12 = 6$. Now the proposition follows from Theorem 8.2, Lemma 8.3, and the following two claims:

(1) $v_i - e_i \leqslant 5$ for each $i > 1$,
(2) $v_i - e_i + \max_{\emptyset \neq \mathsf{F} \subseteq \mathsf{H}} \big(d(\mathsf{F}) - v(\mathsf{F}) + e(\mathsf{F})\big) \leqslant 5$ for each $i \geqslant 1$.

Indeed, provided (1) and (2) hold, Lemma 8.3 shows that H_i is 6-stable for each $i \geqslant 1$, so Theorem 8.2 implies that $N_i \leqslant n^{v_i - e_i} + o(n^6)$ with high probability for each i, so $N \leqslant (1 + o(1))n^6$ with high probability.

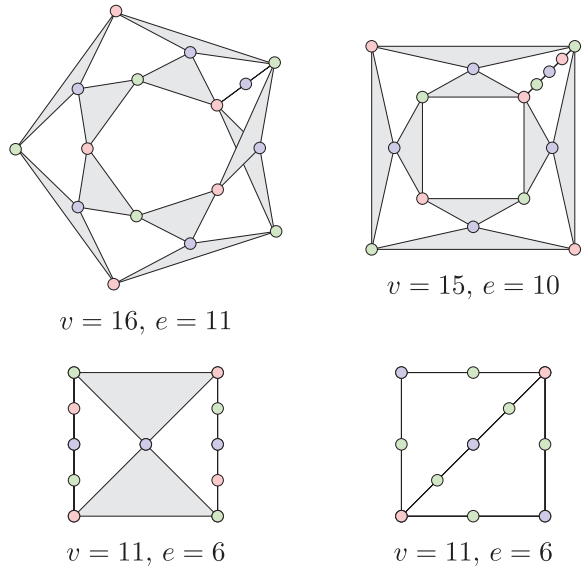

FIGURE 3 Degenerations of H_1 with $v_i - e_i = 5$ and $e_i < 12$. Some triangles are shown flat.

Both claims can be verified by exhaustive search. We find $\mathsf{H}_2, \dots, \mathsf{H}_k$ by starting with H_1 and iteratively identifying pairs of vertices, using breadth-first search. Thus, we verify (1). Now for each H_i we check all subsystems $\mathsf{F} \subseteq \mathsf{H}_i$ and compute $d(\mathsf{F})$ by checking all $\mathsf{F}_0 \subseteq \mathsf{F}$, and thus we verify (2).

It turns out $k = 1206$, and there are 154 distinct isomorphism classes among the degenerations H_i. The quantity in (2) turns out to be at most 4 in all cases except H_1, for which it is 5. There are just eight degenerations H_i (up to isomorphism) for which $v_i - e_i = 5$. Of these, four are just H_1 with a single pair of vertices identified (so $v_i = 17$ and $e_i = 12$). The other four cases are shown in Figure 3. These cases are therefore the dominant contributors to the error term. □

ACKNOWLEDGEMENTS

Sean Eberhard has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (Grant Agreement Number: 803711) and from the Royal Society. Rudi Mrazović is supported in part by the Croatian Science Foundation under the project UIP-2017-05-4129 (MUNHANAP). Freddie Manners is supported by a Sloan Fellowship.

JOURNAL INFORMATION

The *Proceedings of the London Mathematical Society* is wholly owned and managed by the London Mathematical Society, a not-for-profit Charity registered with the UK Charity Commission. All surplus income from its publishing programme is used to support mathematicians and mathematics research in the form of research grants, conference grants, prizes, initiatives for early career researchers and the promotion of mathematics.

REFERENCES

1. J. W. Cooper, D. Král', A. Lamaison, and S. Mohr, *Quasirandom Latin squares*, Random Structures Algorithms **61** (2022), no. 2, 298–308. MR4456029.

2. S. Eberhard, F. Manners, and R. Mrazović, *Additive triples of bijections, or the toroidal semiqueens problem*, J. Eur. Math. Soc. (JEMS) **21** (2019), no. 2, 441–463. MR3896207.
3. S. Eberhard, F. Manners, and R. Mrazović, *An asymptotic for the Hall–Paige conjecture*, Adv. Math. **404** (2022), no. part A, Paper No. 108423, 73. MR4416136.
4. F. Garbe, R. Hancock, J. Hladký, and M. Sharifzadeh, *Theory of limits of sequences of Latin squares*, Acta Math. Univ. Comenian. (N.S.) **88** (2019), no. 3, 709–716. MR4012871.
5. W. T. Gowers, *Quasirandom groups*, Combin. Probab. Comput. **17** (2008), no. 3, 363–387. MR2410393.
6. M. Kwan, *Almost all Steiner triple systems have perfect matchings*, Proc. Lond. Math. Soc. (3) **121** (2020), no. 6, 1468–1495. MR4144368.
7. M. Kwan, A. Sah, M. Sawhney, and M. Simkin, *Substructures in Latin squares*, arXiv:2202.05088, 2022.