

Almost all Steiner triple systems have perfect matchings

Matthew Kwan *

Abstract

We show that for any n divisible by 3, almost all order- n Steiner triple systems have a perfect matching (also known as a *parallel class* or *resolution class*). In fact, we prove a general upper bound on the number of perfect matchings in a Steiner triple system and show that almost all Steiner triple systems essentially attain this maximum. We accomplish this via a general theorem comparing a uniformly random Steiner triple system to the outcome of the triangle removal process, which we hope will be useful for other problems. Our methods can also be adapted to other types of designs; for example, we sketch a proof of the theorem that almost all Latin squares have transversals.

1 Introduction

A *Steiner triple system* of order n is a collection S of size-3 subsets of $[n] = \{1, \dots, n\}$ (that is, a *3-uniform hypergraph* on the vertex set $[n]$), such that every pair of vertices is included in exactly one hyperedge of S . Steiner triple systems are among the most fundamental types of combinatorial designs, and have strong connections to a wide range of different subjects, ranging from group theory, to finite geometry, to experimental design, to the theory of error-correcting codes. See [13] for an introduction to the subject. Observe that a Steiner triple system is actually nothing more than a partition of the edges of the complete graph K_n into triangles (a *triangle-decomposition* of K_n), so Steiner triple systems are natural “symmetric” counterparts to *Latin squares*, which can be defined as triangle-decompositions of the complete tripartite graph $K_{n,n,n}$.

In 1974 Wilson [48] used estimates for the number of Latin squares to prove a coarse estimate for the number of Steiner triple systems. Babai [4] used this estimate to prove that almost all Steiner triple systems have trivial automorphism group (that is to say, a uniformly random order- n Steiner triple system a.a.s.¹ has trivial automorphism group). We believe this is the only nontrivial property known to hold a.a.s. for random Steiner triple systems. Following Erdős and Rényi’s seminal paper [15] on random graphs and Erdős’ popularization of the probabilistic method, there have been great developments in the theory of random combinatorial structures of all kinds, but essentially none of the tools developed seem to be applicable to Steiner triple systems. Steiner triple systems lack independence or any kind of recursive structure, which rules out many of the techniques used to study Erdős–Rényi random graphs and random permutations, and there is basically no freedom to make local changes, which precludes the use of “switching” techniques often used in the study of random regular graphs (see for example [31]). It is not even clear how to study random Steiner triple systems empirically; in an attempt to find an efficient algorithm to generate a random Steiner triple system, Cameron [10] designed a Markov chain on Steiner triple systems, but he was not able to determine whether this chain was connected.

In a recent breakthrough, Keevash [26] proved that for a large class of combinatorial designs generalising Steiner triple systems, “partial” designs satisfying a certain “quasirandomness” condition can be completed into designs. Shortly afterwards [27], he showed that his results could

*Department of Mathematics, Stanford University, Stanford, CA 94305. Email: mattkwan@stanford.edu. This research was done while the author was working at ETH Zurich, and is supported in part by SNSF project 178493.

¹By “asymptotically almost surely”, or “a.a.s.”, we mean that the probability of an event is $1 - o(1)$. Here and for the rest of the paper, asymptotics are as $n \rightarrow \infty$.

be used for approximate enumeration; in particular, matching an upper bound due to Linial and Luria [33] he proved that there are

$$(n/e^2 + o(n))^{n^2/6} \tag{1}$$

Steiner triple systems of order n , as long as n satisfies a necessary divisibility condition (Steiner triple systems can only exist if n is 1 or 3 mod 6).

Of course, this new estimate makes it possible, in theory, to prove new a.a.s. properties of random Steiner triple systems just by giving an estimate asymptotically smaller than (1) for the number of Steiner triple systems not satisfying a certain property. However, for most properties it is not at all clear how to prove such estimates. Instead, we introduce a way to use Keevash’s methods to show that a uniformly random Steiner triple system can in some sense be approximated by the outcome of a random process called the *triangle removal process*. We remark that actually Keevash proved (1) with a randomised construction that involves the triangle removal process, so many properties that hold a.a.s. in the triangle removal process trivially hold a.a.s. in this random construction. Such results have been proved in [34, Proposition 3.1] and [35]. However, the Steiner triple systems obtainable by Keevash’s construction comprise a negligible proportion of the set of Steiner triple systems, and a somewhat more delicate approach is required to study a uniformly random Steiner triple system. In Section 2 we state a general theorem summarising our method.

A *matching* in a hypergraph is a collection of disjoint edges, and a *perfect matching* is a matching covering the entire vertex set. The existence of perfect matchings is one of the most central questions in the theory of graphs and hypergraphs; in particular, some of the most important recent developments in the field are the work of Johansson, Kahn and Vu [24] on perfect matchings in random hypergraphs, and the theory of Keevash and Mycroft [29] characterising when dense hypergraphs have perfect matchings. We should also mention the “nibble” method of Rödl [42] for finding almost-perfect hypergraph matchings, which has had a significant influence throughout combinatorics in the last 30 years. A perfect matching in a Steiner triple system is also called a *parallel class* or *resolution class*, and has particular significance. One of the oldest problems in combinatorics, famously solved in the affirmative by Ray-Chaudhuri and Wilson [41], asks whether for all $n \equiv 3 \pmod{6}$ there exists an order- n Steiner triple system which can be partitioned into hyperedge-disjoint perfect matchings (a *Kirkman triple system*). Alon, Kim and Spencer [2] proved that every Steiner triple system has an almost-perfect matching covering all but $o(\sqrt{n} \log^{3/2} n)$ vertices, and Bryant and Horsley [8] proved that for infinitely many $n \equiv 3 \pmod{6}$ there exist Steiner triple systems with no perfect matching. As an application of our new method, we prove our main result that if $n \equiv 3 \pmod{6}$ (that is, if $3 \mid n$ and an order- n Steiner triple system exists) then almost all order- n Steiner triple systems have many perfect matchings.

Theorem 1.1. *Let $n \equiv 3 \pmod{6}$ and let \mathcal{S} be a uniformly random order- n Steiner triple system. Then a.a.s. \mathcal{S} contains*

$$\left((1 - o(1)) \frac{n}{2e^2} \right)^{n/3}$$

perfect matchings.

We remark that if $n \equiv 1 \pmod{6}$ then obviously no order- n Steiner triple system can have a perfect matching, but exactly the same proof can be used to show that in a random order- n Steiner triple system there is a.a.s. a matching covering all but one vertex.

We prove Theorem 1.1 using our new method combined with the so-called *absorbing method*, which was introduced as a general method by Rödl, Ruciński and Szemerédi [43] (the basic idea had been used earlier, for example by Krivelevich [30]). Basically, we prove the a.a.s. existence of certain substructures that are “flexible” and allow us to complete an almost-perfect matching into a perfect one.

Up to the error term, a random Steiner triple system actually has the maximum possible number of perfect matchings: we also prove the following upper bound.

Theorem 1.2. *Any Steiner triple system S has at most*

$$\left((1 + o(1)) \frac{n}{2e^2} \right)^{n/3}$$

perfect matchings.

The proof of Theorem 1.2 is quite short, and uses the notion of *entropy*. This particular type of argument was introduced by Radhakrishnan [40] and further developed by Linal and Luria [33], among others.

1.1 Latin squares

An order- n Latin square is usually defined as an $n \times n$ array of the numbers between 1 and n (we call these *symbols*), such that each row and column contains each symbol exactly once. As mentioned earlier, this is equivalent to a 3-uniform hypergraph whose hyperedges comprise a triangle-decomposition of the edges of the complete tripartite graph $K_{n,n,n}$ (the three parts correspond to the rows, columns and symbols, so a triangle (i, j, k) corresponds to putting the symbol k in the cell (i, j)). A perfect matching in this hypergraph is called a *transversal* and the property of containing a transversal is of great interest. In particular, the famous Ryser–Brualdi–Stein conjecture speculates that every odd-order Latin square has a transversal, and every even-order Latin square has a partial transversal of size $n - 1$. Although this conjecture remains wide open, there have been many partial results and generalizations; see for example [16, 11, 21, 1, 38]. See [45] for an introduction to the subject of Latin transversals, and [25] for an introduction to Latin squares in general.

The counterpart of Theorem 1.2 for Latin squares, that a Latin square can have no more than $((1 + o(1))n/e^2)^n$ transversals, was first proved by Taranenko [44]. Glebov and Luria [19] gave a simpler entropy-based proof of the same fact and asked whether the counterpart of Theorem 1.1 holds: do almost all Latin squares have essentially the maximum possible number of transversals? Although there do exist a number of techniques for studying random Latin squares (see for example [9, 20, 36, 12, 32]), none of them seem suitable to attack this question.

In the time since the first version of this paper, Keevash [28] has generalised his methods to a number of different classes of designs, including Latin squares. Using a result from this recent work, it is straightforward to adapt the methods used to prove Theorem 1.1 to prove the following theorem, answering Glebov and Luria’s question and proving that the Ryser–Brualdi–Stein conjecture holds for almost all Latin squares.

Theorem 1.3. *Let L be a uniformly random order- n Latin square. Then a.a.s. L contains*

$$\left((1 - o(1)) \frac{n}{e^2} \right)^n$$

transversals.

1.2 Structure of the paper

The structure of this paper is as follows. In Section 2 we present our general theorem for comparing random Steiner triple systems with the triangle removal process. The proof requires a straightforward but necessary generalization of (1), estimating the number of completions of a partial Steiner triple system, and this in turn requires a routine analysis of the triangle removal process. These parts of the proof are deferred to Section 3 and Section 4. In Section 5 we use the theory from Section 2 and the absorbing method to prove Theorem 1.1. In Section 6 we prove Theorem 1.2, and in Section 7 we explain how to adapt our methods to prove Theorem 1.3. Finally, in Section 8 we have some concluding remarks, including a long list of open problems.

1.3 Notation

We use standard asymptotic notation throughout. For functions $f = f(n)$ and $g = g(n)$:

- $f = O(g)$ means there is a constant C such that $|f| \leq C|g|$,
- $f = \Omega(g)$ means there is a constant $c > 0$ such that $f \geq c|g|$,
- $f = \Theta(g)$ means that $f = O(g)$ and $f = \Omega(g)$,
- $f = o(g)$ means that $f/g \rightarrow 0$ as $n \rightarrow \infty$. To say that a.a.s. $f = o(g)$ means that for any $\varepsilon > 0$, a.a.s. $f/g < \varepsilon$.

Also, following [27], the notation $f = 1 \pm \varepsilon$ means $1 - \varepsilon \leq f \leq 1 + \varepsilon$.

We also use standard graph theory notation: $V(G)$ and $E(G)$ are the sets of vertices and (hyper)edges of a (hyper)graph G , and $v(G)$ and $e(G)$ are the cardinalities of these sets. The subgraph of G induced by a vertex subset U is denoted $G[U]$, the degree of a vertex v is denoted $\deg_G(v)$, and the subgraph obtained by deleting v is denoted $G - v$.

For a positive integer n , we write $[n]$ for the set $\{1, 2, \dots, n\}$. For a real number x , the floor and ceiling functions are denoted $\lfloor x \rfloor = \max\{i \in \mathbb{Z} : i \leq x\}$ and $\lceil x \rceil = \min\{i \in \mathbb{Z} : i \geq x\}$. We will however mostly omit floor and ceiling signs and assume large numbers are integers, wherever divisibility considerations are not important. All logarithms are in base e .

Finally, we remark that throughout the paper we adopt the convention that random variables (and random objects more generally) are printed in bold.

1.4 Acknowledgements

The author would like to thank Asaf Ferber and Benny Sudakov for very helpful discussions about random designs and the intricacies of the absorbing method, Patrick Morris and Stefan Glock for pointing out some oversights in earlier versions of this paper, and Lutz Warnke for pointing out some related references.

2 Random Steiner triple systems via the triangle removal process

In this section we describe our method for comparing random Steiner triple systems with the outcome of the triangle removal process. Define

$$N = \binom{n}{2}/3 = (1 + o(1))n^2/6$$

to be the number of hyperedges in a Steiner triple system. We assume throughout this section that n is 1 or 3 mod 6.

Definition 2.1 (partial systems). A *partial Steiner triple system* (or *partial system* for short) is a 3-uniform hypergraph on $[n]$ in which every pair of vertices is included in no more than one hyperedge. Let \mathcal{S}_m be the set of partial systems with m hyperedges. We will also want to consider partial systems equipped with an ordering on their hyperedges. Let \mathcal{O} be the set of ordered Steiner triple systems, and let \mathcal{O}_m be the set of ordered partial systems with m hyperedges. For $S \in \mathcal{O}_m$ and $i \leq m$, let S_i be the ordered partial system consisting of just the first i hyperedges of S . For a (possibly ordered) partial system S , let $G(S)$ be the graph with an edge for every pair of vertices which does not appear in any hyperedge of S . So, if S has m hyperedges, then $G(S)$ has $\binom{n}{2} - 3m$ edges.

Definition 2.2 (quasirandomness). For a graph G with n vertices and m edges, let $d(G) = m/\binom{n}{2}$ denote its density. We say G is (ε, h) -*quasirandom* if for every set A of at most h vertices, we

have $|\bigcap_{w \in A} N_G(w)| = (1 \pm \varepsilon)d(G)^{|A|}n$. Let $\mathcal{S}_m^{\varepsilon, h} \subseteq \mathcal{S}_m$ be the set of partial systems $S \in \mathcal{S}_m$ such that $G(S)$ is (ε, h) -quasirandom, and let $\mathcal{O}_m^{\varepsilon, h} \subseteq \mathcal{O}_m$ be the set of ordered partial systems $S \in \mathcal{O}_m$ such that $S_i \in \mathcal{S}_i^{\varepsilon, h}$ for each $i \leq m$.

Definition 2.3 (the triangle removal process). The triangle removal process is defined as follows. Start with the complete graph K_n and iteratively delete a triangle chosen uniformly at random from all triangles in the remaining graph. If we continue this process for m steps, the deleted triangles (in order) can be interpreted as an ordered partial system in \mathcal{O}_m . It is also possible that the process aborts (because there are no triangles left) before m steps, in which case we say it returns the value “*”. We denote by $\mathbb{R}(n, m)$ the resulting distribution on $\mathcal{O}_m \cup \{*\}$.

Now, we can state a general theorem comparing random Steiner triple systems with the triangle removal process. Basically, if we can show that the first few edges of the triangle removal process (as an ordered partial system) satisfy some property with extremely high probability, then it follows that the first few edges of a uniformly random ordered Steiner triple system satisfy the same property with high probability. Moreover, it suffices to study the triangle removal process conditioned on some “good” event, provided that this event contains the event that our partial system is sufficiently quasirandom.

Theorem 2.4. *Fixing $h \in \mathbb{N}$ and sufficiently small $a > 0$, there is $b = b(a, h) > 0$ such that the following holds. Fix $\alpha \in (0, 1)$, let $\mathcal{P} \subseteq \mathcal{O}_{\alpha N}$ be a property of ordered partial systems, let $\varepsilon = n^{-a}$, let $\mathcal{Q} \supseteq \mathcal{O}_{\alpha N}^{\varepsilon, h}$, let $\mathbf{S} \in \mathcal{O}$ be a uniformly random ordered Steiner triple system and let $\mathbf{S}' \sim \mathbb{R}(n, \alpha N)$. If*

$$\Pr(\mathbf{S}' \notin \mathcal{P} \mid \mathbf{S}' \in \mathcal{Q}) \leq \exp(-n^{2-b})$$

then

$$\Pr(\mathbf{S}_{\alpha N} \notin \mathcal{P}) \leq \exp(-\Omega(n^{1-2a})).$$

Note that (as we prove in Section 4), the triangle removal process is likely to produce quasirandom graphs; that is, $\Pr(\mathbf{S}' \in \mathcal{Q}) \geq \Pr(\mathbf{S}' \in \mathcal{O}_{\alpha N}^{\varepsilon, h}) = 1 - o(1)$. However, as we will see in Section 5.1.3, the conditioning in Theorem 2.4 can still be useful because the probabilities under consideration are so small (it is certainly not true that $\Pr(\mathbf{S}' \notin \mathcal{Q})$ is anywhere near as small as $\exp(-\Omega(n^2))$).

The proof of Theorem 2.4 follows from a sequence of several lemmas. The most important is the following: we can estimate the number of ways to complete a partial system S , and show that it does not vary too much between choices of S .

Lemma 2.5. *For an ordered partial system $S \in \mathcal{O}_m$, let $\mathcal{O}^*(S) \subseteq \mathcal{O}$ be the set of ordered Steiner triple systems S^* such that $S_m^* = S$. Fixing sufficiently large $h \in \mathbb{N}$ and any $a > 0$, there is $b = b(a, h) > 0$ such that the following holds. For any fixed $\alpha \in (0, 1)$, any $\varepsilon = \varepsilon(n) \leq n^{-a}$ and any $S, S' \in \mathcal{O}_{\alpha N}^{\varepsilon, h}$,*

$$\frac{|\mathcal{O}^*(S)|}{|\mathcal{O}^*(S')|} \leq \exp(O(n^{2-b})).$$

Lemma 2.5 can be proved with slight adaptations to proofs of Keevash [27] and Linial and Luria [33] giving lower and upper bounds on the total number of Steiner triple systems. The details are in Section 3.

The point of Lemma 2.5 is that if we can prove some property holds with extremely high probability (say $1 - \exp(-\Omega(n^2))$) in a uniformly random $\mathbf{S} \in \mathcal{O}_{\alpha N}^{\varepsilon, h}$, then it also holds with essentially the same probability in $\mathbf{S}_{\alpha N}$, for a uniformly random $\mathbf{S} \in \mathcal{O}$ conditioned on the event $\mathbf{S}_{\alpha N} \in \mathcal{O}_{\alpha N}^{\varepsilon, h}$. The next step is to show that the event $\mathbf{S}_{\alpha N} \in \mathcal{O}_{\alpha N}^{\varepsilon, h}$ is very likely. In fact, this event occurs a.a.s. for a random ordering of any given Steiner triple system. We prove the following lemma in Section 2.1.

Lemma 2.6. *The following holds for any fixed $h \in \mathbb{N}$, $\alpha \in (0, 1)$ and $a \in (0, 1/2)$. Let $\varepsilon = n^{-a}$, consider any Steiner triple system S , and uniformly at random order its hyperedges to obtain an ordered Steiner triple system $\mathbf{S} \in \mathcal{O}$. Then $\Pr(\mathbf{S}_{\alpha N} \notin \mathcal{O}_{\alpha N}^{\varepsilon, h}) = \exp(-\Omega(n^{1-2a}))$.*

The upshot of Lemmas 2.5 and 2.6 is that if we can prove a property holds with extremely high probability in a uniformly random $\mathbf{S} \in \mathcal{O}_{\alpha N}^{\varepsilon, h}$ for sufficiently small ε and sufficiently large h , then that property also holds a.a.s. in the first αN hyperedges of a uniformly random $\mathbf{S} \in \mathcal{O}$.

Next, the following lemma says that each $S \in \mathcal{O}_{\alpha N}^{\varepsilon, h}$ is roughly equally likely to be produced by the triangle removal process, so that $\mathbb{R}(n, \alpha N)$ approximates the uniform distribution on $\mathcal{O}_{\alpha N}^{\varepsilon, h}$. It is proved in Section 2.2.

Lemma 2.7. *The following holds for any fixed $a \in (0, 2)$ and $\alpha \in [0, 1]$. Let $\varepsilon = n^{-a}$, let $S, S' \in \mathcal{O}_{\alpha N}^{\varepsilon, 2}$ and let $\mathbf{S} \sim \mathbb{R}(n, \alpha N)$. Then*

$$\frac{\Pr(\mathbf{S} = S)}{\Pr(\mathbf{S} = S')} \leq \exp(O(n^{2-a})).$$

We can finally combine everything to prove Theorem 2.4.

Proof of Theorem 2.4. First, note that we can assume h is large enough for Lemma 2.5, because if we increase h we still have $\mathcal{Q} \supseteq \mathcal{O}_{\alpha N}^{\varepsilon, h}$. Let $\mathbf{S}'' \in \mathcal{O}_{\alpha N}^{\varepsilon, h}$ be a uniformly random partial system in $\mathcal{O}_{\alpha N}^{\varepsilon, h}$. By Lemma 2.7, we have

$$\Pr(\mathbf{S}' \notin \mathcal{P} \mid \mathbf{S}' \in \mathcal{O}_{\alpha N}^{\varepsilon, h}) = \exp(O(n^{2-a})) \Pr(\mathbf{S}'' \notin \mathcal{P}).$$

(Recalling the definition of big-oh notation, we emphasise that this formula encapsulates a lower bound as well as an upper bound). Next, let $c = b(a, h)$ in the notation of Lemma 2.5. We similarly have

$$\Pr(\mathbf{S}_{\alpha N} \notin \mathcal{P} \mid \mathbf{S}_{\alpha N} \in \mathcal{O}_{\alpha N}^{\varepsilon, h}) = \exp(O(n^{2-c})) \Pr(\mathbf{S}'' \notin \mathcal{P}).$$

Using Lemma 2.6, it follows that

$$\begin{aligned} \Pr(\mathbf{S}_{\alpha N} \notin \mathcal{P}) &\leq \Pr(\mathbf{S}_{\alpha N} \notin \mathcal{P} \mid \mathbf{S}_{\alpha N} \in \mathcal{O}_{\alpha N}^{\varepsilon, h}) + \Pr(\mathbf{S}_{\alpha N} \notin \mathcal{O}_{\alpha N}^{\varepsilon, h}) \\ &\leq \exp(O(n^{2-c})) \exp(O(n^{2-a})) \Pr(\mathbf{S}' \notin \mathcal{P} \mid \mathbf{S}' \in \mathcal{O}_{\alpha N}^{\varepsilon, h}) + \exp(-\Omega(n^{1-2a})). \end{aligned}$$

But, if a is small enough then $\Pr(\mathbf{S}' \in \mathcal{Q}) \geq \Pr(\mathbf{S}' \in \mathcal{O}_{\alpha N}^{\varepsilon, h}) = 1 - o(1)$, by a standard analysis of the triangle removal process; see Theorem 4.1. So,

$$\begin{aligned} \Pr(\mathbf{S}' \notin \mathcal{P} \mid \mathbf{S}' \in \mathcal{O}_{\alpha N}^{\varepsilon, h}) &= \frac{\Pr(\mathbf{S}' \notin \mathcal{P} \text{ and } \mathbf{S}' \in \mathcal{O}_{\alpha N}^{\varepsilon, h})}{\Pr(\mathbf{S}' \in \mathcal{O}_{\alpha N}^{\varepsilon, h})} \\ &\leq \frac{\Pr(\mathbf{S}' \notin \mathcal{P} \text{ and } \mathbf{S}' \in \mathcal{Q})}{\Pr(\mathbf{S}' \in \mathcal{O}_{\alpha N}^{\varepsilon, h})} \\ &= (1 + o(1)) \frac{\Pr(\mathbf{S}' \notin \mathcal{P} \text{ and } \mathbf{S}' \in \mathcal{Q})}{\Pr(\mathbf{S}' \in \mathcal{Q})} \\ &= (1 + o(1)) \Pr(\mathbf{S}' \notin \mathcal{P} \mid \mathbf{S}' \in \mathcal{Q}). \end{aligned}$$

Choosing b such that $b < \min\{c, a\}$ and $2 - b \geq 1 - 2a$, we then have

$$\Pr(\mathbf{S}_{\alpha N} \notin \mathcal{P}) \leq \exp(-\Omega(n^{1-2a}))$$

as desired. \square

In Section 2.1 we prove Lemma 2.6 and in Section 2.2 we prove Lemma 2.7. Also, in Section 2.3 we prove some lemmas which are useful tools for applying Theorem 2.4 in practice.

2.1 Randomly ordered Steiner triple systems

In this subsection we prove Lemma 2.6.

Proof. Consider $m \leq \alpha N$. Note that \mathcal{S}_m (as an unordered partial system) is a uniformly random subset of m hyperedges of S . Also note that

$$d(G(\mathcal{S}_m)) = \frac{\binom{n}{2} - 3m}{\binom{n}{2}} = 1 - \frac{m}{N}.$$

We can obtain a random partial system almost equivalent to \mathcal{S}_m by including each hyperedge of S with independent probability m/N . Let \mathcal{S}' denote the partial system so obtained, and let $\mathcal{G}' = G(\mathcal{S}')$. Now, fix a set A of at most h vertices. It suffices to prove

$$\left| \bigcap_{w \in A} N_{\mathcal{G}'}(w) \right| = (1 \pm n^{-a}) \left(1 - \frac{m}{N}\right)^{|A|} n, \quad (2)$$

with probability $1 - \exp(-\Omega(n^{1-2a}))$. Indeed, the so-called Pittel inequality (see [22, p. 17]) would imply that the same estimate holds with essentially the same probability if we replace \mathcal{S}' with \mathcal{S}_m (thereby replacing \mathcal{G}' with $G(\mathcal{S}_m)$). We would then be able to finish the proof by applying the union bound over all $m \leq \alpha N$ and all choices of A .

Note that there are at most $\binom{|A|}{2} = O(1)$ hyperedges of S that include more than one vertex in A (by the defining property of a Steiner triple system). Let U be the set of vertices involved in these atypical hyperedges, plus the vertices in A , so that $|U| = O(1)$. Let $\mathbf{N} = \left| \left(\bigcap_{w \in A} N_{\mathcal{G}'}(w) \right) \setminus U \right|$. For every $v \notin U$ and $w \in A$ there is exactly one hyperedge e_v^w in S containing v and w , whose presence in \mathcal{S}' would prevent v from contributing to \mathbf{N} . For each fixed $v \notin U$ the hyperedges e_v^w , for $w \in A$, are distinct, so

$$\Pr \left(v \in \bigcap_{w \in A} N_{\mathcal{G}'}(w) \right) = \left(1 - \frac{m}{N}\right)^{|A|},$$

and by linearity of expectation $\mathbb{E}\mathbf{N} = (1 - m/N)^{|A|}(n - O(1))$. Now, \mathbf{N} is determined by the presence of at most $(n - |U|)|A| = O(n)$ hyperedges in \mathcal{S}' , and changing the presence of each affects \mathbf{N} by at most $2 = O(1)$. So, by the Azuma–Hoeffding inequality (see [22, Section 2.4]),

$$\begin{aligned} \Pr \left(\left| \mathbf{N} - \left(1 - \frac{m}{N}\right)^{|A|} n \right| > n^{-a} \left(1 - \frac{m}{N}\right)^{|A|} n - |U| \right) &\leq \exp \left(-\Omega \left(\frac{\left(n^{-a} (1 - \alpha)^h n \right)^2}{n} \right) \right) \\ &= \exp(-\Omega(n^{1-2a})). \end{aligned}$$

Finally, we recall that $\left| \left(\bigcap_{w \in A} N_{\mathcal{G}'}(w) \right) \right| = \mathbf{N} \pm |U|$, which completes the proof of (2). \square

2.2 Approximate uniformity of the triangle removal process

In this subsection we prove Lemma 2.7. We first make the simple observation that subgraph statistics in a quasirandom graph G can be easily estimated in terms of the density of G . In fact we prove a more general extension lemma about “rooted” subgraphs, which we will also use later in the paper. An *embedding* of a graph H in a graph G is a surjective homomorphism from H into G .

Proposition 2.8. *Let H be a fixed graph with identified vertices u_1, \dots, u_k . Let G be an $(\varepsilon, v(H) - 1)$ -quasirandom graph on n vertices, and let $\phi : \{u_1, \dots, u_k\} \rightarrow G$ be an embedding of $F = H[\{u_1, \dots, u_k\}]$ in G . Then, the number of ways to extend ϕ to an embedding of H in G is*

$$(1 \pm O(\varepsilon)) d(G)^{e(H) - e(F)} n^{v(H) - k}.$$

In particular, taking $k = 0$, the number of copies of H in G is

$$(1 \pm O(\varepsilon))d(G)^{e(H)} \frac{n^{v(H)}}{|\text{Aut}(H)|},$$

where $|\text{Aut}(H)|$ is the number of automorphisms of H .

Proof. Let $U = \{u_1, \dots, u_k\}$; we proceed by induction on the number of vertices in $V(H) \setminus U$. The base case is where $U = V(H)$, which is trivial. Suppose there is a vertex $v \in V(H) \setminus U$; by induction there are

$$(1 \pm O(\varepsilon))d(G)^{e(H)-e(F)-\deg_H(v)} n^{v(H)-|U|-1}$$

embeddings of $H - v$ extending ϕ . For each such embedding, by $(\varepsilon, v(H) - 1)$ -quasirandomness, there are $(1 \pm \varepsilon)d(G)^{\deg_H(v)} n$ ways to choose a vertex of G with the right adjacencies to complete the embedding of H . The desired result follows. \square

Now we are ready to prove Lemma 2.7.

Proof. Each $G(S_i)$ has

$$(1 \pm O(n^{-a})) \left(1 - \frac{i}{N}\right)^3 \frac{n^3}{6}$$

triangles, by $(n^{-a}, 2)$ -quasirandomness and Proposition 2.8. We therefore have

$$\Pr(\mathbf{S} = S) = \prod_{i=0}^{\alpha N - 1} \frac{1}{(1 \pm O(n^{-a}))(1 - i/N)^3 n^3 / 6},$$

and a similar expression holds for $\Pr(\mathbf{S} = S')$. Taking quotients term-by-term gives

$$\begin{aligned} \frac{\Pr(\mathbf{S} = S)}{\Pr(\mathbf{S} = S')} &\leq (1 + O(n^{-a}))^{\alpha N} \\ &\leq \exp(O(n^{2-a})) \end{aligned}$$

as desired. \square

2.3 A coupling lemma and a concentration inequality

In this subsection we prove two lemmas that will be useful in combination with Theorem 2.4. First, after some definitions we will show how to couple the triangle removal process with a simpler random hypergraph distribution.

Definition 2.9. For a partial system S , let $\mathbb{G}(S, p)$ be the random distribution on 3-uniform hypergraphs where each hyperedge not conflicting with S (that is, not intersecting a hyperedge of S in more than 2 vertices) is included with probability p . So, if \emptyset is the empty order- n partial system, then $\mathbb{G}(\emptyset, p) =: \mathbb{G}(n, p)$ is the standard binomial random 3-uniform hypergraph. Let $\mathbb{G}^*(S, p)$ be the distribution on partial systems obtained from $\mathbb{G}(S, p)$ by considering all hyperedges which intersect another hyperedge in more than 2 vertices, and deleting all these hyperedges (and let $\mathbb{G}^*(n, p) = \mathbb{G}^*(\emptyset, p)$). Let $\mathbb{R}(S, m)$ be the partial system distribution obtained with m steps of the triangle removal process starting from $G(S)$.

Note that $\binom{n}{3}/n = (1 + o(1))N$. For small $\alpha > 0$, we can view $\mathbb{G}^*(S, \alpha/n)$ as a ‘‘bite’’ of a ‘‘nibbling’’ process, that should be comparable to $\mathbb{R}(S, \alpha N)$.

Lemma 2.10. *Let \mathcal{P} be a property of unordered partial systems that is monotone increasing in the sense that $S \in \mathcal{P}$ and $S' \supseteq S$ implies $S' \in \mathcal{P}$. Fix $\alpha \in (0, 1)$ and $S \in \mathcal{O}_m$ for some $m \leq N - \alpha N$. Let $\mathbf{S} \sim \mathbb{R}(S, \alpha N)$ and $\mathbf{S}^* \sim \mathbb{G}^*(S, \alpha/n)$. Then*

$$\Pr(S \cup \mathbf{S} \notin \mathcal{P}) = O(1) \Pr(S \cup \mathbf{S}^* \notin \mathcal{P}),$$

where $S \cup S'$ denotes the unordered partial system containing all the edges of S and of S' .

Before proving Lemma 2.10, we remark that we abuse notation slightly and use the conventions that $*$ is a superset of every partial system, that $* \in \mathcal{P}$, and that $S \cup * = *$. That is, if the triangle removal process aborts before reaching αN edges, we still say that it satisfies \mathcal{P} .

Proof. Let $\mathbf{S}^* \sim \mathbb{G}^*(S, \alpha/n)$ be obtained from $\mathbf{G} \sim \mathbb{G}(S, \alpha/n)$. If $e(\mathbf{G}) \leq \alpha N$, then \mathbf{S}^* can be coupled as a subset of \mathbf{S} . Indeed, a random ordering of the edges of \mathbf{G} can be viewed as the first few elements of a random ordering of the set of triangles of $G(S)$, and the triangle removal process with this ordering produces a superset of \mathbf{S}^* . It follows that

$$\Pr(S \cup \mathbf{S} \notin \mathcal{P}) \leq \Pr(S \cup \mathbf{S}^* \notin \mathcal{P} \mid e(\mathbf{G}) \leq \alpha N).$$

Next, note that $e(\mathbf{G})$ has a binomial distribution with mean $\mathbb{E}e(\mathbf{G}) = \binom{n}{3}\alpha/n \leq \alpha N$, so it is easy to see that $\Pr(e(\mathbf{G}) \leq \alpha N) = \Omega(1)$. It follows that

$$\Pr(S \cup \mathbf{S} \notin \mathcal{P}) \leq \Pr(S \cup \mathbf{S}^* \notin \mathcal{P}) / \Pr(e(\mathbf{G}) \leq \alpha N) = O(1) \Pr(S \cup \mathbf{S}^* \notin \mathcal{P}). \quad \square$$

We remark that Lemma 2.10 is similar in spirit to [46, Lemma 6.1]. In this subsection we also state and prove a bounded-differences inequality with Bernstein-type tails which can be used to analyse $\mathbb{G}^*(S, \alpha/n)$. Standard bounded-difference inequalities such as the Azuma–Hoeffding inequality do not provide strong enough tail bounds to apply Theorem 2.4.

Theorem 2.11. *Let $\boldsymbol{\omega} = (\omega_1, \dots, \omega_n)$ be a sequence of independent, identically distributed random variables with $\Pr(\omega_i = 1) = p$ and $\Pr(\omega_i = 0) = 1 - p$. Let $f : \{0, 1\}^n \rightarrow \mathbb{R}$ satisfy the Lipschitz condition $|f(\boldsymbol{\omega}) - f(\boldsymbol{\omega}')| \leq K$ for all pairs $\boldsymbol{\omega}, \boldsymbol{\omega}' \in \{0, 1\}^n$ differing in exactly one coordinate. Then*

$$\Pr(|f(\boldsymbol{\omega}) - \mathbb{E}f(\boldsymbol{\omega})| > t) \leq \exp\left(-\frac{t^2}{4K^2np + 2Kt}\right).$$

Proof. We use Freedman’s inequality (Lemma 4.2), with the Doob martingale $\mathbf{X}(0), \dots, \mathbf{X}(n)$ defined by $\mathbf{X}(i) = \mathbb{E}[f(\boldsymbol{\omega}) \mid \omega_1, \dots, \omega_i]$. With $\Delta\mathbf{X}(i)$ as the one-step change $\mathbf{X}(i+1) - \mathbf{X}(i)$ and with $V(i) = \sum_{j=0}^i \mathbb{E}[(\Delta\mathbf{X}(j))^2 \mid \omega_1, \dots, \omega_j]$, it suffices to show that $V(n) \leq 2K^2np$ with probability 1.

Condition on $\omega_1, \dots, \omega_i$ (thereby conditioning on $\mathbf{X}(i)$). Let X^0 and X^1 be the values of $\mathbf{X}(i+1)$ in the cases $\omega_{i+1} = 0$ and $\omega_{i+1} = 1$, respectively. We have

$$\begin{aligned} \mathbf{X}(i) &= pX^1 + (1-p)X^0, \\ |\mathbf{X}(i) - X^0| &= p|X^1 - X^0| \leq Kp. \end{aligned}$$

So,

$$\begin{aligned} \mathbb{E}[(\Delta\mathbf{X}(i))^2 \mid \omega_1, \dots, \omega_i] &= p(\mathbf{X}(i) - X^1)^2 + (1-p)(\mathbf{X}(i) - X^0)^2 \\ &\leq K^2p + (1-p)K^2p^2 \\ &\leq 2K^2p. \end{aligned}$$

The desired bound on $V(n)$ follows. □

Since the first version of this paper, we learned that Theorem 2.11 is also a direct consequence of [47, Theorem 1.3].

3 Counting completions of Steiner triple systems

In this section we prove Lemma 2.5. This is accomplished with minor adaptations of proofs by Linial and Luria [33] and Keevash [27]. As in Section 2, let $N = \binom{n}{2}/3$ and assume that n is 1 or 3 mod 6.

For a partial system $S \in \mathcal{S}_{\alpha N}^{n-a}$, let $\mathcal{S}^*(S)$ be the number of Steiner triple systems that include S . We want to determine $|\mathcal{O}^*(S)| = (N - \alpha N)! |\mathcal{S}^*(S)|$ up to a factor of $e^{n^{2-b}}$ (for some $b > 0$).

First, we can get an upper bound via the *entropy method*, as used by Linial and Luria [33]. The reader may wish to refer to that paper for more detailed exposition.

Before we begin the proof, we briefly remind the reader of the basics of the notion of entropy. For random elements \mathbf{X}, \mathbf{Y} with supports $\text{supp } \mathbf{X}, \text{supp } \mathbf{Y}$, we define the (base- e) *entropy*

$$H(\mathbf{X}) = - \sum_{x \in \text{supp } \mathbf{X}} \Pr(\mathbf{X} = x) \log(\Pr(\mathbf{X} = x))$$

and the *conditional entropy*

$$H(\mathbf{X} | \mathbf{Y}) = \sum_{y \in \text{supp } \mathbf{Y}} \Pr(\mathbf{Y} = y) H(\mathbf{X} | \mathbf{Y} = y).$$

We will use two basic properties of entropy. First, we always have $H(\mathbf{X}) \leq \log |\text{supp } \mathbf{X}|$, with equality only when \mathbf{X} has the uniform distribution on its support. Second, for any sequence of random elements $\mathbf{X}_1, \dots, \mathbf{X}_n$, we have

$$H(\mathbf{X}_1, \dots, \mathbf{X}_n) = \sum_{i=1}^n H(\mathbf{X}_i | \mathbf{X}_1, \dots, \mathbf{X}_{i-1}).$$

See for example [14] for an introduction to the notion of entropy and proofs of the above two facts.

Theorem 3.1. *For any $a > 0$, any $\alpha \in [0, 1]$, and any $S \in \mathcal{S}_{\alpha N}^{n-a,2}$,*

$$|\mathcal{S}^*(S)| \leq \left(\left(1 + O(n^{-a} + n^{-1/2}) \right) \left(\frac{1-\alpha}{e} \right)^2 n \right)^{N(1-\alpha)}.$$

Proof. Let $\mathbf{S}^* \in \mathcal{S}^*(S)$ be a uniformly random completion of S . We will estimate the entropy $H(\mathbf{S}^*) = \log |\mathcal{S}^*(S)|$ of \mathbf{S}^* .

Let $G = G(S)$. For each $e = \{x, y\} \in G$, let $\{x, y, z_e\}$ be the hyperedge that includes e in \mathbf{S}^* . So, the sequence $(z_e)_{e \in G}$ determines \mathbf{S}^* . For any ordering on the edges of G , we have

$$H(\mathbf{S}^*) = \sum_{e \in G} H(z_e | (z_{e'} : e' < e)). \quad (3)$$

Now, a sequence $\lambda \in [0, 1]^{E(G)}$ with all λ_e distinct induces an ordering on the edges of G , with $e' < e$ when $\lambda_{e'} > \lambda_e$. Let $\mathbf{R}_e(\lambda)$ be an upper bound on $|\text{supp}(z_e | z_{e'} : \lambda_{e'} > \lambda_e)|$ defined as follows. $\mathbf{R}_e(\lambda) = 1$ if $\lambda_{\{x, z_e\}} > \lambda_e$ or $\lambda_{\{y, z_e\}} > \lambda_e$ (because in this case z_e is determined). Otherwise, $\mathbf{R}_e(\lambda)$ is 1 plus the number of vertices $v \notin \{x, y, z_e\}$ such that $\{x, v\}, \{y, v\} \in G$, and $\lambda_{e'} < \lambda_e$ for each of the 6 edges $e' \in G$ included in the hyperedges that include $\{x, v\}$ and $\{y, v\}$ in \mathbf{S}^* . Since $\mathbf{R}_e(\lambda)$ is an upper bound on $|\text{supp}(z_e | z_{e'} : \lambda_{e'} > \lambda_e)|$, we have

$$H(z_e | z_{e'} : \lambda_{e'} > \lambda_e) \leq \mathbb{E}[\log \mathbf{R}_e(\lambda)]. \quad (4)$$

It follows from (3) and (4) that

$$H(\mathbf{S}^*) \leq \sum_{e \in G} \mathbb{E}[\log \mathbf{R}_e(\lambda)].$$

This is true for any fixed λ , so it is also true if λ is chosen randomly, as follows. Let $\boldsymbol{\lambda} = (\lambda_e)_{e \in G}$ be a sequence of independent random variables, where each λ_e has the uniform distribution in $[0, 1]$. (With probability 1 each λ_v is distinct). Then

$$H(\mathcal{S}^*) \leq \sum_{e \in G} \mathbb{E}[\log \mathbf{R}_e(\boldsymbol{\lambda})].$$

Next, for any $S^* \in \mathcal{S}^*(S)$ and $\lambda_e \in [0, 1]$, let

$$R_e^{S^*, \lambda_e} = \mathbb{E}[\mathbf{R}_e(\boldsymbol{\lambda}) \mid \mathcal{S}^* = S^*, \lambda_e = \lambda_e, \boldsymbol{\lambda}_{\{x, z_e\}}, \boldsymbol{\lambda}_{\{y, z_e\}} < \lambda_e].$$

(Note that $\lambda_e = \lambda_e$ occurs with probability zero, so formally we should condition on $\lambda_e = \lambda_e \pm d\lambda_e$ and take limits in what follows, but there are no continuity issues so we will ignore this detail). Now, in G , by $(n^{-a}, 2)$ -quasirandomness x and y have $(1 + O(n^{-a}))(1 - \alpha)^2 n$ common neighbours other than z_e . By the definition of $\mathbf{R}_e(\boldsymbol{\lambda})$ and linearity of expectation, we have

$$R_e^{S^*, \lambda_e} = 1 + (1 + O(n^{-a}))(1 - \alpha)^2 \lambda_e^6 n.$$

By Jensen's inequality,

$$\mathbb{E}[\log \mathbf{R}_e(\boldsymbol{\lambda}) \mid \mathcal{S}^* = S^*, \lambda_e = \lambda_e, \boldsymbol{\lambda}_{\{x, z_e\}}, \boldsymbol{\lambda}_{\{y, z_e\}} < \lambda_e] \leq \log R_e^{S^*, \lambda_e},$$

and

$$\Pr(\boldsymbol{\lambda}_{\{x, z_e\}}, \boldsymbol{\lambda}_{\{y, z_e\}} < \lambda_e \mid \lambda_e = \lambda_e) = \lambda_e^2,$$

so

$$\mathbb{E}[\log \mathbf{R}_e(\boldsymbol{\lambda}) \mid \mathcal{S}^* = S^*, \lambda_e = \lambda_e] \leq \lambda_e^2 \log R_e^{S^*, \lambda_e} + (1 - \lambda_e^2) \log 1 = \lambda_e^2 \log R_e^{S^*, \lambda_e}.$$

We then have

$$\begin{aligned} \mathbb{E}[\log \mathbf{R}_e(\boldsymbol{\lambda}) \mid \mathcal{S}^* = S^*] &\leq \mathbb{E}[\lambda_e^2 \log R_e^{S^*, \lambda_e}] \\ &= \int_0^1 \lambda_e^2 \log(1 + (1 + O(n^{-a}))(1 - \alpha)^2 \lambda_e^6 n) d\lambda_e. \end{aligned}$$

For $C > 0$ we can compute

$$\int_0^1 t^2 \log(1 + Ct^6) dt = \frac{1}{3} \left(\log(1 + C) - 2 + \frac{2 \arctan \sqrt{C}}{\sqrt{C}} \right), \quad (5)$$

so (taking $C = (1 + O(n^{-a}))(1 - \alpha)^2 n$) we deduce

$$\mathbb{E}[\log \mathbf{R}_e(\boldsymbol{\lambda}) \mid \mathcal{S}^* = S^*] \leq \frac{1}{3} \left(\log((1 - \alpha)^2 n) - 2 \right) + O(n^{-a} + n^{-1/2}).$$

We conclude that

$$\begin{aligned} \log |\mathcal{S}^*(S)| &\leq H(\mathcal{S}^*) \\ &\leq \sum_{e \in G} \mathbb{E}[\log \mathbf{R}_e(\boldsymbol{\lambda})] \\ &\leq (N - \alpha N) \left(\log((1 - \alpha)^2 n) - 2 + O(n^{-a} + n^{-1/2}) \right), \end{aligned}$$

which is equivalent to the theorem statement. \square

For the lower bound, we will count ordered Steiner triple systems.

Theorem 3.2. Fixing sufficiently large $h \in \mathbb{N}$ and any $a > 0$, there is $b = b(a, h) > 0$ such that the following holds. For any $\alpha \in (0, 1)$ and any $S \in \mathcal{O}_{\alpha N}^{n^{-a}, h}$,

$$|\mathcal{O}^*(S)| \geq \left(\left(1 - O(n^{-b})\right) \left(\frac{1-\alpha}{e}\right)^2 n \right)^{N(1-\alpha)} (N - \alpha N)!.$$

To prove Theorem 3.2 we will need an analysis of the triangle removal process (which we provide in Section 4) and the following immediate consequence of [27, Theorem 2.1].

Theorem 3.3. There are $h \in \mathbb{N}$, $\varepsilon_0, a \in (0, 1)$ and $n_0, \ell \in \mathbb{N}$ such that if $S \in \mathcal{S}_m^{\varepsilon, h}$ is a partial system with $n \geq n_0$, $d(G(S)) = 1 - m/N \geq n^{-a}$ and $\varepsilon \leq \varepsilon_0 d(G)^\ell$, then S can be completed to a Steiner triple system.

Proof of Theorem 3.2. Let $h \geq 2$, ℓ, ε_0 be as in Theorem 3.3. Let $c > 0$ be smaller than $a \cdot b(a, h)$ in the notation of Theorem 4.1, and smaller than the “ a ” in Theorem 3.3. Let $\varepsilon = n^{-c/\ell}/\varepsilon_0$ and $M = (1 - \varepsilon)N$. Let $\mathbf{S}^* \in \mathcal{O}_M \cup \{*\}$ be the result of running the triangle removal process on $G(S)$ to build a partial system extending S , until there are M hyperedges (that is, $\mathbf{S}^* = S \cup \mathbf{S}$ where $\mathbf{S} \sim \mathbb{R}(S, M - \alpha N)$ in the notation of Section 2.3). Let \mathcal{O}^* be the set of M -hyperedge (ε, h) -quasirandom ordered partial systems $S^* \in \mathcal{O}_M^{\varepsilon, h}$ extending S . The choice of c ensures that by Theorem 4.1 we a.a.s. have $\mathbf{S}^* \in \mathcal{O}^*$, and by Theorem 3.3 each $S^* \in \mathcal{O}^*$ can be completed to an ordered Steiner triple system.

Now, by Proposition 2.8 and quasirandomness, for each $S^* \in \mathcal{O}^*$, the number of triangles in each $G(S_i^*)$ is

$$(1 \pm O(n^{-c}))(1 - i/N)^3 n^3 / 6,$$

so

$$\Pr(\mathbf{S}^* = S^*) \leq \prod_{i=\alpha N}^{M-1} \frac{1}{(1 - O(n^{-c}))(1 - i/N)^3 n^3 / 6}.$$

As discussed, using Theorem 4.1 we have

$$\sum_{S^* \in \mathcal{O}^*} \Pr(\mathbf{S}^* = S^*) = 1 - o(1),$$

so

$$\begin{aligned} |\mathcal{O}^*| &\geq (1 - o(1)) \prod_{i=\alpha N}^{M-1} (1 - O(n^{-c})) \left(1 - \frac{i}{N}\right)^3 \frac{n^3}{6} \\ &= \left((1 - O(n^{-c})) \frac{n^3}{6} \right)^{(1-\alpha)N} \exp\left(3 \sum_{i=\alpha N}^{M-1} \log\left(1 - \frac{i}{N}\right) \right). \end{aligned}$$

Now, note that

$$\sum_{i=\alpha N}^{M-1} \frac{1}{N} \log\left(1 - \frac{i+1}{N}\right) \leq \int_{\alpha}^{(1-\varepsilon)} \log(1-t) dt \leq \sum_{i=\alpha N}^{M-1} \frac{1}{N} \log\left(1 - \frac{i}{N}\right).$$

We compute

$$\begin{aligned} \sum_{i=\alpha N}^M \left(\log\left(1 - \frac{i}{N}\right) - \log\left(1 - \frac{i+1}{N}\right) \right) &= \sum_{i=\alpha N}^M \log\left(1 + \frac{1}{N - (i+1)}\right) \\ &\leq \sum_{i=\alpha N}^M \frac{1}{N - (i+1)} \\ &= O(\log n), \end{aligned}$$

so, noting that $\int \log s \, ds = s(\log s - 1)$,

$$\begin{aligned} 3 \sum_{i=\alpha N}^M \log\left(1 - \frac{i}{N}\right) &= 3N \int_{\alpha}^{(1-\varepsilon)} \log(1-t) \, dt + O(\log n) \\ &= 3N \int_{\varepsilon}^{(1-\alpha)} \log s \, ds + O(\log n) \\ &= 3N((1-\alpha)(\log(1-\alpha) - 1) - \varepsilon(\log \varepsilon - 1)) + O(\log n), \\ \exp\left(3 \sum_{i=\alpha N}^M \log\left(1 - \frac{i}{N}\right)\right) &= \left(\left(1 + O\left(n^{-c/\ell} \log n\right)\right) \frac{1-\alpha}{e}\right)^{3N(1-\alpha)}. \end{aligned}$$

For $b < c/\ell$, it follows that

$$\begin{aligned} |\mathcal{O}^*| &\geq \left(\left(1 - O\left(n^{-b}\right)\right) \frac{n^3(1-\alpha)^3}{6e^3}\right)^{(1-\alpha)N} \\ &= \left(\left(1 - O\left(n^{-b}\right)\right) \left(\frac{1-\alpha}{e}\right)^2 n\right)^{(1-\alpha)N} (N - \alpha N)!. \end{aligned}$$

Recalling that each $S^* \in \mathcal{O}^*$ can be completed, the desired result follows. \square

Now, it is extremely straightforward to prove Lemma 2.5.

Proof. Let $b \leq \min\{a, 1/2\}$ and $h \geq 2$ satisfy Theorem 3.2. By Theorem 3.1 we have

$$|\mathcal{O}^*(S)| \leq |\mathcal{S}^*(S)|(N - \alpha N)! \leq \left(\left(1 + O\left(n^{-b}\right)\right) \left(\frac{1-\alpha}{e}\right)^2 n\right)^{N(1-\alpha)} (N - \alpha N)!,$$

and by Theorem 3.2 we have

$$|\mathcal{O}^*(S')| \geq \left(\left(1 - O\left(n^{-b}\right)\right) \left(\frac{1-\alpha}{e}\right)^2 n\right)^{N(1-\alpha)} (N - \alpha N)!.$$

Dividing these bounds gives

$$\frac{|\mathcal{O}^*(S)|}{|\mathcal{O}^*(S')|} \leq \left(1 - O\left(n^{-b}\right)\right)^{N(1-\alpha)} \leq \exp\left(O\left(n^{2-b}\right)\right). \quad \square$$

4 An analysis of the triangle removal process

The triangle removal process starting from the complete graph has already been thoroughly analysed; see in particular the precise analysis by Bohman, Frieze and Lubetzky [6] and their simplified analysis in [5]. In this paper, we will need an analysis of the triangle removal process starting from a quasirandom graph. This basically follows from the aforementioned work of Bohman, Frieze and Lubetzky, but since the result we need was not stated in a concrete form in their papers, we provide our own (very simplified, and very crude) analysis in this section. We emphasise that this section contains no new ideas.

As in Section 2, let $N = \binom{n}{2}/3$ and assume that n is 1 or 3 mod 6. As previously introduced, the triangle removal process is defined as follows. We start with a graph G with say $3N - 3m$ edges, then iteratively delete (the edges of) a triangle chosen uniformly at random from all triangles in the remaining graph. Let

$$G = \mathbf{G}(m), \mathbf{G}(m+1), \dots$$

be the sequence of random graphs generated by this process. This process cannot continue forever, but we “freeze” the process instead of aborting it: if $\mathbf{G}(\mathbf{M})$ is the first graph in the sequence with no triangles, then let $\mathbf{G}(i) = \mathbf{G}(\mathbf{M})$ for $i \geq \mathbf{M}$.

Our objective in this section is to show that if G is quasirandom then the triangle removal process is likely to maintain quasirandomness and unlikely to freeze until nearly all edges are gone.

Theorem 4.1. *For all $h \geq 2$ and $a > 0$ there is $b(a, h) > 0$ such that the following holds. Let $n^{-a} \leq \varepsilon < 1/2$ and suppose G is a (ε, h) -quasirandom graph with $N - 3m$ edges. Then a.a.s. $\mathbf{M} \geq (1 - \varepsilon^b)N$ and moreover for each $m \leq i \leq (1 - \varepsilon^b)N$, the graph $\mathbf{G}(i)$ is (ε^b, h) -quasirandom.*

Note that K_n is $(O(1/n), h)$ -quasirandom for any fixed h , so in particular when we start the triangle removal process from $G = K_n$ it typically runs almost to completion.

To prove Theorem 4.1, it will be convenient to use Freedman’s inequality [18, Theorem 1.6], as follows. (This was originally stated for martingales, but it also holds for supermartingales with the same proof). Here and in what follows, we write $\Delta X(i)$ for the one-step change $X(i+1) - X(i)$ in a variable X .

Lemma 4.2. *Let $\mathbf{X}(0), \mathbf{X}(1), \dots$ be a supermartingale with respect to a filtration (\mathcal{F}_i) . Suppose that $|\Delta \mathbf{X}(i)| \leq K$ for all i , and let $V(i) = \sum_{j=0}^{i-1} \mathbb{E}[(\Delta \mathbf{X}(j))^2 \mid \mathcal{F}_j]$. Then for any $t, v > 0$,*

$$\Pr(\mathbf{X}(i) \geq \mathbf{X}(0) + t \text{ and } V(i) \leq v \text{ for some } i) \leq \exp\left(-\frac{t^2}{2(v + Kt)}\right).$$

Proof of Theorem 4.1. For a set A of at most h vertices, let $\mathbf{Y}_A(i) = |\bigcap_{w \in A} N_{\mathbf{G}(i)}(w)|$. Let $p(i) = (1 - i/N)$ and let $p^k(i) = (1 - i/N)^k$, so that $p^{|A|}(i)n$ is the predicted trajectory of each $\mathbf{Y}_A(i)$.

Fix some large C and small c to be determined. We will choose $b < c/(C+1)$ so that $e(i) := p(i)^{-C} \varepsilon^c \leq \varepsilon^b$ for $i \leq N(1 - \varepsilon^b)$. This means that if the conditions

$$\begin{aligned} \mathbf{Y}_A(i) &\leq p^{|A|}(i)n(1 + e(i)), \\ \mathbf{Y}_A(i) &\geq p^{|A|}(i)n(1 - e(i)), \end{aligned}$$

are satisfied for all A , then $\mathbf{G}(i)$ is $(e(i), h)$ -quasirandom (therefore (ε^b, h) -quasirandom).

Let \mathbf{T}' be the smallest index $i \geq m$ such that for some A , the above equations are violated (let $\mathbf{T}' = \infty$ if this never happens). Let $\mathbf{T} = \mathbf{T}' \wedge N(1 - \varepsilon^b)$. Define the stopped processes

$$\begin{aligned} \mathbf{Y}_A^+(i) &= \mathbf{Y}_A(i \wedge \mathbf{T}) - p^{|A|}(i \wedge \mathbf{T})n(1 + e(i \wedge \mathbf{T})), \\ \mathbf{Y}_A^-(i) &= -\mathbf{Y}_A(i \wedge \mathbf{T}) + p^{|A|}(i \wedge \mathbf{T})n(1 - e(i \wedge \mathbf{T})). \end{aligned}$$

We want to show that for each A and each $s \in \{+, -\}$, the process $\mathbf{Y}_A^s = (\mathbf{Y}_A^s(i), \mathbf{Y}_A^s(i+1), \dots)$ is a supermartingale, and then we want to use Lemma 4.2 and the union bound to show that a.a.s. each \mathbf{Y}_A^s only takes negative values.

To see that this suffices to prove Theorem 4.1, note that if $i < \mathbf{T}$ then by Proposition 2.8 the number of triangles in $\mathbf{G}(i)$ is

$$\mathbf{Q}(i) = (1 \pm O(e(i))) \frac{p^3(i)n^3}{6} > 0.$$

This means $\mathbf{T} \leq \mathbf{M}$, so the event that each \mathbf{Y}_A^s only takes negative values contains the event that each $\mathbf{G}(i)$ is non-frozen and sufficiently quasirandom for $i \leq N(1 - \varepsilon^b)$.

Let $\mathbf{R}_A(i) = \bigcap_{w \in A} N_{\mathbf{G}(i)}(w)$, so that $\mathbf{Y}_A(i) = |\mathbf{R}_A(i)|$. Fix A , and consider $x \in \mathbf{R}_A(i)$, for $i < \mathbf{T}$. The only way we can have $x \notin \mathbf{R}_A(i+1)$ is if we remove a triangle containing an edge

$\{x, w\}$ for some $w \in A$. Now, for each $w \in A$, the number of triangles in $\mathbf{G}(i)$ containing the edge $\{x, v\}$ is $(1 \pm O(e(i)))p^2(i)n$ by Proposition 2.8. The number of triangles containing x and more than one vertex of A is $O(1)$. So,

$$\begin{aligned} \Pr(x \notin \mathbf{R}_A(i+1)) &= \frac{1}{Q(i)} \left(\sum_{w \in A} (1 \pm O(e(i)))p^2(i)n - O(1) \right) \\ &= (1 \pm O(e(i))) \frac{|A|}{p(i)N}. \end{aligned}$$

For $i < \mathbf{T}$ we have $|\mathbf{R}_A(i)| = (1 \pm e(i))p^{|A|}(i)n$, so by linearity of expectation

$$\begin{aligned} \mathbb{E}[\Delta \mathbf{Y}_A(i) \mid \mathbf{G}(i)] &= -(1 \pm O(e(i))) \frac{|A|p^{|A|-1}(i)n}{N} \\ &= -\frac{|A|p^{|A|-1}(i)n}{N} + O\left(\frac{e(i)p^{|A|-1}(i)}{n}\right). \end{aligned}$$

Note also that we have the bound $\Delta \mathbf{Y}_A(i) \leq 2 = O(1)$ (with probability 1). Also, for fixed k , we have

$$\begin{aligned} \Delta p^k(i) &= \left(1 - \frac{i+1}{N}\right)^k - \left(1 - \frac{i}{N}\right)^k \\ &= \left(1 - \frac{i}{N}\right)^k \left(\left(\frac{N-i-1}{N-i}\right)^k - 1 \right) \\ &= p^k(i) \left(\left(1 - \frac{1}{N-i}\right)^k - 1 \right) \\ &= p^k(i) \left(-\frac{k}{N-i} + O\left(\frac{1}{(N-i)^2}\right) \right) \\ &= -\frac{kp^{k-1}(i)}{N} \left(1 + O\left(\frac{p(i)}{n^2}\right) \right) \\ &= -\frac{kp^{k-1}(i)}{N} + o\left(\frac{e(i)p^{k-1}(i)}{n^2}\right), \end{aligned}$$

and with ep^k denoting the pointwise product $i \mapsto e(i)p^k(i)$, we then have

$$\begin{aligned} \Delta(ep^k)(i) &= \varepsilon^c \Delta p^{k-C}(i) \\ &= \varepsilon^c \Theta\left(\frac{(C-k)p^{k-C-1}(i)}{N}\right) \\ &= \Theta\left(\frac{(C-k)e(i)p^{k-1}(i)}{n^2}\right). \end{aligned}$$

For large C it follows that

$$\mathbb{E}[\Delta \mathbf{Y}_A^+(i) \mid \mathbf{G}(i)] = \mathbb{E}[\Delta \mathbf{Y}_A(i) \mid \mathbf{G}(i)] - \Delta p^{|A|}(i)n - \Delta(ep^{|A|})(i)n \leq 0,$$

and similarly

$$\mathbb{E}[\Delta \mathbf{Y}_A^-(i) \mid \mathbf{G}(i)] \leq 0$$

for $i < \mathbf{T}$. (For $i \geq \mathbf{T}$ we trivially have $\Delta \mathbf{Y}_A^s(i) = 0$) Since each \mathbf{Y}_A^s is a Markov process, it follows that each is a supermartingale. Now, we need to bound $\Delta \mathbf{Y}_A^s(i)$ and $\mathbb{E}[(\Delta \mathbf{Y}_A^s(i))^2 \mid \mathbf{G}(i)]$,

which is easy given the preceding calculations. First, recalling that $\Delta \mathbf{Y}_A(i) = O(1)$ and noting that $\Delta p^k(i), \Delta(ep^k)(i) = O(1/N)$ we immediately have $|\Delta \mathbf{Y}_A^s(i)| = O(1)$. Noting in addition that $\mathbb{E}[\Delta \mathbf{Y}_A(i) | \mathbf{G}(i)] = O(1/n)$, we have

$$\mathbb{E}\left[(\Delta \mathbf{Y}_A^s(i))^2 \mid \mathbf{G}(i)\right] = O(\mathbb{E}[\Delta \mathbf{Y}_A^s(i) \mid \mathbf{G}(i)]) = O\left(\frac{1}{n}\right).$$

Since $T \leq N$, we also have

$$\sum_{i=0}^{\infty} \mathbb{E}\left[(\Delta \mathbf{Y}_A^s(i))^2 \mid \mathbf{G}(i)\right] = O\left(\frac{N}{n}\right) = O(n).$$

Provided $c < 1$ (and recalling that $\varepsilon < 1/2$), applying Lemma 4.2 with $t = e(m)p^{|A|}(m)n - \varepsilon p^{|A|}(m)n = \Omega(n\varepsilon^c)$ and $v = O(n)$ then gives

$$\Pr(\mathbf{Y}_A^s(i) > 0 \text{ for some } i) \leq \exp(-O(n\varepsilon^{2c})).$$

So, if $2c < \log_{\varepsilon} n \leq a$, the union bound over all A, s finishes the proof. \square

5 Perfect matchings via absorbers

In this section we prove Theorem 1.1 using Theorem 2.4 and the absorbing method. First, we define our absorbers, which are small rooted hypergraphs that can contribute to a perfect matching in two different ways. We are careful to make the definition in such a way that absorbers can be shown to appear in $\mathbb{R}(n, \alpha N)$ with extremely high probability.

Definition 5.1. An *absorber* for an ordered triple (x, y, z) is a set of hyperedges of the form

$$\{\{x, x_1, x_2\}, \{y, y_1, y_2\}, \{z, z_1, z_2\}, \{w_x, w_y, w_z\}, \{x_1, y_2, w_z\}, \{y_1, z_2, w_x\}, \{z_1, x_2, w_y\}\}.$$

We call x, y, z the *rooted vertices* and we call the other nine vertices the *external vertices*. Also, we say the three hyperedges containing rooted vertices are *rooted hyperedges*, and the other four hyperedges containing only external vertices are *external hyperedges*. Note that an absorber has a perfect matching on its full set of 12 vertices (we call this the *covering* matching), and it also has a perfect matching on its external vertices (we call this the *non-covering* matching). See Figure 1.

Absorbers are the basic building blocks for a larger structure which will eventually allow us to complete an almost-perfect matching into a perfect matching. The relative positions of the absorbers in this structure will be determined by a “template” with a “resilient matching” property, as follows.

Lemma 5.2. *For any sufficiently large n , there exists a 3-uniform hypergraph T with $10n$ vertices, $120n$ hyperedges and a set Z of $2n$ vertices, such that if we remove any n vertices from Z , the resulting hypergraph has a perfect matching. We call T a resilient template and we call Z its flexible set.*

To prove Lemma 5.2 we use the following lemma of Montgomery [37, Lemma 2.8]

Lemma 5.3. *For any sufficiently large n , there exists a bipartite graph R with vertex parts X and $Y \sqcup Z$, with $|X| = 3n$, $|Y| = |Z| = 2n$, and maximum degree 40, such that if we remove any n vertices from Z , the resulting bipartite graph has a perfect matching.*

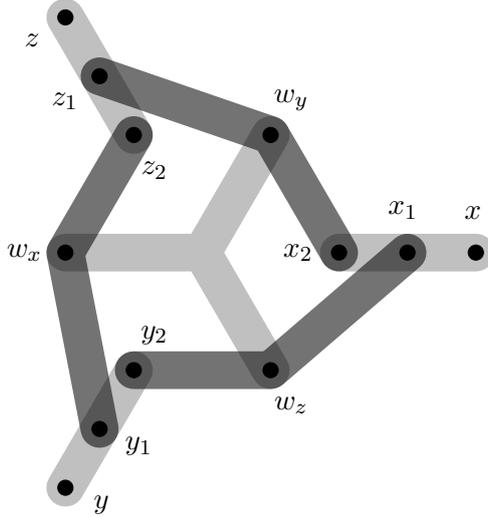


Figure 1. An illustration of an absorber for (x, y, z) . The light hyperedges are the covering matching and the dark hyperedges are the non-covering matching.

Proof of Lemma 5.2. Consider the bipartite graph R from Lemma 5.3 on the vertex set $X \sqcup (Y \sqcup Z)$. Note that R has at most $40|X| = 120n$ edges (we can assume it has exactly $120n$ edges, because adding edges does not affect the resilient matching property). Add a set W of $|X|$ new vertices and put a perfect matching between W and X , to obtain a $10n$ -vertex tripartite graph R' . Now, define a hypergraph T on the same vertex set by putting a hyperedge for each 3-vertex path running through all three parts of R' (we call such paths *special paths*). Note that an edge in R can be uniquely extended to a special path in R' , so T has $120n$ hyperedges. Moreover, a matching in R can always be extended to a vertex-disjoint union of special paths in R' , so T is a resilient template with flexible set Z . \square

Now we can describe our absorbing structure in its entirety.

Definition 5.4. An *absorbing structure* is a 3-uniform hypergraph H of the following form. Consider a resilient template T and put externally vertex-disjoint absorbers on each hyperedge of T , introducing 9 new vertices for each. Then delete the edges of T . That is, the template just describes the relative positions of the absorbers, its hyperedges are not actually in the absorbing structure.

Note that an absorbing structure with a flexible set of size $2n$ has $10n + 9 \times 120n = O(n)$ vertices and $7 \times 120n = O(n)$ hyperedges. An absorbing structure H has the same crucial property as the resilient template T that defines it: if we remove half of the vertices of the flexible set then what remains of H has a perfect matching. Indeed, after this removal we can find a perfect matching M of T , then our perfect matching of H can be comprised of the covering matching of the absorber on each hyperedge of M and the non-covering matching for the absorber on each other hyperedge of T . The existence of an absorbing structure, in addition to some very weak pseudorandomness conditions, allows us to find perfect matchings in a Steiner triple system, as follows.

Lemma 5.5. Consider a Steiner triple system S with vertex set V , $|V| = n \equiv 3 \pmod{6}$, satisfying the following conditions for some $\delta = \delta(n) = o(1/\log n)$ and fixed $\beta > 0$.

1. There is an absorbing structure H in S with a flexible set Z of size $6 \lfloor \delta^2 n \rfloor$.
2. For at most δn of the vertices $v \in V \setminus V(H)$, we have $|\{\{x, y\} \subseteq Z : \{v, x, y\} \in E(S)\}| < 6\delta^5 n$. That is to say, few vertices have unusually low degree into the flexible set Z , in S .

3. Every vertex subset $W \subseteq V$ with $|W| \geq 3\delta^5 n$ induces at least $(1 - \beta)|W|^3/(6n)$ hyperedges.

Then S has

$$\left(\frac{n}{2e^2}(1 - \beta - o(1))\right)^{n/3} \quad (6)$$

perfect matchings.

Proof. We first describe a procedure to build a perfect matching (provided n is sufficiently large), then we count the number of ways to perform this procedure.

Let $U \subseteq V \setminus V(H)$ be the set of vertices with unusually low degree into Z , as per condition 2. The first few hyperedges of our matching will cover U , and will not use any vertices of H . We can in fact choose these hyperedges one-by-one in a greedy fashion: considering each $v \in U$ in any order, note that v is in $(n - 1)/2$ hyperedges of the Steiner triple system S , and at most $3|U| + v(H) \leq 18\delta n + 3(10 + 9 \times 120)\lceil \delta^2 n \rceil = o(n)$ of these hyperedges involve a vertex of H or a vertex in the hyperedges chosen so far.

Now, let $n' \geq n - v(H) - 3|U| \geq n - 19\delta n$ be the number of vertices in $V \setminus V(H)$ remaining unmatched. Note that $v(H) = 3(10 + 9 \times 120)\lceil \delta^2 n \rceil$ is divisible by 3, so n' is divisible by 3 as well. We next use condition 3 to repeatedly choose a hyperedge induced by the remaining unmatched vertices in $V \setminus V(H)$ until there are only $3\lceil \delta^5 n \rceil$ such vertices remaining unmatched. (This means we are choosing $m = (n' - 3\lceil \delta^5 n \rceil)/3 \geq n'/3 - 20\delta n$ hyperedges). We call this step the *main step*.

Next, we greedily extend our matching to cover the remaining vertices in $V \setminus V(H)$. Considering each uncovered $v \in V \setminus V(H)$ in any order, recall that $v \notin U$ so by condition 2 there are at least $6\delta^5 n$ hyperedges of S containing v and two vertices of Z . We can therefore choose such a hyperedge avoiding the (fewer than $2 \times 3\lceil \delta^5 n \rceil$) vertices in Z used so far, to extend our matching. We have now covered all of $V \setminus V(H)$ and $6\lceil \delta^5 n \rceil$ vertices of Z ; we can then repeatedly apply condition 3 to the uncovered vertices in Z to extend our matching to cover half of Z . By the crucial property of an absorbing structure, we can find a perfect matching on the remaining vertices, completing our perfect matching of S .

Now we analyse the number of ways to perform the above procedure. It actually suffices to count the number of ways to make the ordered sequence of choices in the main step, which is at least

$$\begin{aligned} \prod_{i=1}^m (1 - \beta) \frac{(n' - 3i)^3}{6n} &= \left(\frac{(1 - \beta)(n')^3}{6n}\right)^m \exp\left(\sum_{i=1}^m 3 \log\left(1 - 3\frac{i}{n'}\right)\right) \\ &\geq \left((1 - \beta - 3 \times 19\delta) \frac{n'^2}{6}\right)^{n'/3} n^{-20\delta n} \exp\left(n' \sum_{i=1}^m \frac{3}{n'} \log\left(1 - 3\frac{i}{n'}\right)\right). \end{aligned} \quad (7)$$

Now, noting that $\int \log s \, ds = s(\log s - 1)$, we have the Riemann sum approximation

$$\begin{aligned} \sum_{i=1}^m \frac{3}{n'} \log\left(1 - 3\frac{i}{n'}\right) &= \int_0^{m/n'} 3 \log(1 - 3t) \, dt + o(1) \\ &= \int_0^{1/3 - o(1)} 3 \log(1 - 3t) \, dt + o(1) \\ &= \int_{o(1)}^1 \log s \, ds + o(1) \\ &= -1 + o(1). \end{aligned}$$

So, the expression in (7) is at least

$$\left((1 - \beta + O(\delta)) \frac{n^2}{6}\right)^{n/3} n^{O(\delta n)} e^{-n - o(n)} = \left((1 - \beta + o(1)) \frac{n^2}{6e^3}\right)^{n/3}.$$

(Recall that we are assuming $\delta = o(1/\log n)$). This is a lower bound for the number of *ordered* perfect matchings in S . So, we divide by $(n/3)! = ((1 + o(1))n/3e)^{n/3}$ (using Stirling's approximation) to obtain (6). \square

5.1 Absorbing conditions in the triangle removal process

In this section we prove that the conditions in Lemma 5.5 (for say $\delta = 1/\log^2 n$ and arbitrarily small $\beta > 0$) hold in a random Steiner triple system, proving Theorem 1.1. We do this using Theorem 2.4, showing that the same properties hold with probability $1 - \exp(-\tilde{\Omega}(n^2))$ in the triangle removal process. (A tilde over asymptotic notation indicates that polylogarithmic factors are being ignored).

Fix a large constant $h \in \mathbb{N}$ (we will see later exactly how large it should be), and fix small $\alpha > 0$ (which we will assume is small enough to satisfy certain inequalities later in the proof). Fix a set Z of $6\lfloor \delta^2 n \rfloor$ vertices (say $Z = [6\lfloor \delta^2 n \rfloor]$); we will eventually find an absorbing structure with Z as a flexible set.

5.1.1 High degree into the flexible set

If condition 2 is violated, there is a set W of $\lfloor \delta n \rfloor$ vertices outside Z each with degree less than $6\delta^5 n$ into Z . There are then fewer than $6\delta^6 n^2$ hyperedges with one vertex in W and two vertices in Z . We show that it is extremely unlikely that there is a set W with this property.

We will use Lemma 2.10, so let $\mathbf{S}^* \sim \mathbb{G}^*(n, \alpha/n)$ be obtained from $\mathbf{G} \sim \mathbb{G}(n, \alpha/n)$. Consider a set W of $\lfloor \delta n \rfloor$ vertices outside Z . Let \mathbf{Y} be the number of hyperedges of \mathbf{S}^* with one vertex in W and two vertices in Z . That is to say, \mathbf{Y} is the number of such hyperedges in \mathbf{G} that are *isolated* in the sense that they do not intersect any other hyperedge of \mathbf{G} in more than one vertex. There are $\Theta((\delta n)(\delta^2 n)^2) = \Theta(\delta^5 n^3)$ possible hyperedges and each is present and isolated with probability $(\alpha/n)(1 - \alpha/n)^{O(n)} = \Theta(n^{-1})$, so $\mathbb{E}\mathbf{Y} = \Theta(\delta^5 n^2)$. Now, adding a hyperedge to \mathbf{G} can increase \mathbf{Y} by at most 1, and removing a hyperedge can increase \mathbf{Y} by at most 3 (by making three hyperedges isolated). So, by Theorem 2.11,

$$\Pr(\mathbf{Y} \leq 2\delta^5 n^2) \leq \Pr(|\mathbf{Y} - \mathbb{E}\mathbf{Y}| \geq \Theta(\delta^5 n^2)) \leq \exp\left(-\Omega\left(\frac{(\delta^5 n^2)^2}{3^2 \binom{n}{3} \alpha/n + 3\delta^5 n^2}\right)\right) = \exp(-\tilde{\Omega}(n^2)).$$

Since there are no more than 2^n choices for W , we can use the union bound, Lemma 2.10 (with $S = \emptyset$) and Theorem 2.4 (with no conditioning; that is, $\mathcal{Q} = \mathcal{O}_{\alpha N} \cup \{*\}$) to prove that condition 2 of Lemma 5.5 holds a.a.s. in a random Steiner triple system.

5.1.2 Density in subsets

Now we deal with condition 3. It is not immediately clear that one can consider just the first few hyperedges of a random Steiner triple system as in Section 5.1.1, but the key observation is that in a random *ordered* Steiner triple system, by symmetry the first αN hyperedges have the same distribution as the hyperedges corresponding to any other choice of αN indices.

With \mathbf{S}^* and \mathbf{G} as in Section 5.1.1, consider a set $W \subseteq V$ with $|W| \geq 3\delta^5 n$ and redefine \mathbf{Y} to be the number of hyperedges of \mathbf{S}^* included in W (which is the number of such hyperedges in \mathbf{G} that are isolated). There are $(1 + o(1))|W|^3/6$ possible hyperedges, and each is present and isolated in \mathbf{G} with probability $(\alpha/n)(1 - \alpha/n)^{O(n)} = (\alpha/n)(1 - O(\alpha))$. Reasoning as in Section 5.1.1, with probability $1 - \exp(-\tilde{\Omega}(n^2))$ we have $\mathbf{Y} \geq \alpha(1 - O(\alpha))|W|^3/(6n)$. The union bound, Lemma 2.10 and Theorem 2.4 prove that if \mathbf{S} is a random Steiner triple system, then a.a.s. every appropriate subset W induces at least $\alpha(1 - O(\alpha))|W|^3/(6n)$ hyperedges in $\mathbf{S}_{\alpha N}$. By symmetry this property also holds a.a.s. in $\mathbf{S}_{k\alpha N} \setminus \mathbf{S}_{(k-1)\alpha N}$ for each $k \leq 1/\alpha$. So,

a.a.s. every W induces a total of $(1 - O(\alpha))|W|^3/(6n)$ hyperedges in \mathcal{S} . For β a large multiple of α , condition 3 of Lemma 5.5 is then satisfied.

5.1.3 Absorbers

Finally we show how to find an absorbing structure for condition 1, which is much more involved. The first step is to show that there are many absorbers rooted on every triple of vertices. We cannot hope to do this by naively analysing $\mathbb{G}(n, \alpha/n)$ and using Theorem 2.4 as in Sections 5.1.1 and 5.1.2, because the probability that a vertex is isolated is already too large. Instead we must use Theorem 2.4 in its full generality, conditioning on the quasirandomness of the first few steps of the triangle removal process.

Let a be small enough for Theorem 2.4, and let

$$\mathcal{Q} = \{*\} \cup \left\{ S \in \mathcal{O}_{2\alpha N} : S_{\alpha N} \in \mathcal{O}_{\alpha N}^{n^{-a}, h} \right\} \supseteq \mathcal{O}_{2\alpha N}^{n^{-a}, h}.$$

Let $\mathcal{S} \sim \mathbb{R}(n, 2\alpha N)$, and condition on any $\mathcal{S}_{\alpha N} = S \in \mathcal{O}_{\alpha N}^{n^{-a}, h}$. We will use Lemma 2.10 to analyse $\mathcal{S} \setminus \mathcal{S}_{\alpha N} \sim \mathbb{R}(S, \alpha N)$ via $\mathbb{G}^*(S, \alpha/n)$. So, let $\mathcal{S}^* \sim \mathbb{G}^*(S, \alpha/n)$ be obtained from $\mathcal{G} \sim \mathbb{G}(S, \alpha/n)$.

By quasirandomness, every vertex has degree $(1 \pm n^{-a})(1 - \alpha)n$ in $G(S)$, so every vertex is in $(1 \pm n^{-a})\alpha n/2 = \Omega(\alpha n)$ hyperedges of S . Consider vertices x, y, z . Say an *absorber-extension* is a collection of four hyperedges which can be combined with three hyperedges of S incident to x, y, z , to form an absorber on (x, y, z) . (S provides the rooted hyperedges of an absorber, and an absorber-extension provides the external hyperedges). Let \mathbf{Y} be the maximum size of a hyperedge-disjoint collection of absorber-extensions in \mathcal{S}^* ; equivalently, \mathbf{Y} is the maximal size of a collection of disjoint isolated absorber-extensions in \mathcal{G} . This particular choice of random variable is crucial, and allows us to use Theorem 2.11. (The idea comes from a similar random variable used by Bollobás [7]). Adding a hyperedge to \mathcal{G} can increase the size of a maximal collection of hyperedge-disjoint absorber-extensions by at most one, and removing a hyperedge can cause at most three hyperedge-disjoint absorber-extensions to become isolated. So, changing the presence of a hyperedge in \mathcal{G} can change \mathbf{Y} by at most 3, as in Sections 5.1.1 and 5.1.2.

Claim 5.6. *If h is large enough and α is small enough, then $\mathbb{E}\mathbf{Y} = \Omega(n^2)$.*

To prove Claim 5.6 we'll need a simple lemma about quasirandom graphs.

Lemma 5.7. *Suppose G is an n -vertex $(o(1), 2)$ -quasirandom graph with density γ satisfying $\gamma, 1 - \gamma = \Omega(1)$. Consider sets X, Y, Z of disjoint ordered pairs of vertices of G , such that $|X|, |Y|, |Z| = \Omega(n)$. Then there are at least $\gamma^3|X||Y||Z|/8 - o(n^3)$ choices of $(x_1, x_2) \in X$, $(y_1, y_2) \in Y$, $(z_1, z_2) \in Z$ such that $\{x_1, y_2\}$, $\{y_1, z_2\}$ and $\{z_1, x_2\}$ are all edges of G .*

Proof. This quasirandomness condition implies quasirandomness in the Chung–Graham–Wilson sense (see for example [3, Theorem 9.3.2]). In particular, for any (not necessarily disjoint) vertex sets A, B , there are at least $\gamma|A||B|/2 - o(n^2)$ edges between them, and all but $o(n)$ vertices of A have at least $\gamma|B|/2 - o(n)$ neighbours in B . (To prove this latter fact, note that if for any $\varepsilon = \Omega(1)$ there were a set $A' \subseteq A$ of εn vertices that each had $\gamma|B| - \varepsilon n$ neighbours in B , then there would be at most $\gamma|A'||B|/2 - \varepsilon^2 n^2$ edges between A' and B , contradicting quasirandomness).

So, for $|X| - o(n)$ choices of $(x_1, x_2) \in X$, there are subsets $Y_{x_1, x_2} \subseteq Y$, $Z_{x_1, x_2} \subseteq Z$ with $|Y_{x_1, x_2}| = (\gamma|Y|/2 - o(n))$ and $|Z_{x_1, x_2}| = (\gamma|Z|/2 - o(n))$, such that $\{x_1, y_2\}$ and $\{z_1, x_2\}$ are edges of G for each $(y_1, y_2) \in Y_{x_1, x_2}$, $(z_1, z_2) \in Z_{x_1, x_2}$. For each such choice of x_1, x_2 , there are $\gamma|Y_{x_1, x_2}||Z_{x_1, x_2}|/2 - o(n^2)$ choices of $(y_1, y_2) \in Y_{x_1, x_2}$, $(z_1, z_2) \in Z_{x_1, x_2}$ such that $\{y_1, z_2\}$ is an edge of G . In summary, there are at least

$$(|X| - o(n))(\gamma(\gamma|Y|/2 - o(n))(\gamma|Z|/2 - o(n))/2 - o(n^2)) = \gamma^3|X||Y||Z|/8 - o(n^3)$$

suitable choices of $(x_1, x_2) \in X$, $(y_1, y_2) \in Y$, $(z_1, z_2) \in Z$. □

Now we prove Claim 5.6.

Proof of Claim 5.6. Let \mathbf{X} be the total number of isolated absorber-extensions in \mathbf{G} and let \mathbf{Z} be the number of pairs of hyperedge-intersecting absorber-extensions in \mathbf{G} . We can obtain a collection of disjoint isolated absorber-extensions by considering the collection of all isolated absorber-extensions and deleting one from each intersecting pair, so $\mathbf{Y} \geq \mathbf{X} - \mathbf{Z}$ and $\mathbb{E}\mathbf{Y} \geq \mathbb{E}\mathbf{X} - \mathbb{E}\mathbf{Z}$. We first estimate $\mathbb{E}\mathbf{X}$.

First we show that there are $\Theta(\alpha^3 n^6)$ possible absorber-extensions not conflicting with S . To this end, we first show that there are $\Theta((\alpha n)^3)$ ways to choose three disjoint hyperedges $e_x = \{x, x_1, x_2\}, e_y = \{y, y_1, y_2\}, e_z = \{z, z_1, z_2\} \in E(S)$, such that $\{x_1, y_2\}, \{y_1, z_2\}$ and $\{z_1, x_2\}$ are all edges of $G(S)$. Indeed, let X (respectively, Y and Z) be the set of all pairs of vertices $\{x_1, x_2\}$ (respectively, $\{y_1, y_2\}$ and $\{z_1, z_2\}$) such that $\{x, x_1, x_2\} \in E(S)$ (respectively, $\{y, y_1, y_2\} \in E(S)$ and $\{z, z_1, z_2\} \in E(S)$). Note that $|X|, |Y|, |Z| = \Theta(\alpha n)$, arbitrarily choose an ordering for each of the constructed pairs, and apply Lemma 5.7 to show that there are $\Theta((1 - \alpha)^3 (\alpha n)^3) = \Theta((\alpha n)^3)$ ways to choose e_x, e_y, e_z . Then, the number of ways to choose an absorber-extension compatible with e_x, e_y, e_z is precisely the number of copies in $G(S)$ of a certain graph F rooted on the vertices of e_x, e_y, e_z . (Specifically, F is the graph obtained by taking the external hyperedges of the hypergraph in Definition 5.1 and replacing each hyperedge with a triangle on its vertex set). Provided h is large enough, by Proposition 2.8 the number of suitable copies of F is $(1 - \alpha)^{O(1)} n^3 = \Theta(n^3)$. (Note that strictly speaking we are over-counting, because it is possible that an absorber-extension can contribute to multiple different absorbers, but this constant factor will not bother us).

The probability that each possible absorber-extension appears and is isolated in \mathbf{G} is

$$\Theta\left((\alpha/n)^4 (1 - \alpha/n)^{O(n)}\right) = \Theta(\alpha^4 n^{-4}),$$

so $\mathbb{E}\mathbf{X} = \Theta(\alpha^7 n^2)$. Now, we estimate $\mathbb{E}\mathbf{Z}$. It will be convenient to consider *labelled* absorbers and absorber-extensions; for the hypergraph in Definition 5.1 denote its hyperedges (in the same order as in Definition 5.1) by

$$e_x, e_y, e_z, e_*, e_1, e_2, e_3.$$

There are several possibilities for a hyperedge-intersecting pair of distinct absorber-extensions.

- Suppose they intersect in one hyperedge. Each such pair appears with probability $O\left((\alpha/n)^7\right)$.
 - Suppose the intersecting hyperedge is e_* for one of the absorber-extensions (say the second). There are $O\left((\alpha n)^6 n^3\right)$ possibilities for such a pair of absorber-extensions, as follows. Choose the first absorber-extension in one of $O\left((\alpha n)^3 n^3\right)$ ways, and choose one of its hyperedges which will intersect with the second absorber-extension. Then, the second absorber-extension is determined by its choices for e_x, e_y, e_z .
 - Suppose the intersecting hyperedge is say e_1 for the second absorber-extension. There are $O\left((\alpha n)^4 n^5\right)$ possible such pairs, as follows. After choosing the first absorber-extension, and choosing its hyperedge which will be intersecting, the choices for e_x and e_y for the second absorber-extension are already determined (if a suitable choice exists at all), because in S each pair of vertices is included in at most one hyperedge. One of the vertices of e_* is already also determined, so the second absorber-extension is determined by a choice of e_z and two vertices of e_* .
- Suppose they intersect in two hyperedge. Each such pair appears with probability $O\left((\alpha/n)^6\right)$.

- Suppose the intersecting hyperedges are say e_* and e_1 for the second absorber-extension. There are $O((\alpha n)^4 n^3)$ possibilities for such a pair: after choosing the first absorber-extension and its hyperedges which will be intersecting, the second absorber-extension is determined by its choice for e_z .
- Suppose the intersecting hyperedges are say e_1 and e_2 for the second absorber-extension. There are $O((\alpha n)^3 n^4)$ possibilities for such a pair: after choosing the first absorber-extension and its hyperedges which will be intersecting, the second absorber-extension is determined by a single vertex for e_* .
- Note that choosing three of e_1, e_2, e_3, e_* determines the other, so there are only $O((\alpha n)^3 n^3)$ possibilities for a pair of absorber-extensions intersecting in three hyperedges. Each such pair appears with probability $O((\alpha/n)^5)$.

In summary (for small α), we have

$$\begin{aligned} \mathbb{E}Z &= O\left(\left((\alpha n)^6 n^3 + (\alpha n)^4 n^5\right)(\alpha/n)^7 + \left((\alpha n)^4 n^3 + (\alpha n)^3 n^4\right)(\alpha/n)^6 + (\alpha n)^3 n^3 (\alpha/n)^5\right) \\ &= O(\alpha^{11} n^2). \end{aligned}$$

So, $\mathbb{E}Y \geq \Theta(\alpha^7 n^2) - O(\alpha^{11} n^2) = \Theta(\alpha^7 n^2)$. \square

As in Sections 5.1.1 and 5.1.2, Theorem 2.11 proves that $Y = \Omega(n^2)$ with probability $1 - \exp(-\Omega(n^2))$. Note that if there are $\Omega(n^2)$ hyperedge-disjoint absorber-extensions then there must in fact be $\Omega(n)$ externally vertex-disjoint absorbers rooted on x, y, z . We can find these greedily; each vertex is involved in only $O(n)$ hyperedges of S , so removing $O(1)$ vertices from consideration results in at most $O(n)$ hyperedges being removed from consideration. By the union bound, Lemma 2.10 and Theorem 2.4 (with \mathcal{Q} as defined at the beginning of this subsection), it follows that in a random Steiner triple system, a.a.s. every triple of vertices has $\Omega(n)$ externally vertex-disjoint absorbers.

If a Steiner triple system has this property, then we can very straightforwardly greedily build an absorbing structure with flexible set Z , as follows. Choose a resilient template T on the vertex set of S , such that the flexible set is Z . Consider each hyperedge (x, y, z) of T in any order, and greedily choose an absorber in S rooted on (x, y, z) , each of whose external vertices is disjoint to the template and all absorbers chosen so far. We have proved that condition 1 of Lemma 5.5 holds a.a.s. in a random Steiner triple system, completing the proof of Theorem 1.1.

6 An upper bound for the number of perfect matchings

In this section we prove Theorem 1.2, with the entropy method. Recall the definition of entropy and its basic properties from Section 3.

Proof of Theorem 1.2. Let \mathcal{M} be the set of perfect matchings in S . Consider a uniformly random $M \in \mathcal{M}$, so that $H(M) = \log |\mathcal{M}|$ is the entropy of M . Let M_v be the hyperedge of M containing the vertex v , so that the sequence $(M_v)_{v \in [n]}$ determines M . For any ordering on the vertices of S ,

$$H(M) = \sum_{v \in V} H(M_v \mid M_{v'} : v' < v). \quad (8)$$

Now, a sequence $\lambda \in [0, 1]^n$ with all λ_v distinct induces an ordering on $[n]$, with $v' < v$ when $\lambda_{v'} > \lambda_v$. Let $R_v(\lambda)$ be 1 plus the number of hyperedges $e \neq M_v$ containing v in S such that $\lambda_{v'} < \lambda_v$ for all $v' \in (\bigcup_{z \in e} M_z) \setminus \{v\}$. (In particular, $R_v(\lambda) = 1$ if $\lambda_{v'} > \lambda_v$ for some

$v' \in \mathbf{M}_v \setminus \{v\}$, in which case \mathbf{M}_v is determined by the information $(\mathbf{M}_{v'} : \lambda_{v'} > \lambda_v)$. Note that $\mathbf{R}_v(\lambda)$ is an upper bound on $|\text{supp}(\mathbf{M}_v | \mathbf{M}_{v'} : \lambda_{v'} > \lambda_v)|$, and therefore

$$H(\mathbf{M}_v | \mathbf{M}_{v'} : \lambda_{v'} > \lambda_v) \leq \mathbb{E}[\log \mathbf{R}_v(\lambda)]. \quad (9)$$

Let $\boldsymbol{\lambda} = (\lambda_v)_{v \in [n]}$ be a sequence of independent random variables, where each λ_v has the uniform distribution in $[0, 1]$. It follows from (8) and (9) that

$$H(\mathbf{M}) \leq \sum_{v \in [n]} \mathbb{E}[\log \mathbf{R}_v(\boldsymbol{\lambda})].$$

Next, for any $M \in \mathcal{M}$ and $\lambda_v \in [0, 1]$, let

$$R_v^{M, \lambda_v} = \mathbb{E}[\mathbf{R}_v(\boldsymbol{\lambda}) | \mathbf{M} = M, \lambda_v = \lambda_v, \lambda_{v'} < \lambda_v \text{ for all } v' \in \mathbf{M}_v \setminus \{v\}].$$

Now, there are $(n-1)/2$ hyperedges in S containing v , and for each such hyperedge $e = \{x, y, v\}$ other than M_v , note that $M_x \neq M_y$ (because e and M_x are different hyperedges of a Steiner triple system and can therefore intersect in at most one vertex). So, $|(\bigcup_{z \in e} M_z) \setminus M_v| = 6$ and by linearity of expectation,

$$R_v^{M, \lambda_v} = 1 + ((n-1)/2 - 1)\lambda_v^6.$$

By Jensen's inequality,

$$\mathbb{E}[\log \mathbf{R}_v(\boldsymbol{\lambda}) | \mathbf{M} = M, \lambda_v = \lambda_v, \lambda_{v'} < \lambda_v \text{ for all } v' \in \mathbf{M}_v \setminus \{v\}] \leq \log R_v^{M, \lambda_v},$$

and

$$\Pr(\lambda_{v'} < \lambda_v \text{ for all } v' \in \mathbf{M}_v \setminus \{v\} | \lambda_v = \lambda_v) = \lambda_v^2,$$

so

$$\mathbb{E}[\log \mathbf{R}_v(\boldsymbol{\lambda}) | \mathbf{M} = M, \lambda_v = \lambda_v] \leq \lambda_v^2 \log R_v^{M, \lambda_v} + (1 - \lambda_v^2) \log 1 = \lambda_v^2 \log R_v^{M, \lambda_v}.$$

It follows that

$$\begin{aligned} \mathbb{E}[\log \mathbf{R}_v(\boldsymbol{\lambda}) | \mathbf{M} = M] &\leq \mathbb{E}[\lambda_v^2 \log R_v^{M, \lambda_v}] \\ &= \int_0^1 \lambda_v^2 \log(1 + ((n-1)/2 - 1)\lambda_v^6) d\lambda_v \\ &= \frac{1}{3}(\log(n/2) - 2) + o(1), \end{aligned}$$

using (5) (from the proof of Theorem 3.1 in Section 3). We conclude that

$$\begin{aligned} \log |\mathcal{M}| &= H(\mathbf{M}) \\ &\leq \sum_{v \in [n]} \mathbb{E}[\log \mathbf{R}_v(\boldsymbol{\lambda})] \\ &\leq \frac{n}{3}(\log(n/2) - 2 + o(1)), \end{aligned}$$

which is equivalent to the theorem statement. \square

7 Latin squares

In this section we sketch how one should adapt the methods in this paper to prove Theorem 1.3.

A partial Steiner triple system is a collection of edge-disjoint triangles in K_n , whereas a partial Latin square is a collection of edge-disjoint triangles in the complete tripartite graph $K_{n,n,n}$. Let $V = V_1 \sqcup V_2 \sqcup V_3$ be the tripartition of $K_{n,n,n}$. We say a subgraph $G \subseteq K_{n,n,n}$ with m edges between each pair of parts is (ε, h) -quasirandom if for each $i \in \{1, 2, 3\}$, every set $A \subseteq V \setminus V_i$ with $|A| \leq h$ has $(1 \pm \varepsilon)\binom{m/n^2}{|A|}$ common neighbours in V_i . The following result is a special case of [28, Theorem 1.5].

Theorem 7.1. *There are $h \in \mathbb{N}$, $\varepsilon_0, a \in (0, 1)$ and $n_0, \ell \in \mathbb{N}$ such that the following holds. Suppose $n \geq n_0$, $m/n^2 \geq n^{-a}$ and $\varepsilon \leq \varepsilon_0(m/n^2)^\ell$, and suppose that $G \subseteq K_{n,n,n}$ is an (ε, h) -quasirandom graph with m edges between each pair of parts, that arises as the set of edges of $K_{n,n,n}$ not covered by the triples of some partial Latin square. Then the edges of G can be decomposed into triangles.*

With our new notion of quasirandomness and Theorem 7.1 playing the role of Theorem 3.3 we can then adapt the proofs in Sections 3 and 4 in a straightforward manner to prove the obvious Latin squares counterpart to Lemma 2.5. This allows us to prove the Latin squares counterpart to Theorem 2.4.

Now we outline what should be adapted from the arguments in Section 5 for the Latin squares case. The definition of an absorber can remain the same, noting that the hypergraph in Definition 5.1 is tripartite (and the tripartition can be chosen to have x, y, z in different parts). The definition of a resilient template should be adapted slightly: we define a resilient template to be a tripartite hypergraph H (with tripartition $V(H) = V_1(H) \sqcup V_2(H) \sqcup V_3(H)$, say) with a flexible set Z , such that each $Z_i = V_i(H) \cap Z$ has the same size, and such that if half the vertices of each Z_i are removed, then the remaining hypergraph has a perfect matching. To prove a counterpart of Lemma 5.2 we can just use three vertex-disjoint copies of the tripartite hypergraph in the proof of Lemma 5.2 (one copy for each Z_i). The counterpart of Lemma 5.5 is as follows (with virtually the same proof).

Lemma 7.2. *Consider an order- n Latin square L (with tripartition $V = V_1 \sqcup V_2 \sqcup V_3$) satisfying the following properties for some $\delta = \delta(n) = o(1/\log n)$ and fixed $\beta > 0$.*

1. *There is an absorbing structure H in L with at most δn vertices and a flexible set Z such that each $Z_i = V_i \cap Z = 2\lfloor \delta^2 n \rfloor$.*
2. *For at most δn of the vertices $v \in V_1$, we have $|\{(x, y) \in Z_2 \times Z_3 : (v, x, y) \in E(L)\}| < 6\delta^5 n$, and the analogous statements hold for degrees of $v \in V_2$ and $v \in V_3$ into $Z_1 \times Z_3$ and $Z_1 \times Z_2$ respectively.*
3. *For any choice of $W_i \subseteq V_i$ such that each $|W_i| \geq \delta^5 n$, there are at least $(1 - \beta)|W_1||W_2||W_3|/n$ hyperedges in $W_1 \times W_2 \times W_3$.*

Then L has

$$\left(\frac{n}{e^2}(1 - \beta - o(1))\right)^n$$

transversals.

One can then use Lemma 7.2 to prove Theorem 1.3 in basically the same way as the proof of Theorem 1.1 in Section 5.1. We remark that in Lemma 5.7 we used an implication between certain different notions of graph quasirandomness; to straightforwardly adapt this lemma to the Latin squares case we would need an analogous implication between notions of ‘‘bipartite quasirandomness’’. Such an implication is well-known to hold (see for example [3, Exercise 9.10]).

8 Concluding remarks

In this paper we introduced a new method for analysing random Steiner triple systems, and we used it to prove that almost all Steiner triple systems have many perfect matchings. There are many interesting open questions that remain.

- We believe the most interesting problem that seems approachable by our methods is to prove that almost all Steiner triple systems (and Latin squares) can be *decomposed* into disjoint perfect matchings (transversals). Since the first version of this paper, together with

Ferber we [17] proved an approximate version of this fact: namely, almost all Steiner triple systems have $(1 - o(1))n/2$ disjoint perfect matchings. As mentioned in the introduction, a Steiner triple system that can be decomposed into perfect matchings is called a Kirkman triple system, and even proving the existence of Kirkman triple systems was an important breakthrough. For Latin squares, the property of being decomposable into transversals is equivalent to the important property of having an *orthogonal mate*, which has a long history dating back to Euler. More details can be found in [45].

- A (q, r, λ) -*design* ($q > r$) of order n is a q -uniform hypergraph on the vertex set $[n]$ such that every r -set of vertices is included in exactly λ hyperedges. A (q, r) -*Steiner system* is a $(q, r, 1)$ -design (so, a Steiner triple system is a $(3, 2, 1)$ -design or equivalently a $(3, 2)$ -Steiner system, and a d -regular graph is a $(2, 1, d)$ -design). The methods in Section 2 generalise to (q, r, λ) -designs with mainly notational changes. Note that a 3-uniform perfect matching is actually a $(3, 1)$ -Steiner system, so as a sweeping generalization of Theorem 1.1 we might ask for which $r' \leq r$ and $\lambda' \leq \lambda$ do (q, r, λ) -designs typically contain spanning (q, r', λ') -designs of the same order. We note that in the case of regular graphs a much stronger phenomenon occurs: there is a sense in which a random $(d_1 + d_2)$ -regular graph is “asymptotically the same” as a random d_1 -regular graph combined with a random d_2 -regular graph (see [22, Section 9.5]).
- Another interesting question about random Steiner triple systems is whether they contain Steiner triple subsystems on fewer vertices. McKay and Wanless [36] proved that almost all Latin squares have many small Latin subsquares (see also [32]), but it was conjectured by Quackenbush [39] that most Steiner triple systems do not have proper subsystems. It seems unlikely that the methods in this paper will be able to prove or disprove this conjecture without substantial new ideas. Actually, by consideration of the random 3-uniform hypergraph $\mathbb{G}(n, 1/n)$ we suspect the expected number of 7-vertex Steiner triple subsystems (Fano planes) in a random Steiner triple system is $\Theta(1)$, and that the distribution of this number is asymptotically Poisson.
- We could ask more generally about containment and enumeration of subgraphs. Which hypergraphs H appear a.a.s. in a random Steiner triple system? Can we show that for all such H the number of copies of H is concentrated? The methods in this paper can probably be used to prove a lower bound for the number of copies of H when every subgraph of H has at least 2 more vertices than hyperedges, but due to the “infamous upper tail” issue (see [23]), an upper bound for the number of copies of H is likely to be more difficult.
- One of the most fundamental properties of random graphs and hypergraphs is that they have low *discrepancy*, meaning that every sufficiently large subset of vertices has about the expected number of (hyper)edges. In Section 5.1.2 we effectively proved a very weak one-sided discrepancy bound, but it is not clear how to use our methods to reach anywhere near optimal discrepancy. See [34] for some theorems and conjectures about discrepancy of Latin squares, many of which have natural counterparts for Steiner triple systems.

References

- [1] R. Aharoni and E. Berger, *Rainbow matchings in r -partite r -graphs*, Electronic Journal of Combinatorics **16** (2009), no. 1, R119.
- [2] N. Alon, J.-H. Kim, and J. Spencer, *Nearly perfect matchings in regular simple hypergraphs*, Israel Journal of Mathematics **100** (1997), no. 1, 171–187.
- [3] N. Alon and J. H. Spencer, *The probabilistic method*, John Wiley & Sons, 2004.

- [4] L. Babai, *Almost all Steiner triple systems are asymmetric*, Annals of Discrete Mathematics **7** (1980), 37–39.
- [5] T. Bohman, A. Frieze, and E. Lubetzky, *A note on the random greedy triangle-packing algorithm*, Journal of Combinatorics **1** (2010), 477–488.
- [6] T. Bohman, A. Frieze, and E. Lubetzky, *Random triangle removal*, Advances in Mathematics **280** (2015), 379–438.
- [7] B. Bollobás, *The chromatic number of random graphs*, Combinatorica **8** (1988), no. 1, 49–55.
- [8] D. Bryant and D. Horsley, *Steiner triple systems without parallel classes*, SIAM Journal on Discrete Mathematics **29** (2015), no. 1, 693–696.
- [9] P. J. Cameron, *Almost all quasigroups have rank 2*, Discrete Mathematics **106** (1992), 111–115.
- [10] P. J. Cameron, *A generalization of t -designs*, Discrete Mathematics **309** (2009), no. 14, 4835–4842.
- [11] P. J. Cameron and I. M. Wanless, *Covering radius for sets of permutations*, Discrete mathematics **293** (2005), no. 1, 91–109.
- [12] N. J. Cavenagh, C. Greenhill, and I. M. Wanless, *The cycle structure of two rows in a random Latin square*, Random Structures & Algorithms **33** (2008), no. 3, 286–309.
- [13] C. J. Colbourn and A. Rosa, *Triple systems*, Oxford University Press, 1999.
- [14] T. M. Cover and J. A. Thomas, *Elements of information theory*, John Wiley & Sons, 2012.
- [15] P. Erdős and A. Rényi, *On random graphs I*, Publ. Math. Debrecen **6** (1959), 290–297.
- [16] P. Erdős and J. Spencer, *Lopsided Lovász local lemma and Latin transversals*, Discrete Applied Mathematics **30** (1991), no. 2, 151–154.
- [17] A. Ferber and M. Kwan, *Almost all Steiner triple systems are almost resolvable*, arXiv preprint arXiv:1907.06744 (2019).
- [18] D. A. Freedman, *On tail probabilities for martingales*, the Annals of Probability (1975), 100–118.
- [19] R. Glebov and Z. Luria, *On the maximum number of Latin transversals*, Journal of Combinatorial Theory, Series A **141** (2016), 136–146.
- [20] R. Häggkvist and J. C. Janssen, *All-even Latin squares*, Discrete Mathematics **157** (1996), no. 1, 199–206.
- [21] P. Hatami and P. W. Shor, *A lower bound for the length of a partial transversal in a Latin square*, Journal of Combinatorial Theory, Series A **115** (2008), no. 7, 1103–1113.
- [22] S. Janson, T. Łuczak, and A. Ruciński, *Random graphs*, Cambridge University Press, 2000.
- [23] S. Janson and A. Ruciński, *The infamous upper tail*, Random Structures & Algorithms **20** (2002), no. 3, 317–342.
- [24] A. Johansson, J. Kahn, and V. Vu, *Factors in random graphs*, Random Structures & Algorithms **33** (2008), no. 1, 1–28.
- [25] A. D. Keedwell and J. Dénes, *Latin squares and their applications*, second ed., Elsevier, 2015.
- [26] P. Keevash, *The existence of designs*, arXiv preprint arXiv:1401.3665 (2014).
- [27] P. Keevash, *Counting designs*, Journal of the European Mathematical Society **20** (2018), no. 4, 903–927.
- [28] P. Keevash, *The existence of designs II*, arXiv preprint arXiv:1802.05900 (2018).
- [29] P. Keevash and R. Mycroft, *A geometric theory for hypergraph matching*, Memoirs of the American Mathematical Society **233** (2015), no. 1098.

- [30] M. Krivelevich, *Triangle factors in random graphs*, *Combinatorics, Probability and Computing* **6** (1997), no. 03, 337–347.
- [31] M. Krivelevich, B. Sudakov, V. H. Vu, and N. C. Wormald, *Random regular graphs of high degree*, *Random Structures & Algorithms* **18** (2001), no. 4, 346–363.
- [32] M. Kwan and B. Sudakov, *Intercalates and discrepancy in random Latin squares*, *Random Structures & Algorithms* **52** (2018), no. 2, 181–196.
- [33] N. Linial and Z. Luria, *An upper bound on the number of Steiner triple systems*, *Random Structures & Algorithms* **43** (2013), no. 4, 399–406.
- [34] N. Linial and Z. Luria, *Discrepancy of high-dimensional permutations*, *Discrete Analysis* (2016), 1–8.
- [35] A. Lubotzky, Z. Luria, and R. Rosenthal, *Random Steiner systems and bounded degree coboundary expanders of every dimension*, *Discrete & Computational Geometry* (2018).
- [36] B. D. McKay and I. M. Wanless, *Most Latin squares have many subsquares*, *Journal of Combinatorial Theory, Series A* **86** (1999), no. 2, 323–347.
- [37] R. Montgomery, *Embedding bounded degree spanning trees in random graphs*, arXiv preprint arXiv:1405.6559 (2014).
- [38] A. Pokrovskiy, *An approximate version of a conjecture of Aharoni and Berger*, *Advances in Mathematics* **333** (2018), 1197–1241.
- [39] R. W. Quackenbush, *Algebraic speculations about Steiner systems*, *Annals of Discrete Mathematics* **7** (1980), 25–35.
- [40] J. Radhakrishnan, *An entropy proof of Brégman’s theorem*, *Journal of Combinatorial Theory, Series A* **77** (1997), no. 1, 161–164.
- [41] D. K. Ray-Chaudhuri and R. M. Wilson, *Solution of Kirkman’s schoolgirl problem*, *Proceedings of Symposia in Pure Mathematics*, vol. 19, 1971, pp. 187–203.
- [42] V. Rödl, *On a packing and covering problem*, *European Journal of Combinatorics* **6** (1985), no. 1, 69–78.
- [43] V. Rödl, A. Ruciński, and E. Szemerédi, *Perfect matchings in large uniform hypergraphs with large minimum collective degree*, *Journal of Combinatorial Theory, Series A* **116** (2009), no. 3, 613–636.
- [44] A. Taranenko, *Multidimensional permanents and an upper bound on the number of transversals in Latin squares*, *Journal of Combinatorial Designs* **23** (2015), no. 7, 305–320.
- [45] I. M. Wanless, *Transversals in Latin squares: a survey*, *Surveys in Combinatorics 2011* (R. Chapman, ed.), *London Mathematical Society Lecture Note Series*, vol. 392, Cambridge University Press, 2011, pp. 403–437.
- [46] L. Warnke, *The C_ℓ -free process*, *Random Structures & Algorithms* **44** (2014), no. 4, 490–526.
- [47] L. Warnke, *On the method of typical bounded differences*, *Combinatorics, Probability and Computing* **25** (2016), no. 2, 269–299.
- [48] R. M. Wilson, *Nonisomorphic Steiner triple systems*, *Mathematische Zeitschrift* **135** (1974), no. 4, 303–313.