Received: 23 August 2022 | Revised: 4 December 2024 | Accepted: 11 February 2025

DOI: 10.1112/jlms.70116

RESEARCH ARTICLE

Journal of the London Mathematical Society

Size-Ramsey numbers of graphs with maximum degree three

Nemanja Draganić[1] | **Kalina Petrova**[2]

[1]Mathematical Institute, University of Oxford, Oxford, UK

[2]Institute of Science and Technology Austria (ISTA), Klosterneuburg, Austria

Correspondence
Nemanja Draganić, Mathematical Institute, University of Oxford, Oxford, UK.
Email:
nemanja.draganic@maths.ox.ac.uk

Funding information
SNSF, Grant/Award Numbers: 217926, 200021_196965, CRSII5173721; European Commission, Grant/Award Number: 101034413; ETH Zürich

Abstract

The size-Ramsey number $\hat{r}(H)$ of a graph H is the smallest number of edges a (host) graph G can have, such that for any red/blue colouring of G, there is a monochromatic copy of H in G. Recently, Conlon, Nenadov and Trujić showed that if H is a graph on n vertices and maximum degree three, then $\hat{r}(H) = O(n^{8/5})$, improving upon the upper bound of $n^{5/3+o(1)}$ by Kohayakawa, Rödl, Schacht and Szemerédi. In this paper, we show that $\hat{r}(H) \leqslant n^{3/2+o(1)}$. While the previously used host graphs were vanilla binomial random graphs, we prove our result using a novel host graph construction. Our bound hits a natural barrier of the existing methods.

MSC 2020

05D10, 05D40, 05C80 (primary)

1 | INTRODUCTION

Almost a century ago, Ramsey [37] showed a result which gave rise to one of the most important notions in combinatorics. His theorem, which was followed by extensive research, states that for every two integers k and ℓ, there exists $r = r(k, \ell)$ referred to as their *Ramsey number*, which is the smallest integer such that in any colouring of the edges of K_r in red and blue, there is either a red K_k or a blue K_ℓ. Determining the value $r(k, \ell)$ for general k and ℓ has turned out to be difficult, and despite decades of research, there is still an exponential gap between the best-known lower and upper bounds [10, 18, 20, 41, 43, 44]. A natural generalisation of this concept is the Ramsey number $r(H_1, H_2)$ of two graphs H_1 and H_2, which is the minimum r such that any 2-colouring of the edges of K_r contains a red H_1 or a blue H_2; we also write $r(k)$ for $r(k,k)$ and $r(H)$ for $r(H,H)$.

For instance, a classic result of Gerencsér and Gyárfás [22] shows that the Ramsey number of a path P_n on n edges satisfies $r(P_n) = m$ for $m = \lceil \frac{3n+1}{2} \rceil$. That is, however we colour the edges of K_m, there is a monochromatic copy of P_n in it. Notice that the coloured graph K_m has quadratically many edges in n, while P_n only has linearly many edges. Is there a graph with much fewer edges than K_m such that any 2-colouring of its edges again gives a monochromatic copy of P_n? Already in 1983, answering a \$100 question of Erdős, Beck [4] showed that there is such a graph with only linearly many edges, which is evidently best possible.

To give a general framework for questions of this type, Erdős, Faudree, Rousseau and Schelp introduced the notion of *size-Ramsey numbers* [19]. Namely, given a graph H, the size-Ramsey number $\hat{r}(H)$ is the minimum number of edges a graph G can have, such that G is *Ramsey for* H, that is, any 2-colouring of G contains a monochromatic copy of H. We refer to G as the *host graph* for H.

The concept of size-Ramsey numbers allows us to study the minimality of the host graph more precisely. It is always possible to take a large enough complete graph as the host graph, hence $\binom{r(H)}{2}$ is a trivial upper bound for $\hat{r}(H)$. This is also tight when H is complete [19], but for other graphs H the optimal host graph is often much sparser. Indeed, for certain graph classes, one can even show that the size-Ramsey number is linear in the number of vertices of H, a significant improvement over the trivial upper bound, which is always at least quadratic. Namely, in addition to the aforementioned result by Beck [4] that $\hat{r}(P_n) = O(n)$, in their very elegant paper Friedman and Pippenger [21] proved that for every tree T of bounded degree on n vertices, $\hat{r}(T) = O(n)$. Furthermore, Haxell, Kohayakawa and Łuczak [25] established that for the cycle on n vertices, it holds that $\hat{r}(C_n) = O(n)$. Moreover, a linear upper bound of the size-Ramsey number was recently proved for long subdivisions of bounded degree graphs [17] and for bounded degree graphs with bounded treewidth [5, 26]. For further recent results on size-Ramsey numbers of (hyper)graphs, see [8, 9, 11, 13, 16, 23, 24, 32].

Considering the mentioned results, one may suspect that the size-Ramsey number of every bounded degree graph is linear in its number of vertices. In fact, Beck [4] asked this question before most of these 'positive examples' were discovered, but the answer was given much later and, perhaps surprisingly, was negative. Indeed, in 2000, Rödl and Szemerédi [40] showed that for every n, there are n-vertex graphs H of maximum degree 3 with $\hat{r}(H) \geqslant cn(\log n)^{\frac{1}{60}}$ for some constant c. In the same paper, it was conjectured that this can be improved to $n^{1+\varepsilon}$ for some constant $\varepsilon > 0$. Until very recently, $cn(\log n)^{\frac{1}{60}}$ was still the best known lower bound on the size-Ramsey number of bounded degree graphs. As the first version of this article, Tikhomirov [45] showed that there are n-vertex graphs H of maximum degree three with size-Ramsey number $\hat{r}(H) \geqslant cn\exp(c\sqrt{\log n})$ for a universal constant c. However, Rödl and Szemerédi's conjectured $n^{1+\varepsilon}$ still remains out of sight.

On the other hand, there have been some more recent important developments on the upper bound side. The baseline to be improved upon here is given by a classic result by Chvatál, Rödl, Szemerédi and Trotter [7], which shows that the Ramsey number of bounded degree graphs is linear in their number of vertices. This in turn gives a trivial quadratic upper bound for their size-Ramsey number. In 2011, Kohayakawa, Rödl, Schacht and Szemerédi [29] were able to show that every n-vertex graph H with maximum degree Δ satisfies $\hat{r}(H) \leqslant n^{2-1/\Delta+o(1)}$, thus bounding it away from quadratic. In the special case of H being triangle-free and $\Delta \geqslant 5$, this result was improved to $n^{2-\frac{1}{\Delta-1/2}+o(1)}$ by Nenadov [34].

Turning to particular instances of Δ, note that for graphs H of maximum degree 2, the size-Ramsey number is linear. Indeed, such graphs have bounded treewidth, and so by [26], we have that $\hat{r}(H) \leqslant O(n)$. Complementing the lower bound of Rödl and Szemerédi for size-Ramsey numbers of cubic graphs H, and giving the first improvement over the general upper bound by Kohayakawa, Rödl, Schacht and Szemerédi, recently Conlon, Nenadov and Trujić [14] showed that $\hat{r}(H) = O(n^{8/5})$ for all cubic graphs H. With the additional assumption that H is triangle-free, they further improved this bound to $\hat{r}(H) = O(n^{11/7})$, whereas when H is bipartite, they proved that $\hat{r}(H) = O(n^{14/9})$.

In this paper, we show that the size-Ramsey number of cubic graphs is at most $n^{3/2+o(1)}$. As we will discuss below, this bound hits a natural barrier of the existing methods underlying previous work in this direction.

Theorem 1.1. *The size-Ramsey number $\hat{r}(H)$ of every n-vertex graph H with maximum degree 3 satisfies $\hat{r}(H) \leqslant n^{3/2+o(1)}$.*

In general, when it comes to size-Ramsey numbers, a natural candidate for a host graph is the binomial random graph $G(N, p)$[†]. Indeed, most of the size-Ramsey number upper bounds so far are achieved with $G(N, p)$ as a host graph. Typically, one proves that a graph sampled from $G(N, p)$ is with high probability[‡] Ramsey for H, which gives an upper bound of $O(N^2 p)$ on the size-Ramsey number of H, as $G(N, p)$ with high probability has $O(N^2 p)$ edges.

However, there are some limits to what can be done with $G(N, p)$ as a host graph, coming from the fact that it is typically locally sparse. In particular, the upper bound of $O(n^{8/5})$ for cubic graphs H achieved in [14] is the best one can hope for using a vanilla binomial random graph as a host graph. For $p \ll N^{-2/5}$, the graph $G(N, p)$ is with high probability not even Ramsey for K_4 [38, 39], which can be a subgraph of H. We overcome this barrier by using a different host graph and additional new ideas for embedding H into a monochromatic subgraph of the host graph.

Our result pushes the known tools to their limit (up to the $o(1)$ term), most notably due to regularity inheritance, which is a property also used in the previous upper bounds on size-Ramsey numbers of bounded degree graphs. The employed approach which exploits this property breaks down when the number of edges in the host graph is asymptotically smaller than $n^{3/2}$, at least if they are 'uniformly' distributed. To move past this limitation, it seems that entirely new ideas are required.

We also note here that our proof yields a *universality-type* result, meaning that for every red/blue colouring of our host graph G, there is a monochromatic subgraph of G which contains all cubic graphs H on n vertices. In fact, as observed in [1], any graph which contains all cubic graphs must have at least $\Omega(n^{4/3})$ edges (even without the colouring condition). Following [29], we say that a graph G is *partition universal* for a class of graphs $\mathcal{F}$ if for every 2-colouring of the edges of G, there exists a monochromatic subgraph of G which contains a copy of every graph in $\mathcal{F}$. Hence, the proof of our main result shows that an optimal partition universal graph for all n-vertex graphs with maximum degree three has at most $n^{3/2+o(1)}$ edges, complementing the aforementioned lower bound of $\Omega(n^{4/3})$.

The rest of the paper is structured as follows. In Section 2, we give an overview of our approach. Section 3 provides some technical tools which we later use. In Section 4, we show a decomposition

[†] The binomial random graph $G(N, p)$ on N vertices is obtained by adding each potential edge independently at random with probability p.

[‡] A property is said to hold *with high probability* if it holds with probability tending to 1 as $N \to \infty$.

result for the cubic graph H, which prescribes the embedding process. Section 5 is where we give the construction of our host graph and prove a number of useful properties of it. These set the stage for the proof of Theorem 1.1 in Section 6. Finally, in Section 7, we make some concluding remarks.

2 PROOF OUTLINE

In this section, we present the main ideas of our approach. We first summarise the methods of Conlon, Nenadov and Trujić [14] that give an upper bound of $O(n^{8/5})$, which we use as a starting point.

Given a cubic graph H on cn vertices for some constant c, and an arbitrary 2-colouring of $G \sim G(n, p)$ with $p = \Omega(n^{-2/5})$, they use the regularity method to find 20 linear-sized sets of vertices in G, all pairs of which are regular in say blue, with some minimum density of order p. Next, to find a monochromatic copy of H in G, they decompose H into a number of vertex-disjoint parts $B_1, \dots, B_t$, and then embed the parts one by one in the blue subgraph induced by these 20 sets. Each part is either an induced path or an induced cycle of length at least 4, and the decomposition is *1-degenerate*, that is, each vertex $v \in B_i$ can have at most one neighbour u_v in $B_1 \cup \dots \cup B_{i-1}$. Hence, when they want to embed B_i, the 'candidate set' for v is the blue neighbourhood of u_v. The candidate sets can, with some care, be guaranteed to be of size $\Omega(np)$, and as $p = \Omega(n^{-1/2} \log n)$ one can ensure that each pair of candidate sets inherits regularity (as stated in Lemma 3.9, which we borrow from [14]).

There are two reasons why this approach reaches its limit when the host graph has $\Theta(n^{8/5})$ edges. The first one, as already mentioned in the introduction, is that for $p \ll n^{-2/5}$, the graph $G(n, p)$ is with high probability not Ramsey for K_4 [38, 39] and so is not a suitable host graph. The second bottleneck is in the embedding of the B_i's which are induced cycles. If $p = o(n^{-2/5})$, a copy of C_4 can no longer be embedded into the candidate sets as desired. Indeed, the technique used to embed the copy of C_4 relies on the KŁR conjecture (which is a theorem by now [3, 12, 35, 42]), and breaks down at the mentioned threshold for p.

More generally, any decomposition which contains short cycles as parts is an obstacle for constructing sparser host graphs. In particular, if the length of the shortest induced cycle in the decomposition is L, then with the technique at hand one needs at least $n^{\frac{3}{2}+\frac{1}{4L-6}}$ edges in the host graph. Thus, to overcome this barrier and make the host graph sparser, a new decomposition of H is needed, one that does not have short induced cycles as parts.

We deal with these two hurdles by using a different host graph and a different decomposition of H, as well as introducing some other techniques. The first ingredient we need is a random graph model which is locally dense but globally sparse, thereby being ideal for dealing with the aforementioned issue with copies of K_4. We make use of the fact that for every constant C, due to the existence of designs and subject to some divisibility conditions between n and C, the edges of K_n can be partitioned into copies of K_C [47]. Choosing a large enough constant C, we pick one such partition and then define the random graph model $G^C(n, p)$, in which the edges of each copy of K_C in the partition are taken to be present independently from all other copies with probability p (Definition 5.2). Provided that $C \geqslant r(4, 4)$, typically, every 2-colouring of $G^C(n, p)$ will have many monochromatic copies of K_4. This random graph model is an important building block of our host graph, but is not all there is to it, as our new decomposition requires a more involved host graph construction.

To address the second obstacle mentioned above, we need a new decomposition of H which avoids short induced cycles. To construct a host graph on $n^{3/2+o(1)}$ edges, one has to be able to avoid all cycles of length less than L, for any arbitrary constant L, as discussed above. Hence, in Section 4 we show Lemma 4.4, a decomposition result for cubic graphs similar in spirit to the one in [14], yet fundamentally different from it. One part in our decomposition is a graph with bounded treewidth, and the other parts are *long* induced cycles. The decomposition has the same 1-degeneracy condition as the one in [14]. The new part, the graph with bounded treewidth, is more complex, and thus requires different techniques to be embedded. In particular, it is known that bounded degree graphs with bounded treewidth have linear size-Ramsey numbers [5, 26], hence it is possible to construct a host graph on just a linear number of edges to embed the part of bounded treewidth from our decomposition.

It may seem like we are close to being done now, as we know how to handle each of the parts in the decomposition of H. But this is not really the case, as there are many difficulties to overcome. The main challenge lies in the fact that the host graphs we need for induced cycles on the one hand, and for the bounded treewidth part on the other, are very different, as are the embedding strategies used in these two cases. We cannot simply, for example, take a union of these two host graphs, as the adversary may colour one of them in blue and the other in red, and then we could not find all parts in the same colour. Instead, we need to carefully intertwine the two host graphs into a new host graph in such a way that if one part of the decomposition cannot be found in say blue, this guarantees that *all* parts can be found in red.

The host graph for bounded degree bounded treewidth graphs constructed by Kamčev, Liebenau, Wood and Yepremyan [26] is a blow-up of a third power of a random regular graph, in which each vertex of the random regular graph is replaced by a constant-sized clique, and each edge, by a complete bipartite graph. We make use of a slight modification of this host graph, where we use a blow-up of a binomial random graph $G(n', \frac{\log n'}{n'})$ instead, and take many copies of it (each with fresh randomness), superimposed in a particular way, as our host graph. Namely, we first sample a graph $G \sim G^C(n, p)$ with $p = n^{-1/2+\delta}$, and then partition almost all of its cliques of size C into a number of almost perfect packings (disjoint cliques which cover almost all vertices), see Lemma 5.13. On top of each such packing, we add a freshly generated copy of a blow-up of the third power of $G(n', \frac{\log n'}{n'})$, such that the blow-up of each vertex of $G(n', \frac{\log n'}{n'})$ is mapped to one of the cliques in the packing. The union of $G^C(n, p)$ with all these blow-ups of third powers of $G(n', \frac{\log n'}{n'})$ is our host graph Γ (see Definition 5.14).

We prove that if the decomposition part of bounded treewidth is not present in blue in the host graph Γ, it is present in red (Lemma 6.4) *and* we can find 21 linear-sized sets, where each two of them form a dense red regular pair (Proposition 6.7, Lemmas 5.9 and 5.22). We then use those sets to embed the long induced cycles in red similarly to before (Theorem 3.11). Let us note that finding the 21 sets with the required density is another important and technically involved part of the proof, relying on Turán's theorem, results from [26], and a number of careful counting arguments.

With all this at hand, we give a high-level overview of the proof of Theorem 1.1. We start by decomposing the graph H into an ordered collection of one bounded treewidth induced subgraph $\mathcal{T}$ and a number of induced cycles of length at least L for a big constant L (Lemma 4.4). Additionally, we make sure that each vertex has the property that it has at most degree 1 to the previous subgraphs in the collection.

We have different approaches for embedding $\mathcal{T}$, and for embedding the induced cycles. We embed the subgraphs from the decomposition one by one in order, starting from $\mathcal{T}$. We denote our host graph by Γ. Consider the largest subset U of the vertex set of Γ such that either the blue or the red subgraph of $\Gamma[U]$ does not contain $\mathcal{T}$. We distinguish two cases depending on whether the set U is larger than ιn for a small constant ι or not. The constant ι is chosen such that in the latter case $\mathcal{T}$ can be embedded in one of the sets given by an application of the sparse regularity lemma.

In the latter case, we apply the regularity lemma to the coloured Γ to obtain 21 sets which are pairwise (ε, p)-regular and where each pair has density at least γp in say red for appropriate constants ε and γ. After an appropriate cleaning process (Lemmas 3.8 and 3.9) which ensures that the red neighbourhood of each vertex behaves nicely, those sets will still be much larger than U. We then embed $\mathcal{T}$ into one of those 21 sets in red, which is particularly convenient because this way the neighbourhoods of the vertices of $\mathcal{T}$ into the other 20 sets (in red) are large. This enables us to successfully embed the remaining parts from the decomposition. Namely, what is left is to use the remaining 20 sets to embed the induced cycles. The candidate sets for each vertex in those induced cycles are of size at least of order np each. Now we use the technique developed in [14] to embed the cycles in the regular pairs (Theorem 3.11).

In the former case, we have a reasonably large subgraph $\Gamma' := \Gamma[U]$ with no copy of $\mathcal{T}$ in one of the two colours, say blue. We then use the result from [26] to show that, as there is no blue copy of $\mathcal{T}$ in Γ', then a red copy must exist in each large subgraph of Γ' (Lemma 6.4). Now we apply the sparse regularity lemma to the red subgraph of Γ', and using the fact that there are no blue copies of $\mathcal{T}$ in Γ', we conclude as discussed above that there exists a collection of 21 linear-sized sets, all pairs of which are regular and have enough density in red, which is the most technical part of our proof (Proposition 6.7, Lemmas 5.9 and 5.22). Finally, we embed the parts of H as before, by first embedding $\mathcal{T}$ in one of the linear-sized sets, and then embedding the induced cycles in the remaining 20 sets.

3 | PRELIMINARIES

In this section, we introduce our notation and state and prove some results used in our proof.

Notation. We use standard graph and set theoretic notation. For a graph $G = (V, E)$ and not necessarily disjoint vertex sets $A, B \subseteq V$, we denote by $e_G(A, B)$ the number of edges with one endpoint in A and another endpoint in B (edges with both endpoints in $A \cap B$ are counted only once). The neighbourhood $N_G(v, A)$ of a vertex v in A is the set of vertices adjacent to v in G, and the cardinality, $d_G(v, A)$, of that set is referred to as the degree of v in A. We also write $N_G(v)$ for $N_G(v, V(G))$ and $d_G(v)$ for $d_G(v, V(G))$. For an integer $k \geqslant 1$, we denote by G^k the graph obtained by adding an edge between every two vertices that are at distance at most k in the graph. An r-uniform hypergraph $\mathcal{H} = (V, E)$ consists of a vertex set V together with a collection E of r-subsets of V called hyperedges. For a subset $S \subseteq V$, we define the degree $d_{\mathcal{H}}(S)$ to be the number of hyperedges in $\mathcal{H}$ which contain S.

For simplicity, we employ the following conventions. We omit rounding of real numbers to nearest integers whenever it is not of vital importance. For two constants a, b, we use $a \ll b$ to indicate that b is large enough as a function of a so that our proofs go through. For example, we often use inequality chains like $a \gg b \gg c \gg d$, which also implies that in particular $a \gg bcd$. Furthermore, we write $\{a_1, \ldots, a_k\} \gg \{b_1, \ldots, b_m\}$ to abbreviate that $a_i \gg b_j$ for all i and j. For two functions $f, g : \mathbb{N} \to \mathbb{R}$, we write $f \approx g$ to express that $\lim_{n\to\infty} \frac{f(n)}{g(n)} = 1$. We denote by I_C the independent set on C vertices and by K_C the clique on C vertices.

3.1 | Concentration inequalities

We make use of the following standard concentration bounds for random variables.

Theorem 3.1 (McDiarmid's inequality [33]). *Consider the product* $(\Omega, \Pr)$ *of* N *probability spaces* $(\Omega_1, \Pr_1), \ldots, (\Omega_N, \Pr_N)$. *For a random variable* $X : \Omega \to \mathbb{R}$, *the* effect *of the ith coordinate is defined to be at most* c *if for every pair* $\omega, \omega' \in \Omega$ *that agree on all but the ith coordinate, it holds that* $|X(\omega) - X(\omega')| \leqslant c$. *Let* $X : \Omega \to \mathbb{R}$ *be a random variable such that for each* $i \in [N]$, *the effect of the ith coordinate on* X *is at most* c_i. *Then for all* $t \geqslant 0$, *we have*

$$\Pr\left[|X - \mathbb{E}[X]| \geqslant t\right] \leqslant e^{-\frac{2t^2}{\sum_{i=1}^{N} c_i^2}}.$$

The next theorem is a form of Chernoff's inequality.

Theorem 3.2 [2, Theorem A.1.19]. *For every* $C > 0$ *and* $\varepsilon > 0$, *there exists* $\delta > 0$ *so that the following holds: Let* X_i, $1 \leqslant i \leqslant n$ *for an arbitrary* n, *be independent random variables with* $\mathbb{E}[X_i] = 0$, $|X_i| \leqslant C$ *and* $\mathrm{Var}[X_i] = \sigma_i^2$. *Set* $X = \sum_{i=1}^{n} X_i$ *and* $\sigma^2 = \sum_{i=1}^{n} \sigma_i^2$ *so that* $\mathrm{Var}[X] = \sigma^2$. *Then for* $0 < a \leqslant \delta\sigma$, *it holds that*

$$\Pr[X > a\sigma] < e^{-\frac{a^2}{2}(1-\varepsilon)}.$$

We will also need the following slightly altered version of Chernoff's inequality.

Lemma 3.3 (Chernoff bound, weighted version). *Let* $C, \gamma > 0$ *and let* X_i *for* $1 \leqslant i \leqslant n$ *be independent random variables with* $X_i = C_i$ *with probability* p *and* $X_i = 0$ *otherwise, where* $0 < C_i \leqslant C$ *and* $0 < p \leqslant \frac{1}{2}$. *Then for* $X := \sum_{i=1}^{n} X_i$ *and* $\mathbb{E}[X] \to \infty$, *we have*

$$\Pr\left[|X - \mathbb{E}[X]| > \gamma\mathbb{E}[X]\right] \leqslant e^{-\Theta(\mathbb{E}[X])}.$$

Proof. Note that $\mathbb{E}[X] = \sum_{i=1}^{n} C_i p = \Theta(np)$. Let $Y_i := X_i - pC_i$ and $Z_i = -Y_i$ so that Y_i, Z_i fulfill the conditions of Theorem 3.2, setting $Y := \sum_{i=1}^{n} Y_i$ and $Z := \sum_{i=1}^{n} Z_i$. We have $\mathrm{Var}[Y_i] = \mathrm{Var}[Z_i] = \mathbb{E}[Y_i^2] = \mathbb{E}[Z_i^2] = C_i^2(1-p)p$. Thus, $\mathrm{Var}[Y] = \mathrm{Var}[Z] = \sum_{i=1}^{n} C_i^2(1-p)p = \Theta(np)$ and $\sigma = \Theta(\sqrt{np})$. We pick $a = \min\{\frac{\gamma\mathbb{E}[X]}{\sigma}, \sigma\delta_{3.2}\} = \Theta(\sqrt{np})$ where $\delta_{3.2}$ is the δ given by Theorem 3.2 applied with $C := \max_{i\in[n]}\{C_i\}$ and $\varepsilon := \frac{1}{2}$. Thus, by Theorem 3.2, we get

$$\Pr[X > \mathbb{E}[X](1+\gamma)] \leqslant \Pr[Y > a\sigma] \leqslant e^{-\Theta(np)} = e^{-\Theta(\mathbb{E}[X])}$$

$$\Pr[X < \mathbb{E}[X](1-\gamma)] \leqslant \Pr[Z > a\sigma] \leqslant e^{-\Theta(np)} = e^{-\Theta(\mathbb{E}[X])}.$$ □

3.2 | Regularity method

One of the main tools we use in our proof is a sparse version of Szemerédi's regularity lemma, and to state it we need the following two definitions.

Definition 3.4. For a graph G and disjoint subsets $V_1, V_2 \subseteq V(G)$, the pair (V_1, V_2) is said to be (ε, p)-*regular* for some $0 < \varepsilon, p \leqslant 1$ if, for every $U_1 \subseteq V_1, U_2 \subseteq V_2$ with $|U_1| \geqslant \varepsilon |V_1|, |U_2| \geqslant \varepsilon |V_2|$, it holds that

$$|d_G(U_1, U_2) - d_G(V_1, V_2)| \leqslant \varepsilon p,$$

where $d_G(A, B) = e_G(A, B)/(|A||B|)$ is the density of the pair (A, B).

Definition 3.5. A graph $G = (V, E)$ is said to be (γ, p)-upper-*uniform* if for all $U, W \subseteq V$ with $U \cap W = \emptyset$ and $|U|, |W| \geqslant \gamma|V|$, $e_G(U, W) \leqslant (1 + \gamma)p|U||W|$.

The next standard lemma follows directly from Definition 3.4 and shows that large enough subsets of regular pairs still constitute regular pairs.

Lemma 3.6. *Consider constants $0 < \varepsilon_1 < \varepsilon_2 \leqslant 1/2$ and $p \in (0, 1)$, and an (ε_1, p)-regular pair (X, Y). Every two subsets $X' \subseteq X$ and $Y' \subseteq Y$ of size $|X'| \geqslant \varepsilon_2|X|$ and $|Y'| \geqslant \varepsilon_2|Y|$ constitute an $(\varepsilon_1/\varepsilon_2, p)$-regular pair with $d(X', Y') \in d(X, Y) \pm \varepsilon_1 p$.*

In what follows, we state the sparse regularity lemma, which is an adaptation of Szemerédi's regularity lemma for sparse graphs. For a set S, we call $S_1, \ldots, S_t$ an *equipartition* of S if for all $i, j \in [t]$, we have $|S_i| = |S_j| \pm 1$. We call an equipartition $V_1, \ldots, V_t$ of the vertices of some graph G an (ε, p)-*regular equipartition* if all but at most $\varepsilon\binom{t}{2}$ pairs (V_i, V_j) are (ε, p)-regular.

Theorem 3.7 (Sparse regularity lemma [27, 28]). *For every $\varepsilon > 0$ and every integer $t_0 > 0$, there is $\gamma > 0$ and an integer $T \geqslant t_0$ such that every (γ, p)-upper-uniform graph G admits an (ε, p)-regular equipartition $V_1, \ldots, V_t$ of its vertices with $t_0 \leqslant t \leqslant T$.*

Note that the (sparse) regularity lemma was originally stated to give an exceptional set V_0 of size at most $\varepsilon|V(G)|$ as part of the equipartition. One function of this exceptional set is to be able to take all other sets to be of precisely the same size. As we do not need that and can afford to instead have differences of one in size between the parts, we can distribute the exceptional set between the other sets as equally as possible. The regularity property then still holds, but with a larger ε.

The next two lemmas help us 'clean up' a regular partition in such a way that all vertices have large neighbourhoods in each set of the regular partition and all pairs of neighbourhoods that belong to regular pairs are also regular. Note that in Lemma 3.9, we replaced $\tilde{n}d/4$ with $\tilde{n}d/20$, which is just a constant change that can be compensated for in the proof in [14] by taking β even smaller.

Lemma 3.8 [14, Lemma 3.3]. *For every $\Delta \in \mathbb{N}$ and $\gamma > 0$, there exists $\varepsilon_0 > 0$ such that the following holds for any $0 < \varepsilon \leqslant \varepsilon_0$ and $p \in (0, 1)$. Let H be a graph with maximum degree Δ and let $\{V_i\}_{i \in V(H)}$ be a family of subsets of some graph G such that (V_i, V_j) is (ε, p)-regular of density $d \geqslant \gamma p$ (with respect to G) for every $ij \in H$. Then, for every $i \in V(H)$, there exists $V_i' \subseteq V_i$ of order $|V_i'| \geqslant (1 - \Delta\varepsilon)|V_i|$ such that $d_G(v, V_j') \geqslant d|V_j|/2$ for every $v \in V_i'$ and all $ij \in H$.*

Lemma 3.9 [14, Lemma 3.5]. *For all $\varepsilon', \alpha, \gamma, \beta > 0$, there exist $\varepsilon_0 = \varepsilon_0(\varepsilon', \gamma, \beta)$ and $K = K(\varepsilon', \alpha, \gamma)$ such that, for every $0 < \varepsilon \leqslant \varepsilon_0$ and $p \geqslant K(\log n/n)^{1/2}$, the random graph $\Gamma \sim G(n, p)$ with high probability has the following property.*

Suppose $G \subseteq \Gamma$ and $V_1, V_2 \subseteq V(\Gamma)$ are disjoint subsets of order $\tilde{n} = \alpha n$ such that (V_1, V_2) is (ε, p)-regular of density $d \geqslant \gamma p$ with respect to G. Then there exists $B \subseteq V(\Gamma)$ of order $|B| \leqslant \beta\tilde{n}$ such that for each $v, w \in V(\Gamma) \setminus (V_1 \cup V_2 \cup B)$ (not necessarily distinct) the following holds: for any two subsets $N_v \subseteq N_\Gamma(v, V_1)$ and $N_w \subseteq N_\Gamma(w, V_2)$ of order $\tilde{n}d/20$, both (N_v, V_2) and (N_v, N_w) are (ε', p)-regular of density $(1 \pm \varepsilon')d$ with respect to G.

3.3 | Embedding cycles into a regular partition

The following lemma, which we borrow from [14], is used to embed long induced cycles into regular pairs.

Lemma 3.10 [14, Lemma 4.2]. *For every $\alpha, \gamma > 0$ and $\ell \geqslant 3$, there exist $c(\alpha, \gamma, \ell)$, $\varepsilon_0(\gamma, \ell) > 0$ and $K(\alpha, \ell, \gamma) > 0$ such that, for every $0 < \varepsilon \leqslant \varepsilon_0$ and $p \geqslant Kn^{-(\ell-2)/(2\ell-3)}$, the random graph $G \sim G(n, p)$ with high probability has the following property.*

Let C be a cycle of length $t \in [\ell, cn]$. Let $G' \subseteq G$ and, for each $v \in V(C)$, let $s_v \in V(G)$ be a uniquely chosen vertex. Then, for any collection of subsets $N_v \subseteq N_{G'}(s_v)$ of order αnp such that (N_v, N_w) is (ε, p)-regular of density $d \geqslant \gamma p$ (with respect to G') for each $vw \in C$, there exists a copy of C in G' which maps each $v \in V(C)$ to N_v.

In what follows, we make some minor modifications to the proof of the main theorem in [14], to adapt it to our application. The goal is, similarly to [14], to embed a '1-degenerate' collection of long induced cycles into a host graph with 20 vertex sets, all pairs among which are regular and dense enough. The main difference is that in our case some vertices already have predefined candidate sets before the embedding process begins. The proof remains almost unchanged, but we provide it here for completeness.

Theorem 3.11 (Proof of Theorem 1.2 in [14]). *For every $c, \delta, \gamma, \varepsilon, \mu > 0$ and $K(\mu, \ell, \gamma)$ such that $\varepsilon \ll \gamma < 1$, and $c \ll \{\mu, \gamma, \delta\} < 1$, and $\{c^{-1}, \varepsilon^{-1}\} \gg \ell \geqslant 3$ such that $\delta = \frac{1}{4\ell-6}$, the following holds with high probability for $G \sim G(n, p)$ with $p \geqslant Kn^{-\frac{1}{2}+\delta}$. Let $V_1, \dots, V_{20}$ be disjoint subsets of $V(G)$, each of size $|V_i| \geqslant \tilde{n} := \mu n$, and let G' be a subgraph of G.*

Let F be a cubic graph of cn vertices with vertex partition $F_1 \cup \dots \cup F_f$, such that each F_i is an induced cycle of length at least ℓ, and such that for each $v \in F_i$, either v has precisely one neighbour in $F_1 \cup \dots \cup F_{i-1}$, or v has no neighbours in $F_1 \cup \dots \cup F_{i-1}$, but there is a unique vertex $u_v \in V(G) - (V_1 \cup \dots \cup V_{20})$ and a 'candidate set' $C_v \subseteq N_{G'}(u_v)$ for v such that for all j we have $|C_v \cap V_j| \geqslant \frac{\tilde{n}d}{4}$, where $d := \gamma p$. In addition, suppose no u_v is chosen more than three times. Suppose also that

1. *$d_{G'}(v, V_i) \geqslant \frac{\tilde{n}d}{4}$ for each $v \in V_j$ and $i \neq j \in [20]$,*
2. *for all distinct i, j, h, g (but possibly $h = g$), for each $v \in V_h, w \in V_g$, and each $a, b \in V(F)$ with associated C_a, C_b, and any*
   - *$N_1 \subseteq N_{G'}(v, V_i)$ or $N_1 \subseteq C_a \cap V_i$ and*
   - *$N_2 \subseteq N_{G'}(w, V_j)$ or $N_2 \subseteq C_b \cap V_j$*

   of size $|N_1| = |N_2| = \frac{\tilde{n}d}{20}$, (N_1, N_2) and (N_1, V_j) are (ε, p)-regular of density at least $\frac{d}{2}$ in G'.

Then there exists a copy of F in $G'[V_1 \cup \dots \cup V_{20}]$, where each vertex is mapped to its candidate set C_v, if it has one assigned to it.

Proof. For each $i \leqslant 10$, denote by V_i^0 the set V_i, and denote by V_i^1 the set V_{i+10}. Before we start, assign to each vertex v in F a number $\varphi(v)$ from $\{1, \dots, 10\}$, such that each two vertices at distance at most 2 in F get a different number. This can be done easily by a greedy assignment. Now we embed the parts $F_1, \dots, F_f$ into G', one at a time, in the given order. Suppose we already embedded $F_1, \dots, F_{i-1}$ and let us show how to embed F_i.

(i) For every vertex $v \in F_i$, we define the vertex $u_v \in G'$ as follows. If v has a neighbour a_v in $F_1, \dots F_{i-1}$, we set u_v to be the image of that neighbour in G'. Otherwise, u_v is the vertex u_v from the statement of the theorem.
(ii) Let b_v be the smallest number in $\{0, 1\}$, such that the set $V_{\varphi(v)}^{b_v}$ has at least $\frac{\tilde{n}d}{20}$ vertices in $N_{G'}(u_v)$ not occupied by the embedding of $F_1, \dots, F_{i-1}$. If there is no such b_v, stop the procedure, and otherwise let S_v be the set of these at least $\frac{\tilde{n}d}{20}$ many non-occupied vertices.
(iii) Now, as $|S_v| \geqslant \tilde{n}d/20$ for every $v \in F_i$, we have that for those sets, using the second property from the theorem, the conditions of Lemma 3.10 are satisfied, so there exists a copy of the induced cycle F_i in G', which maps every vertex $v \in F_i$ to S_v (note that here we used that for every w adjacent to v in F_i, it holds that $\varphi(a_v) \neq \varphi(w) \neq \varphi(v)$).

If our procedure did not stop in Step (ii), we found the required copy of F in G', so it is enough to show that we did not stop early.

Suppose we are at the point of the algorithm where we want to embed F_i, and let us show that we can successfully do that. As F has cn vertices, the set X of all occupied vertices (at that moment) in $V_1^0 \cup \dots \cup V_{10}^0$ is clearly also of size at most cn. We denote by B_1 the set of vertices $v \in F_1 \cup \dots \cup F_{i-1}$ for which $b_v = 1$, and denote $U = \{u_v : v \in B_1\}$; note that $|B_1|$ is the number of vertices which are embedded into sets $V_1^1, \dots, V_{10}^1$. If we now show that $|B_1| = O(1/p)$, then we would be done, as every vertex u_v has at least $\tilde{n}d/4$ neighbours in each V_i^1, so our procedure would not stop early.

Suppose for contradiction that $|B_1| \geqslant 3C/p$ for a sufficiently large constant C, which implies also that $|U| \geqslant C/p$. Hence, as each vertex in U has at least $\tilde{n}d/8$ neighbours in X, and accounting for possible double-counting, we get that $e_{G'}(U, X) \geqslant \frac{|U|\tilde{n}d}{16}$. Meanwhile, using a Chernoff bound, a union bound and that C is sufficiently large, we conclude that with high probability it holds that for every set U' on at least C/p vertices and X' on cn vertices, there are at most $2|U'||X'|p$ edges in G between U' and X'. Therefore, by considering a superset X' of size cn of our set X, we get that $2|U|cnp \geqslant e_{G'}(U, X') \geqslant \frac{|U|\tilde{n}d}{16}$, which gives the required contradiction because $c \ll \mu\gamma$ (recalling that $\tilde{n} = \mu n$ and $d = \gamma p$). □

3.4 | Other auxiliary and classic results

In this subsection, we state some graph theoretic results which will come in handy in our proof.

The first lemma guarantees that almost perfect hypergraph matchings preserve some properties when a subset of their elements is considered.

Lemma 3.12. *Let M be a hypergraph matching on the vertex set $[n]$, where each edge is of size C, and M covers at least $(1-\gamma)n$ vertices. If $S \subseteq [n]$ is of size at least $|S| \geqslant 4\gamma n$, then there exist at least $|S|/2C$ edges $h \in M$ with $|h \cap S| \geqslant \gamma C/2$.*

Proof. Suppose for contradiction that for some set S, at most $|S|/2C$ edges satisfy $|h \cap S| \geqslant \gamma C/2$. Then the number of vertices in S is at most

$$\frac{|S|}{2C} \cdot C + \frac{n}{C} \cdot \frac{\gamma C}{2} + \gamma n = \frac{|S|}{2} + \frac{3n\gamma}{2} < |S|$$

where the first term bounds the number of vertices of S in edges h with $|h \cap S| \geqslant \gamma C/2$, the second one the vertices in edges which do not satisfy this condition, and the third one counts the vertices in S which are not in an edge in M. This gives the required contradiction. □

Next, we state the classic theorem of Turán.

Theorem 3.13 (Turán's theorem [46]). *Suppose G is a graph on n vertices with no K_{r+1} as a subgraph. Then $|E(G)| \leqslant (1 - \frac{1}{r})\frac{n^2}{2}$.*

We also make use of another well-known extremal result for bipartite graphs.

Theorem 3.14 (Kövári, Sós, Turán [31]). *The maximum number of edges in an n-vertex graph with no $K_{\ell,\ell}$ subgrpah is less than $(\ell - 1)^{\ell} n^{2-1/\ell} + \ell n + 1$.*

The next lemma concerns partitioning the edges of a hypergraph into matchings. The chromatic index $q(\mathcal{H})$ of a hypergraph $\mathcal{H}$ is the smallest integer q such that the set of edges $\mathcal{H}$ can be partitioned into q matchings.

Lemma 3.15 [36]. *For an integer $r \geqslant 2$ and $\gamma > 0$, there exists $\beta = \beta(r, \gamma) > 0$ so that the following holds. If an r-uniform hypergraph $\mathcal{H}$ has the following properties for some t.*

(1) $(1 - \beta)t < d(v) < (1 + \beta)t$ *holds for all vertices v.*
(2) $d(u, v) < \beta t$ *for all distinct pairs of vertices u, v.*

then $q(\mathcal{H}) \leqslant (1 + \gamma)t$.

4 | GRAPH DECOMPOSITION

In this section, we show that the vertices of every cubic graph G can be decomposed into sets which induce long cycles and other bounded treewidth graphs, in a way convenient for the embedding we use to prove Theorem 1.1.

We start by stating the following result from [30], which shows that every graph without long induced cycles and with bounded maximum degree has bounded treewidth. This result (albeit with weaker constants) had already been shown in [6].

Lemma 4.1 [30]. *Any graph G without induced cycles of length at least k and with maximum degree Δ has treewidth at most $(k - 1)(\Delta - 1) + 2$.*

The *strong product* $G \boxtimes H$ of graphs G and H is the graph with vertex set $V(G) \times V(H)$ in which (v_1, u_1) is adjacent to (v_2, u_2) if $v_1 = v_2$ and $\{u_1, u_2\} \in E(H)$, or $\{v_1, v_2\} \in E(G)$ and $u_1 = u_2$, or

$\{v_1, v_2\} \in E(G)$ and $\{u_1, u_2\} \in E(H)$. When H is a complete graph, we refer to $G \boxtimes H$ as a blow-up of G.

The following lemma states that every graph of bounded treewidth and bounded maximum degree is contained in a sufficiently large blow-up of a tree.

Lemma 4.2 [15, 48]. *Let G be a graph of treewidth w and maximum degree d. Then G is a subgraph of $T \boxtimes K_{18wd}$ for some tree T with maximum degree $18wd^2$.*

To state our decomposition result, it will be convenient for us to have the following definition.

Definition 4.3. Let G be a graph and let $\mathcal{S} = \{S_1, \dots S_t\}$ be a partition of its vertex set. Then we say that $\mathcal{S}$ is a *1-degenerate* partition of G if every vertex in S_i is adjacent to at most one vertex in $S_1 \cup \dots \cup S_{i-1}$ for all $2 \leqslant i \leqslant t$.

We are now ready to state our decomposition result for cubic graphs.

Lemma 4.4. *Let $\ell \geqslant 5$ and let G be a graph with $\Delta(G) \leqslant 3$. Then G admits a decomposition into subgraphs J and $F_1, \dots, F_g$ such that $V(J), V(F_1), \dots, V(F_g)$ is a 1-degenerate partition of $V(G)$ and the following hold.*

(A) *Each F_i is an induced cycle C_L with $L \geqslant \ell$.*
(B) *J has treewidth at most 2ℓ and is hence a subgraph of $T \boxtimes K_{400\ell}$ for some tree T of maximum degree 400ℓ.*

Proof. To obtain the graphs $F_1, \dots, F_g$, take out induced cycles of length at least ℓ from G one by one, each time removing all vertices from G which lie on the removed cycle, and taking a new induced cycle in the obtained graph. When there are no more long induced cycles to be removed, we are left with a subcubic graph J which contains no induced cycle of length at least ℓ, and hence by Lemma 4.1 has treewidth bounded by 2ℓ. By Lemma 4.2, J is also a subgraph of $T \boxtimes K_{400\ell}$ for some tree T of maximum degree 400ℓ. Notice that $V(J), V(F_1), \dots, V(F_g)$ is indeed 1-degenerate, as each vertex in some cycle F_i can have at most one edge going out of F_i, while J is the first part of the decomposition, so there are no additional conditions for it. $\square$

5 | THE HOST GRAPH

In this section, we describe our random graph model, and prove several results about its properties. Our model consists of a number of random cliques and random complete bipartite graphs, chosen in a particular way.

5.1 | The cliques

To describe our random graph model, we first recall the by now standard definition of a *Steiner system*. A Steiner system with parameters t, k, n, denoted by $S(t, k, n)$, is an n-element set S together with a collection of k-element subsets of S (called blocks) with the property that each t-element subset of S is contained in exactly one block. In particular, we will use Steiner systems with $t = 2$;

by a well-known result of Wilson [47] we have that for every large enough n, if $k(k-1)$ divides $n-1$, then a $S(2,k,n)$ exists. This immediately implies the following.

Corollary 5.1 [47]. *For any integer $C>0$, and large enough n, with $n-1$ divisible by $C(C-1)$, there exists a partition of the edges of K_n into cliques of size C.*

To define our random graph model, for each C and n as in the corollary above we fix one (arbitrary) Steiner system $S(2,C,n)$ and call it a *canonical* Steiner system. In the rest of the paper, we denote by $S(2,C,n)$ the set of cliques (which we also refer to as blocks) in K_n from the canonical $S(2,C,n)$. Sometimes we treat a block $B \in S(2,C,n)$ as a set of its vertices. We refer to the edges in a block B as $E(B)$. We are now ready to define our random graph model.

Definition 5.2. Given p, and C and n as above, let $G^C(n,p)$ be the random graph obtained by, independently for each block in a canonical $S(2,C,n)$, including all edges induced by that block with probability p.

Note that $G^2(n,p)$ has the same distribution as $G(n,p)$. Also note that by definition each edge in K_n is in precisely one block of $S(2,C,n)$. From now on we assume n is chosen large enough and so that $C(C-1)$ divides $n-1$.

The next result is needed in order to apply the sparse regularity lemma to $G^C(n,p)$. Recall that a graph $G=(V,E)$ is said to be (γ,p)-*upper-uniform* if for all $U,W\subseteq V$ with $U\cap W=\varnothing$ and $|U|,|W|\geqslant \gamma|V|$, $e_G(U,W)\leqslant(1+\gamma)p|U||W|$ (Definition 3.5).

Lemma 5.3. *For $0<\gamma,\delta<1/2$, there exists $K>0$ such that the following holds. The graph $G\sim G^C(n,p)$ with $p\geqslant K/n$ is with high probability (γ,p)-upper-uniform.*

Proof. Let U and W be disjoint subsets of $V(G)$ of size at least γn. Let $\mathcal{B}$ be the collection of blocks in $S(2,C,n)$ with at least one vertex in both U and W. For each block B in $\mathcal{B}$, let X_B be the random variable counting the edges in G between U and W contained in B. Note that $X_B = |U\cap B|\cdot|W\cap B|$ with probability p and $X_B=0$ otherwise. Therefore, by Lemma 3.3, the number of edges between U and W

$$X := \sum_{B\in\mathcal{B}} X_B$$

is at most $(1+\gamma)|U||W|p$ with probability at least $1-e^{-\Theta(|U||W|p)}$. Indeed, we have $\mathbb{E}[X]=|U||W|p$ and for each block B in $\mathcal{B}$, it holds that $X_B\leqslant\binom{C}{2}$, so the conditions of Theorem 3.3 are satisfied. By a union bound over all exponentially many choices of U and W, the probability that G is not (γ,p)-upper-uniform is at most

$$2^{2n}e^{-\Theta(|U||W|p)}\leqslant 2^{2n}e^{-\Theta(n^2p)}=o(1),$$

provided that K is large enough. □

By a similar argument (involving Chernoff bounds) as in the proof of Lemma 5.3, we also have the following.

Lemma 5.4. *For every $0 < \gamma, \delta < 1/2$, there is a $K > 0$ such that the graph $G \sim G^C(n,p)$ with $p \geqslant K \log n/n^2$ with high probability satisfies*

$$e(G) \leqslant (1+\gamma)\binom{n}{2}p \leqslant n^2 p.$$

In the proof of Theorem 1.1, it will be important for us to consider a random subsampling of $G^C(n,p)$, so that the obtained subgraph has the same distribution as $G(n,\tilde{p})$ for $\tilde{p} \approx \frac{p}{\binom{C}{2}}$.

Definition 5.5. In the remainder of the paper, we denote by G the random graph $G \sim G^C(n,p)$ for $p = n^{-1/2+\delta}$ where $\delta > 0$, and we define $\tilde{G}$ as follows. For each block $B \in S(2,C,n)$ with $E(B) \subseteq E(G)$, we sample a non-empty subset of the edges $B' \subseteq E(B)$ to be present in $\tilde{G}$ with the following probability

$$\Pr\left[E(B) \cap E(\tilde{G}) = B'\right] = \frac{\tilde{p}^{|B'|}(1-\tilde{p})^{\binom{C}{2}-|B'|}}{p}$$

where $\tilde{p}$ is given by $p = 1-(1-\tilde{p})^{\binom{C}{2}}$, that is, $\tilde{p} \approx \frac{n^{\delta-1/2}}{\binom{C}{2}}$.

The following simple lemma confirms that $\tilde{G}$ has the same distribution as $G(n,\tilde{p})$.

Lemma 5.6. *$\tilde{G}$ is distributed as $G(n,\tilde{p})$.*

Proof. Let $H \subseteq K_n$. We will show that $\Pr[\tilde{G} = H] = \Pr[G(n,\tilde{p}) = H]$. On the one hand,

$$\Pr[G(n,\tilde{p}) = H] = \tilde{p}^{e(H)}(1-\tilde{p})^{\binom{n}{2}-e(H)}.$$

For each $B \in S(2,C,n)$, denote by E_B the event that $\tilde{G}[B] = H[B]$. Notice that $\tilde{G} = H$ if E_B holds for each $B \in S(2,C,n)$, and that the events E_B are independent. Therefore,

$$\begin{aligned}
\Pr\left[\tilde{G} = H\right] &= \prod_{B\in S(2,C,n)} \Pr[E_B] \\
&= \prod_{\substack{B\in S(2,C,n)\\ e(H[B])=0}} (1-p) \prod_{\substack{B\in S(2,C,n)\\ e(H[B])>0}} \Pr[E_B | E(B) \subseteq E(G)]\Pr[E(B) \subseteq E(G)] \\
&= \prod_{\substack{B\in S(2,C,n)\\ e(H[B])=0}} (1-\tilde{p})^{\binom{C}{2}} \prod_{\substack{B\in S(2,C,n)\\ e(H[B])>0}} \frac{\tilde{p}^{e(H[B])}(1-\tilde{p})^{\binom{C}{2}-e(H[B])}}{p} p \\
&= \tilde{p}^{e(H)}(1-\tilde{p})^{\binom{n}{2}-e(H)},
\end{aligned}$$

where the last equality holds because $S(2,C,n)$ is a partition of the edges of K_n. □

We also need the following observation.

Lemma 5.7. *Let G be any outcome of $G^C(n,p)$. Then for any $B \in S(2,C,n)$ whose edges are in G, and any $e \in E(B)$, we have* $\Pr[e \in E(\tilde{G})] = \frac{\tilde{p}}{p} \approx \frac{1}{\binom{C}{2}}$.

Proof. We have

$$\Pr\left[e \in E(\tilde{G})\right] = \sum_{B' \subseteq E(B), e \in B'} \Pr\left[E(B) \cap E(\tilde{G}) = B'\right] \qquad = \sum_{s=1}^{\binom{C}{2}} \binom{\binom{C}{2}-1}{s-1} \frac{\tilde{p}^s (1-\tilde{p})^{\binom{C}{2}-s}}{p}$$

$$= \frac{\tilde{p}}{p} \sum_{s=0}^{\binom{C}{2}-1} \binom{\binom{C}{2}-1}{s} \tilde{p}^s (1-\tilde{p})^{\binom{C}{2}-1-s} \qquad = \frac{\tilde{p}}{p}. \qquad \square$$

The next lemma guarantees that for any equipartition of the vertices of K_n and any choice of a large enough subset $B_0 \subset B$ for each $B \in S(2,C,n)$, there are many sets B_0 which are roughly equally distributed between most of the parts.

Lemma 5.8. *Let t be an integer with $t \gg 1$ and let $0 < \beta < 1$. Then for $\rho \ll \beta$, and for $C \gg \{t, \beta^{-1}, \rho^{-1}\}$, the following holds.*

Consider $S(2,C,n)$ with n large enough. Suppose each $B \in S(2,C,n)$ has an associated $B_0 \subseteq B$ with $|B_0| = (1-\rho)C$. For any equipartition of $[n]$ into sets $V_1, \dots, V_t$, there is a collection $\mathcal{B}$ which contains a $(1-\beta)$-fraction of the blocks of $S(2,C,n)$, such that for each $B \in \mathcal{B}$, at least $(1-\beta)t$ sets V_i satisfy $|B_0 \cap V_i| \geqslant \frac{C}{200t}$.

Proof. We show the lemma by considering K_n, and counting *internal* edges in the B_0's, that is, edges with both endpoints in the same V_i, in two ways. First, the number of internal edges is at most $\binom{n/t}{2} t \leqslant \frac{n^2}{2t}$. Next, we count the internal edges in another way, by summing over two types of blocks, the *bad* blocks B, for which there are many sets V_i such that the intersection of B_0 and V_i is small, and the remaining *good* blocks.

Consider one block $B \in S(2,C,n)$. For each $i \in [t]$, let $a_i^B = |B_0 \cap V_i|$. Then the number of internal edges in B_0 is $\sum_{i=1}^t \binom{a_i^B}{2}$ and $C(1-\rho) = \sum_{i=1}^t a_i^B$. We say that a block B is *bad*, if there exists a set of βt sets V_i such that $a_i^B < \frac{C}{200t}$; notice that the remaining $(1-\beta)t$ sets must then contain a total of at least $(1-\rho-\beta/200)C$ vertices from B_0. The total number of internal edges in a bad block is minimised when each of those $(1-\beta)t$ sets has approximately the same number of vertices from B_0 (by the convexity of the binomial coefficient function $\binom{x}{2}$). The number of internal edges for a bad block is hence at least

$$(1-\beta)t \binom{\frac{(1-\rho-\beta/200)C}{(1-\beta)t}}{2} \geqslant (1+\beta/2)\frac{C^2}{2t}.$$

On the other hand, the number of internal edges for a good block is at least

$$t \binom{\frac{(1-\rho)C}{t}}{2} \geqslant (1-3\rho)\frac{C^2}{2t}$$

as again the number of internal edges is minimised if all vertices in a good block are distributed equally among the sets V_i. Suppose for contradiction that at least $\beta\frac{n(n-1)}{C(C-1)}$ of the blocks in $S(2,C,n)$ are bad. Then the total number of internal edges we get is at least

$$\beta\frac{n(n-1)}{C(C-1)}\cdot(1+\beta/2)\frac{C^2}{2t}+(1-\beta)\frac{n(n-1)}{C(C-1)}\cdot(1-3\rho)\frac{C^2}{2t}>\frac{n^2}{2t},$$

contradicting the upper bound on all internal edges. Therefore, at least $(1-\beta)\frac{n(n-1)}{C(C-1)}$ of the blocks in $S(2,C,n)$ are good. These form precisely the desired collection $\mathcal{B}$. □

The following lemma states that, given a set of vertices R in G, if a large enough fraction of the blocks present in G intersect R in a large set with no blue clique of a certain size, then R contains 21 linear-sized sets, all pairs of which are regular with large density in red. This serves an important purpose in the proof of Theorem 1.1, ensuring that the long induced cycles from the decomposition of H can be embedded in the same colour as the bounded treewidth part.

Lemma 5.9. *Let $C^{-1}\ll\rho\ll\alpha\leqslant\frac{1}{2}$ and $\{\tau,\varepsilon\}\ll\{C'^{-1},\alpha\}$, as well as $\{C^{-1},\mu\}\ll\varepsilon$. Let $R\subseteq V(G)$ with $|R|=\alpha n$. Then with high probability the following holds for $G\sim G^C(n,p)$.*

For every red/blue colouring of $E(G)$ such that at least an $\alpha/32$ fraction of the blocks B of $S(2,C,n)$ present in G contain a subset $B'\subseteq B\cap R$ with $|B'|=\rho C$ such that there is no blue $K_{C'}$ in $(R\cap B)-B'$, the following holds. There are disjoint subsets of vertices $V_1,\dots,V_{21}\subseteq R$, each of size $m\geqslant\mu n$, such that each pair (V_i,V_j) is (ε,p)-regular in the red subgraph of G and has density at least τp.

Proof. We start by giving an overview of the proof. We first apply the sparse regularity lemma to the red subgraph of $G[R]$, yielding sets $V_1,\dots,V_t$ such that most pairs of sets are regular. Next, with the help of Lemma 5.8, we get a collection $\mathcal{B}$ of many blocks of $S(2,C,n)$ present in G, each of which has a large blue $K_{C'}$-free intersection with most sets V_i. After that, Claim 1 shows via a double-counting argument that there is a collection $\mathcal{V}\subseteq\{V_1,\dots,V_t\}$ of r sets for a large constant r, all pairs of which are regular and such that there are many blocks in $\mathcal{B}$ that have a large blue $K_{C'}$-free intersection with *all* sets in $\mathcal{V}$. Now, the absence of a blue $K_{C'}$ and Turán's theorem imply that every b-tuple of sets in $\mathcal{V}$ contains at least one pair of sets with high red density in between (Claim 2). As r is chosen large enough, the lack of b many sets with pairwise low red density implies by Ramsey's theorem the existence of 21 sets with pairwise high red density, as required by the statement of the lemma. We continue with the detailed proof.

Let $b=2C'$, $\beta=\alpha/3200$ and $\xi=\frac{1}{200}$. Choose $r=r(21,b)$ to be the Ramsey number for 21 and b, and note that as r only depends on C', we may assume $r\ll\varepsilon^{-1}$. Let $t_0\gg\{r,b\}$. Let T be the maximum number of sets given to us by Theorem 3.7 (the sparse regularity lemma) with constants ε and t_0, and note that we can safely assume that $C\gg T$, as T is only a function of C',ε, and we have $C\gg\{C',\varepsilon^{-1}\}$ by assumption.

By a standard application of the Chernoff bound, with high probability G contains $(1\pm\beta)\frac{n(n-1)p}{C(C-1)}$ blocks of $S(2,C,n)$. Thus, by assumption, at least $\frac{\alpha}{32}(1-\beta)\frac{n(n-1)p}{C(C-1)}\geqslant\frac{\alpha}{64}\frac{n(n-1)p}{C(C-1)}$ blocks B present in G have a subset $B'\subseteq B$ with $|B'|=\rho C$ such that there is no blue $K_{C'}$ in $(R\cap B)-B'$.

Set G' to be the red subgraph of $G[R]$ and $n':=|V(G')|\geqslant\alpha n$. By Lemma 5.3, for any constant $\gamma>0$, the graph G, and therefore also G', is with high probability (γ,p)-upper-uniform (as G' has a linear in n number of vertices). We apply Theorem 3.7 to G', with ε and t_0. We get vertex sets $V_1,\dots,V_t$ with $t_0\leqslant t\leqslant T$ such that each V_i has size $m:=\frac{n'}{t}$ and such that all but $\varepsilon\binom{t}{2}$ many pairs

V_i, V_j are (ε, p)-regular. Partition $V(G) \setminus R$ arbitrarily into $t' \leqslant t(\frac{1}{\alpha} - 1)$ sets $V_{t+1}, \dots, V_{t+t'}$ each also of size m.

Next, we apply Lemma 5.8 to $V_1, \dots, V_{t+t'}$ with $t := t + t'$, β and ρ, where for each B, we either assign $B_0 := B - B'$, where B' is as given above (if it exists, which holds for at least $\frac{\alpha}{64} \frac{n(n-1)p}{C(C-1)}$ blocks B), or B_0 is an arbitrary subset $B_0 \subset B$ of size precisely $(1-\rho)C$. We thus get a collection $\mathcal{B} \subseteq S(2, C, n)$ of size at least $(1-\beta)\frac{n(n-1)}{C(C-1)}$ such that for each $B \in \mathcal{B}$, at least $(1-\beta)(t+t')$ sets V_i satisfy $|B_0 \cap V_i| \geqslant \xi \frac{C}{t+t'}$. At least $(1-2\beta)\frac{n(n-1)p}{C(C-1)}$ of the blocks in $\mathcal{B}$ are present in G by a standard application of the Chernoff bound and a union bound over all partitions $V_1, \dots, V_t$. In other words, among all of the blocks present in G, with high probability at most $3\beta \frac{n(n-1)p}{C(C-1)}$ are not in $\mathcal{B}$.

Hence, we have at least $(\frac{\alpha}{64} - 3\beta)\frac{n(n-1)p}{C(C-1)} \geqslant \beta \frac{n(n-1)p}{C(C-1)}$ blocks B from $\mathcal{B}$ in G that contain a subset B_0 of size $(1-\rho)C$ with no blue $K_{C'}$ in $R \cap B_0$. Call these blocks *blue-avoiding*. Note that for each blue-avoiding block B, the number of sets V_i among $V_1, \dots, V_t$ with $|B_0 \cap V_i| \geqslant \xi \frac{C}{t+t'}$ is at least

$$(1-\beta)(t+t') - t' = (1-\beta)t - \beta t' \geqslant (1-\beta)t - \beta t\left(\frac{1}{\alpha} - 1\right) \geqslant \frac{3}{4}t.$$

For a collection $\mathcal{V}' \subseteq \{V_1, \dots, V_t\}$, we say that a blue-avoiding block B *intersects* $\mathcal{V}'$ *nicely* if $|B_0 \cap V_i| \geqslant \xi \frac{C}{t+t'}$ for each $V_i \in \mathcal{V}'$. We next show a claim that provides us with a collection $\mathcal{V} \subseteq \{V_1, \dots, V_t\}$ such that many blue-avoiding blocks intersect $\mathcal{V}$ nicely.

Claim 1. There is a collection $\mathcal{V} \subseteq \{V_1, \dots, V_t\}$ with $|\mathcal{V}| = r$ such that each pair $V_i, V_j \in \mathcal{V}$ is a regular pair and at least $f\frac{n(n-1)p}{C(C-1)}$ blue-avoiding blocks B intersect $\mathcal{V}$ nicely, where $f = \frac{\beta}{10^r}$.

Proof. Suppose for contradiction there is no such collection $\mathcal{V}$. Note that the number of r-tuples with at least one irregular pair is at most $\varepsilon\binom{t}{2}\binom{t-2}{r-2}$. We count in two ways the number N of pairs $(\mathcal{V}, B)$, consisting of an r-tuple $\mathcal{V}$ with no irregular pairs, and a blue-avoiding block B that intersects $\mathcal{V}$ nicely. On the one hand, each of the blue-avoiding blocks intersects nicely at least $\binom{\frac{3t}{4}}{r}$ many r-tuples. On the other hand, we assumed each r-tuple with no irregular pairs is intersected nicely by less than $f\frac{n(n-1)p}{C(C-1)}$ blue-avoiding blocks. Therefore,

$$\beta\frac{n(n-1)p}{C(C-1)}\left[\binom{\frac{3t}{4}}{r} - \varepsilon\binom{t}{2}\binom{t-2}{r-2}\right] \leqslant N \leqslant \binom{t}{r} f \frac{n(n-1)p}{C(C-1)},$$

and now, using that $(a/b)^b \leqslant \binom{a}{b} \leqslant (ae/b)^b$ for all $a \geqslant b$, and the fact that $t \gg r$, we have

$$\beta\frac{t^r}{r^r}\left[(3/4)^r - \varepsilon e^r r^2\right] \leqslant \frac{e^r t^r}{r^r} f \quad \Rightarrow \quad \frac{\beta 3^r}{2 \cdot 4^r} \leqslant e^r f,$$

where the implication holds as $\varepsilon \ll r^{-1}$. This gives a contradiction by our choice of f, showing the existence of the desired collection $\mathcal{V}$. □

Without loss of generality, assume that the collection $\mathcal{V}$ given by Claim 1 is $\mathcal{V} = \{V_1, \dots, V_r\}$. Consider the complete graph $\mathcal{R}$ on the vertex set $V_1, \dots, V_r$ and call an edge V_i, V_j red if there

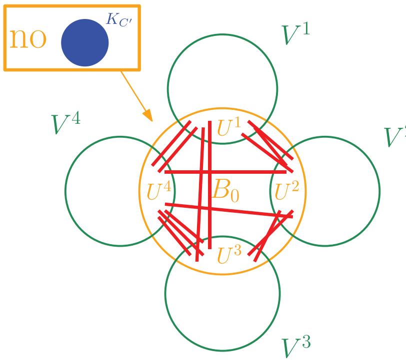

FIGURE 1 An illustration of the argument in Claim 2, showing that the absence of a blue clique implies high density in red. Here we have $b = 4$ and the sets $V^1, \dots, V^b$ are given in green. The intersections of these sets with B_0 (in orange) are $U^1, \dots, U^b$. Their union U has no blue $K_{C'}$, implying by Turán's theorem that there are many red edges in between $V^1, \dots, V^b$.

are at least $\frac{f\xi^2 n^2 p}{16b(t+t')^2}$ red edges between V_i and V_j in G, and blue otherwise. We next show the following claim, which will later allow us to conclude that there is a large red clique in $\mathcal{R}$.

Claim 2. There is no blue K_b in $\mathcal{R}$.

Proof. Consider some b-tuple of sets $V^1, \dots, V^b \in \mathcal{V}$ and some blue-avoiding block B that intersects $\mathcal{V}$ nicely. Recall that B_0 contains no $K_{C'}$ in blue in R and that $|B_0 \cap V^i| \geqslant \xi \frac{C}{t+t'}$ for each $i \in [b]$. Denote by U^i a subset of size precisely $\xi \frac{C}{t+t'}$ of $B_0 \cap V^i$, for each $i \in [b]$. Letting $U := \cup_{i=1}^b U^i$, note that U induces a clique in G with no blue $K_{C'}$ in it. Then by Turán's theorem, at least $\frac{1}{C'+1}\binom{|U|}{2}$ of the edges of U are red. Note that U has $\binom{\frac{b\xi C}{t+t'}}{2}$ edges, out of which $b\binom{\frac{\xi C}{t+t'}}{2}$ are internal, meaning that they are contained in some U^i. Thus, at least $\frac{1}{C'+1}\binom{\frac{b\xi C}{t+t'}}{2} - b\binom{\frac{\xi C}{t+t'}}{2} \geqslant \frac{1}{4b}\binom{\frac{b\xi C}{t+t'}}{2}$ of the edges of U are red and are not internal (see Figure 1 for an illustration of this argument).

Because there are at least $\frac{fn(n-1)p}{C(C-1)}$ blue-avoiding blocks B that intersect $\mathcal{V}$ nicely, it follows that there are at least $\frac{fn(n-1)p}{C(C-1)} \frac{1}{4b}\binom{\frac{b\xi C}{t+t'}}{2} \geqslant \frac{fb\xi^2 n^2 p}{16(t+t')^2}$ red edges that are not internal in $V^1 \cup \dots \cup V^b$. Thus, at least one pair among $\{V^i\}_{i\in[b]}$, say V^i and V^j, has at least $\frac{f\xi^2 n^2 p}{16b(t+t')^2}$ red edges between them. Therefore, the edge V^i, V^j is red in $\mathcal{R}$. Such an edge exists for any arbitrary b-tuple of sets in $\mathcal{V}$, from which the claim follows. $\square$

As $r = r(21, b)$ and by Claim 2 there is no blue K_b in $\mathcal{R}$ among $V_1, \dots, V_r$, there must be a red K_{21}, say given by $V_1, \dots, V_{21}$. Note that for each $i, j \in [21]$, $\{V_i, V_j\}$ is an (ε, p)-regular pair with at least $\frac{f\xi^2 n^2 p}{16b(t+t')^2}$ red edges. Thus, the (red) density of this (ε, p)-regular pair is at least $\frac{f\xi^2 p}{32C'}$. Hence, $V_1, \dots, V_{21}$ are the desired sets, concluding the proof of the lemma. $\square$

The next lemma allows us to conclude that if 21 linear-sized sets of vertices are pairwise regular and dense enough in some colour in G, as given by the conclusion of Lemma 5.9, then the same holds for these 21 sets in $\tilde{G}$.

Lemma 5.10. *Let $K, C \in \mathbb{N}$ with $C \geqslant 2$, $\mu > 0$ and let $0 < \gamma' \ll \varepsilon \ll \gamma$. Let G be an outcome of $G^C(n, p)$ which is (γ', p)-upper-uniform. Suppose there is a red/blue colouring of its edges, such that for disjoint sets of vertices $V_1, \dots, V_K$ with $|V_i| = \mu n$, the red subgraph of $G[V_i, V_j]$ is (ε, p)-regular with density at least γp, for each $i \neq j$. Then with high probability $\tilde{G}$ is such that the red subgraph of $\tilde{G}[V_i, V_j]$ is $(4\varepsilon, \tilde{p})$-regular with density at least $\frac{\gamma}{2}\tilde{p}$, for each $i \neq j$.*

Proof. Let U, W be distinct sets among $V_1, \dots, V_K$. We show the statement of the lemma for U and W, which then by a union bound holds for all such pairs. We refer to the red subgraphs of G and $\tilde{G}$ as G_r and $\tilde{G}_r$, respectively. First we show the following claim.

Claim 3. Let $\nu > 0$ and let $U' \subseteq U$ and $W' \subseteq W$ with $|U'| \geqslant \varepsilon|U|$ and $|W'| \geqslant \varepsilon|W|$. Then with high probability it holds that $d_{\tilde{G}_r}(U', W') = (1 \pm \nu)\frac{\tilde{p}}{p} d_{G_r}(U', W')$.

Proof. Let $\mathcal{B}$ be the collection of blocks from $S(2, C, n)$ which appear in G. For each block B in $\mathcal{B}$, let X_B be the random variable counting the red edges in $\tilde{G}$ between U' and W' contained in B. Note that $0 \leqslant X_B \leqslant \binom{C}{2}$. Denote by $X := \sum_{B \in \mathcal{B}} X_B$ the number of red edges between U' and W' in $\tilde{G}$. By linearity of expectation, and using that for each $B \in \mathcal{B}$,

$$\mathbb{E}[X_B] = \sum_{e \in E(B) \cap E_{G_r}(U', W')} \Pr\big[e \in \tilde{G}\big],$$

we have

$$\mathbb{E}[X] = \sum_{B \in \mathcal{B}} \mathbb{E}[X_B] = e_{G_r}(U', W') \Pr\big[e \in \tilde{G} \text{ for a fixed } e \in E(B)\big] = e_{G_r}(U', W')\frac{\tilde{p}}{p},$$

where the last equality follows by Lemma 5.7. Therefore, $\mathbb{E}[X] = \Omega(n^2\tilde{p})$, as $e_{G_r}(U', W') = \Omega(n^2 p)$ because $G_r[U, W]$ is (ε, p)-regular with density at least γp. Applying Theorem 3.1 with $t = \nu\mathbb{E}[X]$, we get that

$$\Pr[|X - \mathbb{E}[X]| \geqslant \nu\mathbb{E}[X]] \leqslant 2e^{-\frac{2t^2}{|\mathcal{B}|\binom{C}{2}^2}} = e^{-\Omega\left(\frac{n^4\tilde{p}^2}{n^2}\right)} = e^{-\omega(n)}.$$

The claim then follows from a union bound over all 2^{2n} possible choices of U', W'. □

By Claim 3, we have that $d_{\tilde{G}_r}(U, W) \geqslant \gamma\tilde{p}/2$. Let $U' \subseteq U$ and $W' \subseteq W$ be such that $|U'| \geqslant \varepsilon|U|$ and $|W'| \geqslant \varepsilon|W|$. By the triangle inequality and Claim 3 with $\nu := \varepsilon$, we have

$$\begin{aligned}
&\left|d_{\tilde{G}_r}(U', W') - d_{\tilde{G}_r}(U, W)\right| \\
&\leqslant \left|d_{\tilde{G}_r}(U', W') - \frac{\tilde{p} d_{G_r}(U', W')}{p}\right| + \left|\frac{\tilde{p} d_{G_r}(U', W')}{p} - \frac{\tilde{p} d_{G_r}(U, W)}{p}\right| + \left|\frac{\tilde{p} d_{G_r}(U, W)}{p} - d_{\tilde{G}_r}(U, W)\right| \\
&\leqslant \varepsilon\frac{\tilde{p} d_{G_r}(U', W')}{p} + \varepsilon\tilde{p} + \varepsilon\frac{\tilde{p} d_{G_r}(U, W)}{p} \leqslant 2\varepsilon(1 + \gamma')\tilde{p} + \varepsilon\tilde{p} \leqslant 4\varepsilon\tilde{p},
\end{aligned}$$

where we used that G, and therefore also G_r, is (γ', p)-upper-uniform (for $\gamma' \ll \varepsilon$) in the penultimate inequality, and that $G_r[U, W]$ is (ε, p)-regular for bounding the middle term in the second line. □

5.2 | The complete bipartite graphs

We begin the section by showing that with high probability the blocks in G can be partitioned into block matchings, defined below.

Definition 5.11 (A block matching in G). Let $\mathcal{S}$ be the set of blocks in G. A subset $M \subseteq \mathcal{S}$ is *a block matching* if all blocks in M are pairwise vertex-disjoint.

We will also want each of the matchings to cover almost all the vertices, that is, each matching will cover all but ηn vertices, where $\eta > 0$ is small enough and we specify it later. Essentially, in the proof we will take η to be the smallest of all constants we use.

Definition 5.12 (Collection of almost perfect block matchings). Let $\mathcal{S}$ be the set of blocks in G. A *collection of almost perfect block matchings* is a family $\mathcal{M}$ of pairwise disjoint block matchings $M_1, \dots, M_z$ with $\frac{(1-\eta)(n-1)p}{C-1} \leqslant z \leqslant \frac{(1+\eta)(n-1)p}{C-1}$, such that each M_i covers all but at most ηn vertices.

Note that as with high probability G contains at most $\frac{(1+\eta)n(n-1)p}{C(C-1)}$ many blocks, every collection of almost perfect block matchings is such that the number of blocks not appearing in any block matching in it is with high probability at most $\frac{3\eta n(n-1)p}{C(C-1)}$.

Lemma 5.13. *With high probability, there exists a collection of almost perfect block matchings $\mathcal{M}$ of the blocks in G.*

Proof. Consider the hypergraph H with $V(H) = V(G)$ where the set of hyperedges $\mathcal{S}$ is the set of all blocks in G. We will show the existence of z disjoint matchings in H that each cover all but at most ηn vertices, where z is as in Definition 5.12. We apply Lemma 3.15 to H with $r := C$, $\gamma := \eta^3$, $\beta \ll \gamma, C^{-1}$ and $t := \frac{(n-1)p}{C-1} > \sqrt{n}$. We can do this as for each vertex v in H, we have $\mathbb{E}[d_H(v)] = \frac{(n-1)p}{C-1}$ and by using a Chernoff bound and a union bound over all vertices, we have with high probability that $(1-\beta/2)t \leqslant d_H(v) \leqslant (1+\beta/2)t$, and for every pair of vertices $u \neq v$ in H, their codegree satisfies $d_H(u,v) \leqslant 1 < \beta t$, as every pair of vertices is in at most one block in $S(2, C, n)$. Therefore, there is a partition of the hyperedges of H into matchings $M'_1, \dots, M'_{z'}$ for some $z' \leqslant (1+\gamma)t = \frac{(1+\gamma)(n-1)p}{C-1}$.

Suppose for contradiction fewer than $(1-\eta)t$ of the matchings $M'_1, \dots, M'_{z'}$ cover each at least $(1-\eta)n$ of the vertices of H. As each vertex has degree in H at least $(1-\beta/2)t > (1-\gamma)t$, by double counting pairs (v, M'_i) such that M'_i covers v, we get

$$(1-\gamma)tn \leqslant (1-\eta)tn + (1+\gamma-(1-\eta))t(1-\eta)n,$$

which implies $2\gamma \geqslant \gamma\eta + \eta^2$, contradicting our choice of γ. To finish, we choose the hypergraph matchings that each cover at least $(1-\eta)n$ of the vertices to be the required sets $M_1, \dots, M_z$. □

As our final random graph construction builds on top of $G^C(n,p)$, we will now fix an outcome of $G^C(n,p)$ which has some useful properties which hold with high probability in that random graph model. From now on, for each C and all large enough n, we assume G is an outcome of $G^C(n,p)$

for which the conclusions of Lemmas 5.4, 5.9 and 5.13 hold, and additionally the following holds[†]: if we now take $\tilde{G}$ to be subsampled from G as in Definition 5.5, then with probability at least 0.9, the conclusions of both Lemma 3.9 and Theorem 3.11 hold for $\tilde{G}$, and $\tilde{G}$ is $(\zeta, \tilde{p})$-upper-uniform for all constant $\zeta > 0$. Note that such a G exists as the conclusions of Lemmas 5.4, 5.9 and 5.13 hold with high probability for $G^C(n, p)$, and the conclusions of Lemma 3.9 and Theorem 3.11, as well as $(\zeta, \tilde{p})$-upper-uniformity hold with high probability for $G(n, \tilde{p})$. Recall that by Lemma 5.6 the outcome $\tilde{G}$ of the two-step process of sampling $G \sim G^C(n, p)$ and then subsampling it to get $\tilde{G}$ is distributed as $G(n, \tilde{p})$. Fubini's theorem thus implies that with probability at least $1/2$, the outcome of the first process is a graph G for which the following is true. If we subsample from G to get $\tilde{G}$, then with probability at least 0.9 we have that $\tilde{G}$ is $(\zeta, \tilde{p})$-upper-uniform for every constant $\zeta > 0$ and is such that Lemma 3.9 and Theorem 3.11 hold.

Definition 5.14 (The host graph). Let $\mathcal{M} = (M_1, \ldots, M_z)$ be a collection of almost perfect block matchings for the graph G as in Definition 5.12. For each $i \in [z]$, let $n_i := |M_i|$ and $G_i \sim G(n_i, p')$ with probability[‡] $p' = \frac{\log n}{n}$, where each vertex in G_i corresponds to a block in M_i. Let $A_1, \ldots, A_z$ be the collection of random graphs defined as $A_i = G_i \boxtimes I_C$, that is, blow-ups of G_i by an independent set of size C. Similarly, let $A'_1, \ldots, A'_z$ be defined as $A'_i = G_i^3 \boxtimes I_C$. We identify the sets I_C in each A_i and A'_i with the corresponding blocks in M_i, thus defining each A_i and each A'_i on the vertex set of G (see Figure 2). Now we are ready to define our *host graph*, namely it is the union of graphs $\Gamma = G \cup A'_1 \cup \ldots \cup A'_z$.

Lemma 5.15. *The host graph Γ with high probability has at most $n^{\frac{3}{2}+2\delta}$ edges.*

Proof. We have that $E(\Gamma) = E(G) \cup E(A'_1) \cup \ldots \cup E(A'_z)$. By Lemma 5.4, we have $e(G) \leqslant n^{\frac{3}{2}+\delta}$.

For each $i \in [z]$, we have $n_i \leqslant \frac{n}{C}$ as M_i is a block matching. Thus, by the Chernoff bound and the union bound, with high probability $\Delta(G_i) \leqslant 2Cn_i \frac{\log n}{n} \leqslant 2 \log n$, so with high probability $\Delta(G_i^3) \leqslant 8 \log^3 n$. Therefore, $e(G_i^3) \leqslant 8n \log^3 n$ and $e(A'_i) \leqslant 8C^2 n \log^3 n$. As $z \leqslant \frac{2np}{C}$ by Lemma 5.13, we have $e(A'_1) + \cdots + e(A'_z) \leqslant n^2 p \log^4 n \leqslant n^{\frac{3}{2}+2\delta}/2$. □

In a manner similar to Definition 5.5, we can subsample from the graphs A_i to get subgraphs of binomial random graphs.

Definition 5.16. We define the collection of random graphs $\tilde{A}_1, \ldots, \tilde{A}_z$ in the following way. For each $i \in [z]$ and for each copy D of $K_{C,C}$ in A_i that corresponds to an edge in G_i, we sample a non-empty subset of the edges $D' \subseteq E(D)$ to be present in $\tilde{A}_i$ with the following probability

$$\Pr\left[E(D) \cap E(\tilde{A}_i) = D'\right] = \frac{\tilde{p}'^{|D'|}(1 - \tilde{p}')^{C^2 - |D'|}}{p'},$$

where $\tilde{p}'$ is given by $p' = 1 - (1 - \tilde{p}')^{C^2}$, that is, $\tilde{p}' \approx \frac{p'}{C^2} = \frac{\log n}{nC^2}$.

[†] Note that here we did not yet specify the explicit constants which we use in those lemmas, but observe that any choice of constants which satisfies the relations in the lemmas works when n is large enough. We specify the constants later in the proof, when we explicitly call the mentioned lemmas.

[‡] Here we could have chosen any function $\alpha(n)$ growing to infinity with n instead of $\log n$, and our arguments still would go through.

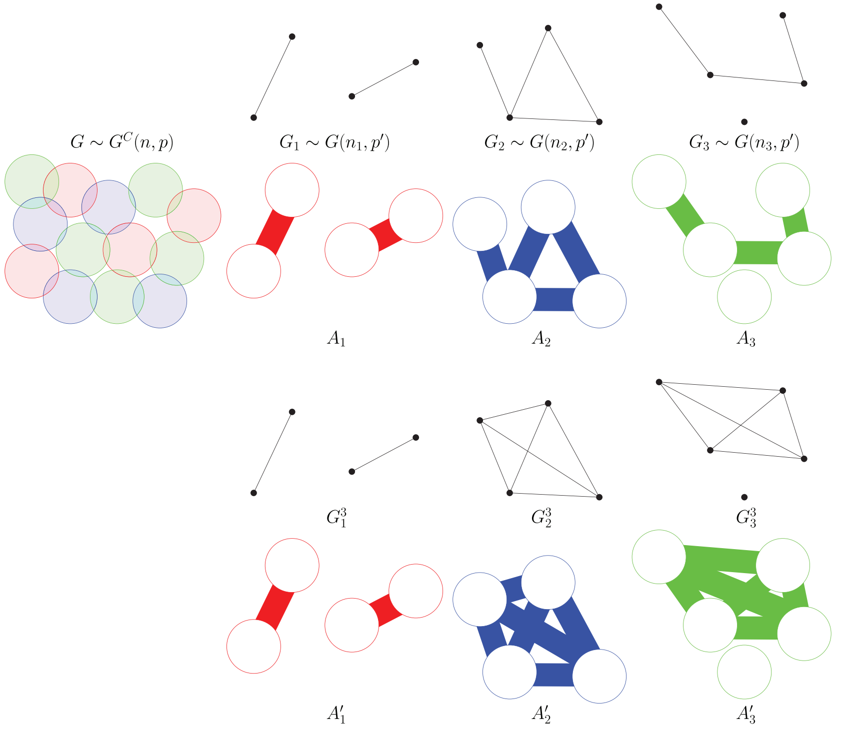

FIGURE 2 Obtaining the host graph: G contains disjoint almost perfect block matchings M_1, M_2, M_3, each of which covers almost all vertices. Each graph A_i is obtained from a binomial random graph G_i by replacing each vertex with a copy of I_C, and identifying these I_C's with the blocks in M_i. Each A_i' is obtained in the same way, but from G_i^3 instead of G_i.

In the remainder of the paper, we additionally use the probabilities $p'' = 1-(1-p')^z \approx \frac{p\log n}{C}$ and $\tilde{p}'' = 1-(1-\tilde{p}')^z \approx \frac{p\log n}{C^3}$, whose meaning will become apparent later in this section. For clarity of presentation, we provide a table with all edge probabilities we use, along with their definitions and asymptotic behaviour.

Lemma 5.17. *For $i \in [z]$, let $A_i = G_i \boxtimes I_C$ for any outcome of $G_i \sim G(n_i, p')$. Then for any copy D of $K_{C,C}$ whose edges are present in A_i, and any $e \in E(D)$, we have $\Pr[e \in E(\tilde{A}_i)] = \frac{\tilde{p}'}{p'} \approx \frac{1}{C^2}$.*

Proof. We have

$$\Pr\left[e \in E(\tilde{A}_i)\right] = \sum_{D' \subseteq E(D), e \in D'} \Pr\left[E(D) \cap E(\tilde{A}_i) = D'\right]$$

$$= \sum_{s=1}^{C^2} \binom{C^2-1}{s-1} \frac{\tilde{p}'^s (1-\tilde{p}')^{C^2-s}}{p'} = \frac{\tilde{p}'}{p'} \sum_{s=0}^{C^2-1} \binom{C^2-1}{s} \tilde{p}'^s (1-\tilde{p}')^{C^2-1-s} = \frac{\tilde{p}'}{p'}. \qquad \square$$

The next lemma considers the two-step process of first sampling $A_1, \ldots, A_z$, and then subsampling it to get $\tilde{A}_1, \ldots, \tilde{A}_z$, and shows that $\tilde{A}_1 \cup \cdots \cup \tilde{A}_z$ behaves as a subgraph of a binomial random graph.

Lemma 5.18. *The graph $\tilde{A}_1 \cup \cdots \cup \tilde{A}_z$ can be viewed as a subgraph of $G(n, \tilde{p}'')$, where $\tilde{p}'' = 1 - (1-\tilde{p}')^z \approx z\tilde{p}' \approx \frac{p \log n}{C^3}$.*

Proof. We couple the sampling process of $\mathcal{A} := \tilde{A}_1 \cup \cdots \cup \tilde{A}_z$ with that of $F \sim G(n, \tilde{p}'')$ in such a way that $\mathcal{A} \subseteq F$. We sample F using multiple exposure by first sampling $F_i \sim G(n, \tilde{p}')$ independently for each $i \in [z]$, and then taking $F := F_1 \cup \cdots \cup F_z$. We then define L_i in terms of F_i in the following way. For each $\{u, v\} \in F_i$, add $\{u, v\}$ to L_i if and only if there are two distinct blocks $B, B' \in M_i$ such that $u \in B$ and $v \in B'$. Take $L := L_1 \cup \cdots \cup L_z$.

As $L_i \subseteq F_i$ and therefore $L = L_1 \cup \cdots \cup L_z \subseteq F_1 \cup \cdots \cup F_z = F$, it is enough to show that L_i is indeed distributed as $\tilde{A}_i$. To do this, consider some $E' \subseteq E(K_n)$. First, if some $\{u, v\} \in E'$ is such that there are no $B, B' \in M_i$ with $u \in B$ and $v \in B'$, then $\Pr[E(L_i) = E'] = \Pr[E(\tilde{A}_i) = E'] = 0$. Next, as in both $\tilde{A}_i$ and L_i the edges between different pairs of blocks are sampled independently, it is sufficient to show that for each complete bipartite graph D between two blocks $B, B' \in M_i$ and for each $D' \subseteq E(D)$,

$$\Pr\big[E(\tilde{A}_i) \cap E(D) = D'\big] = \Pr\big[E(L_i) \cap E(D) = D'\big].$$

We have that the edges of $L_i \cap D$ behave precisely as in a random graph with edge probability $\tilde{p}'$, so

$$\Pr\big[E(L_i) \cap E(D) = D'\big] = \tilde{p}'^{|D'|}(1-\tilde{p}')^{C^2 - |D'|}.$$

On the other hand, for $\tilde{A}_i \cap D$ we analyse two cases depending on whether D' contains at least one edge or not. If $D' = \emptyset$, then, using that $1 - p' = (1-\tilde{p}')^{C^2}$,

$$\Pr\big[E(\tilde{A}_i) \cap E(D) = \emptyset\big] = \Pr[E(A_i) \cap E(D) = \emptyset] = (1-\tilde{p}')^{C^2} = \Pr[E(L_i) \cap E(D) = \emptyset].$$

If $D' \neq \emptyset$, then

$$\Pr\big[E(\tilde{A}_i) \cap E(D) = D'\big] = \Pr\big[E(\tilde{A}_i) \cap E(D) = D' | E(D) \subseteq E(A_i)\big] \cdot \Pr[E(D) \subseteq E(A_i)]$$

$$= \frac{\tilde{p}'^{|D'|}(1-\tilde{p}')^{C^2-|D'|}}{p'} \cdot p' = \tilde{p}'^{|D'|}(1-\tilde{p}')^{C^2-|D'|} = \Pr\big[E(L_i) \cap E(D) = D'\big]. \qquad \square$$

Lemma 5.19. *For every $0 < \gamma < 1$, with high probability the graph $A_1 \cup \cdots \cup A_z$ is (γ, p'')-upper-uniform, where $p'' = 1 - (1-p')^z \approx zp' \approx \frac{p \log n}{C}$.*

Proof. Consider some U and W with $|U|, |W| \geqslant \gamma n$ and $U \cap W = \emptyset$. Fix some $i \in [z]$. Let X_i be the random variable which counts the edges between U and W in A_i. For each copy D of $K_{C,C}$ whose two parts are two blocks from M_i and which has at least one edge between U and W, denote by X_D the number of edges from D contained in A_i with one endpoint in U and the other endpoint in W. Note that $X_i = \sum_D X_D$. Observe that the variables X_D are independent and take values between 0 and C^2. As we are showing an upper bound, we can without loss of generality assume that each pair of vertices in (U, W) is covered by some D, as otherwise we can add random variables X_D to the sum where each new D is the bipartite graph corresponding to just one uncovered pair of vertices. By Lemma 3.3 applied to X_i, we get

$$\Pr\left[X_i > \mathbb{E}[X_i] + \frac{\gamma}{2}|U||W|p'\right] \leqslant e^{-\Theta(\mathbb{E}[X_i])} \leqslant e^{-\Theta(n\log n)}.$$

It follows from a union bound over all 2^{2n} possible choices of U, W and $z = \Theta(np)$ choices of $i \in [z]$ that with high probability for all $i \in [z]$, the upper bound $X_i \leqslant \mathbb{E}[X_i] + \frac{\gamma}{2}|U||W|p'$ holds. Thus, as

$$\mathbb{E}[X_i] = \sum_{u\in U, w\in W} \Pr[uw \in E(A_i)] \leqslant |U||W|p',$$

we have that

$$e_{A_1\cup\dots\cup A_z}(U,W) \leqslant \sum_{i=1}^{z} X_i \leqslant \sum_{i=1}^{z}\left(\mathbb{E}[X_i] + \frac{\gamma}{2}|U||W|p'\right) \leqslant \left(1+\frac{\gamma}{2}\right)|U||W|p'z.$$

On the other hand, $p'' = 1-(1-p')^z \geqslant zp' - z^2(p')^2 = zp' - o(zp')$, so

$$e_{A_1\cup\dots\cup A_z}(U,W) \leqslant (1+\gamma)|U||W|p''. \qquad \square$$

The next technical lemma shows that the same edge is never in many A_i's.

Lemma 5.20. *With high probability, there are at most $n^{\frac{3}{2}}$ edges that occur in more than one A_i and there are no edges that occur in at least five of them.*

Proof. We have that for any two vertices u, w, $\Pr[uw \in E(A_i)] \leqslant p'$, so

$$\Pr[uw \text{ is in at least } k \text{ of the } A_i\text{'s}] \leqslant \binom{z}{k}(p')^k \leqslant (zp')^k = O\left(n^{(2\delta-\frac{1}{2})k}\right).$$

Let X_k be the number of edges that are in at least k of the A_i's. Then $\mathbb{E}[X_k] = O(n^{2-\frac{k}{2}+2\delta k})$. Setting $k = 5$, we get $\mathbb{E}[X_5] = O(n^{10\delta-0.5}) = o(1)$, so by the first moment method with high probability $X_5 = 0$.

For $k = 2$, we apply McDiarmid's inequality (Theorem 3.1) with $t := n^{3/2}/2$, where the coordinates that influence X_2 are all the $O(n^2 z) = O(n^{2.5+\delta})$ potential complete bipartite graphs in the A_i's (which are sampled independently). Note that each coordinate's effect is at most C^2. Then, as $\mathbb{E}[X_2] \leqslant n^{3/2}/2$

$$\Pr\left[X_2 > n^{\frac{3}{2}}\right] \leqslant \Pr[X_2 - \mathbb{E}[X_2] \geqslant t] \leqslant e^{-\frac{2t^2}{O(n^{2.5+\delta})}} = e^{-\Omega(n^{0.5-\delta})}. \qquad \square$$

We now introduce the concept of a *densifier* of some A_i, which is a structure that guarantees some density in one of the colours in A_i. The next lemma shows how to infer from the existence of sufficiently many densifiers in say red, that there are 21 linear-sized sets in Γ, all pairs of which are regular and dense in red. This is a key ingredient in the proof of Theorem 1.1 which shows

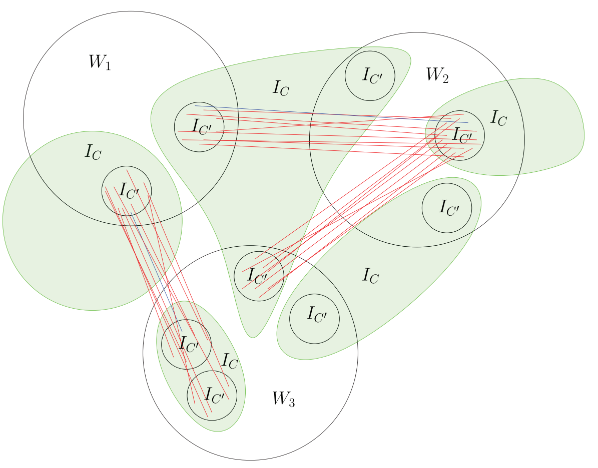

FIGURE 3 A red densifier: Each family W_k consists of a number of copies of $I_{C'}$, each of which is a subset of some copy of I_C in A_i. Whenever two copies of I_C are connected by a complete bipartite graph in A_i, each subgraph of that graph induced between two copies of $I_{C'}$ in different families W_k and $W_{k'}$ contains no blue copy of $K_{s,s}$ and so most of its edges are red.

that the induced cycles from the decomposition of H can be embedded in the same colour as the bounded treewidth part.

Definition 5.21 (A coloured (C', γ, s, q)-densifier). Let C', s, q be integers with $s \leqslant C' \leqslant C$, and let $\gamma > 0$. Let A_i be coloured in red and blue, and $S \subseteq V(A_i)$. Denote by $\mathcal{I} = \{I_1, I_2, \dots, I_{n_i}\}$ the copies of the independent set I_C on C vertices in A_i that correspond to the blocks in M_i. Then a *red (C', γ, s, q)-densifier* of $A_i[S]$ consists of q disjoint families $W_1, \dots, W_q$ each containing $\gamma \frac{n}{C'}$ independent sets, so that each $I \in W_k$ is such that $I \subset I_j \cap S$ for some $j \in [n_i]$ with $|I| = C'$, and all the I's are vertex disjoint across all the W_k's. Furthermore, for every pair $I \in W_k, I' \in W_{k'}$ with $k \neq k'$, there is no blue $K_{s,s}$ between I and I' in A_i (see Figure 3).

Lemma 5.22. *Let $q \geqslant 1000$ and $s \in \mathbb{N}$. For any $\mu \ll \varepsilon \ll \gamma \ll \{\alpha, q^{-1}\}$ and $C \gg C' \gg \{s, \mu^{-1}\}$, with high probability the following holds. Let $R \subset V(\Gamma)$ be a set of size $|R| = \alpha n$ and consider a colouring of Γ such that for at least half of the A_i's there exists a red (C', γ, s, q)-densifier of $A_i[R]$. Then there are disjoint subsets of vertices $V_1, \dots, V_{21} \subset R$, each of size $m \geqslant \mu n$, such that each pair (V_i, V_j) is (ε, p'')-regular in the red subgraph of $A_1 \cup \dots \cup A_z$ with density at least $\tau p''$ for $\tau = \tau(\gamma, \alpha, q)$.*

Proof. Let $A := A_1 \cup \dots \cup A_z$. We first give an overview of the proof. We start by applying the sparse regularity lemma (Theorem 3.7) to the red subgraph of $A[R]$. We get sets $V_1, \dots, V_t$, most pairs of which are regular. Next, we show by double-counting in Claim 4 that for each $i \in [z]$, and for each part W of the densifier of A_i (if such exists), at least some fraction of the sets V_k with $k \in [t]$ have a relatively large intersection with W. We then show with the help of Claim 4 there exist 21 sets $V^1, \dots, V^{21}$ among $V_1, \dots, V_t$ such that all pairs V^a, V^b are regular in red and at least a fraction of the A_i's respective densifiers each contain sets $W^1, \dots, W^{21}$ such that for each j, V^j and W^j have a large intersection (Claim 6). This kind of 'alignment' between the V^j's and the W^j's then allows us to conclude that all pairs V^a, V^b also have high density in red. The reason is that by Claim 5, many independent sets $I^a \in W^a$ and $I^b \in W^b$ each have a large intersection with V^a

and V^b, respectively. As there is no blue $K_{s,s}$ between I^a and I^b by the definition of a densifier, there must be many edges in red whenever a complete bipartite graph is present between I^a and I^b (Theorem 3.14). A concentration inequality lets us conclude that many such complete bipartite graphs exist in A_i, giving us the required density in red.

Let $\tau := \left(\frac{q\gamma^2}{800\alpha^2 e}\right)^{21}\frac{\gamma^2}{10^5\alpha^2}$. Let $t_0 \gg 1$ and let $T = T_{3.7}(t_0, \varepsilon)$ be the upper bound on the number of sets given by the sparse regularity lemma.

By Lemma 5.19, for all $\gamma' > 0$ the graph A is with high probability (γ', p'')-upper-uniform, which implies that the red subgraph of $A[R]$ is also (γ', p'')-upper-uniform for every fixed $\gamma' > 0$. We apply Theorem 3.7 to the red subgraph of $A[R]$ with ε and t_0 and get an equipartition $V_1, \dots, V_t$ of R with $t_0 \leqslant t \leqslant T$ such that all but $\varepsilon\binom{t}{2}$ pairs V_i, V_j are (ε, p'')-regular in red, and such that each V_i has size $|V_i| = m$.

Without loss of generality, assume that $A_1, \dots, A_{z/2}$ have an associated red densifier. Let $\mathcal{W}_i = \{W_1, \dots, W_q\}$ denote the (C', γ, s, q)-densifier of A_i with $i \in [z/2]$ and let $V(W_j)$ be the vertices in independent sets in W_j. We now show several simple counting claims.

Claim 4. Let $\beta = \frac{\gamma}{3\alpha}$, suppose $i \in [z/2]$, and let $W_j \in \mathcal{W}_i$. Then at least βt of the V_k's have the property that $|V_k \cap V(W_j)| \geqslant \beta |V_k|$.

Proof. Suppose this is not the case. Note that $|W_j| = \frac{\gamma n}{C'}$, so the total number of vertices contained in sets from W_j is equal to $v(W_j) := |V(W_j)| = \gamma n$. Then

$$\gamma n \leqslant \beta t m + (1-\beta) t \beta m,$$

where the right-hand side is an upper bound on the number of vertices in W_j, as the first term bounds the vertices in sets V_k with intersection with W_j at least βm, and the second term, all the others. This implies

$$\gamma \leqslant \beta\alpha + (1-\beta)\beta\alpha \leqslant 2\beta\alpha,$$

where we used $mt = \alpha n$, thus contradicting our choice of β. □

Claim 5. Suppose $i \in [z/2]$, let $W_j \in \mathcal{W}_i$ and let V_k be such that $|V(W_j) \cap V_k| \geqslant \beta |V_k| = \beta m$. Then at least $\beta m/2$ of the vertices of V_k each belong to some $I \in W_j$ such that $|I \cap V_k| \geqslant \frac{C'}{12t}$.

Proof. Suppose for contradiction that this is not the case. Then at least $\frac{\beta m}{2}$ of the vertices of V_k belong to some $I \in W_j$ such that $|I \cap V_k| < \frac{C'}{12t}$. To cover those $\beta m/2$ vertices of V_k, the number of such $I \in W_j$ that are required is at least

$$\frac{\beta m/2}{C'/(12t)} = \frac{6\beta\alpha n}{C'} = \frac{2\gamma n}{C'} > |W_j|,$$

contradicting the number of available $I \in W_j$. □

For any 21 sets $V^1, \dots, V^{21}$ from the regularity equipartition, and $i \in [z/2]$, we say that the pair $\left(\{V^1, \dots, V^{21}\}, A_i\right)$ is *nice* if all pairs V^j, V^k are regular in the red subgraph A^r of A and there are distinct sets $W^1, \ldots, W^{21} \in \mathcal{W}_i$, such that for every $j \in [21]$ we have $|V^j \cap V(W^j)| \geqslant \beta m$. Note that the number of *irregular* 21-tuples $V^1, \ldots, V^{21}$ (i.e. the tuples for which there is at least one pair V^j, V^k that is not regular in A^r) is at most εt^{21}.

Claim 6. For $\lambda := \left(\frac{q\gamma^2}{800\alpha^2 e}\right)^{21}$, there is a 21-tuple $\{V^1, \ldots, V^{21}\}$ such that for at least $\lambda z/2$ of the A_i's, $(\{V^1, \ldots, V^{21}\}, A_i)$ is nice.

Proof. Suppose for contradiction there is no such 21-tuple. As there are $\binom{t}{21}$ tuples in total, the number of nice pairs is then at most $\binom{t}{21}\lambda z/2$. Let us now show a lower bound on the number of nice pairs $\left(\{V^1, \ldots, V^{21}\}, A_i\right)$.

First, note that in each of the $z/2$ considered A_i's, there are $\binom{q}{21}$ tuples $W^1, \ldots, W^{21} \in \mathcal{W}_i$. Fix such a tuple $W^1, \ldots, W^{21}$. By Claim 4, for each W^j there are at least βt of the V_k's with $|V_k \cap V(W_j)| \geqslant \beta m$. Thus, there are at least $\binom{\beta t}{21}$ tuples $V^1, \ldots, V^{21}$, such that $|V(W^j) \cap V^j| \geqslant \beta m$ for each $j \in [21]$. Among these, at most εt^{21} tuples are irregular in A^r. This gives rise to at least $\frac{z}{2}\binom{q}{21}\left(\binom{\beta t}{21} - \varepsilon t^{21}\right)$ nice pairs, but note that for each $i \in [z/2]$, we have potentially counted each tuple $V^1, \ldots, V^{21}$ multiple times. Namely, each tuple $V^1, \ldots, V^{21}$ which forms a nice pair with A_i is counted at most $1/\beta^{21}$ many times because each V^k can have an intersection of size at least βm with at most $1/\beta$ many W_j's in $\mathcal{W}_i$. Thus, there are at least $\frac{z}{2}\binom{q}{21}\left(\binom{\beta t}{21} - \varepsilon t^{21}\right)\beta^{21}$ many nice pairs. Comparing this to the upper bound from above, we get

$$\frac{z}{2}\binom{q}{21}\left(\binom{\beta t}{21} - \varepsilon t^{21}\right)\beta^{21} \leqslant \binom{t}{21}\lambda z/2.$$

Thus, as $\beta \gg \varepsilon$, this implies $\frac{q^{21}}{2\cdot 21^{42}} t^{21}\beta^{42} \leqslant \frac{\lambda e^{21} t^{21}}{21^{21}}$, which boils down to $\left(\frac{q\gamma^2}{9\alpha^2 21 e}\right)^{21} \leqslant 2\lambda$, contradicting our choice of λ. □

Pick $\{V^1, \ldots, V^{21}\}$ such that for at least $\lambda z/2$ of the A_i's, the pair $(\{V^1, \ldots, V^{21}\}, A_i)$ is nice (assume without loss of generality these are $A_1, \ldots, A_{\lambda z/2}$). We now finish the proof of the lemma by showing that these $V^1, \ldots V^{21}$ are as desired.

Consider some V^a, V^b with $a, b \in [21]$. Note that V^a, V^b is an (ε, p'')-regular pair in A^r. We show a lower bound for $e_{A^r}(V^a, V^b)$. Consider some A_i such that $(\{V^1, \ldots, V^{21}\}, A_i)$ is nice. There must be some $W^a, W^b \in \mathcal{W}_i$ such that for each $x \in \{a, b\}$, $|V^x \cap V(W^x)| \geqslant \beta m$. By Claim 5, this implies that at least $\frac{\beta m}{2}$ of the vertices in V^x are in some $I \in W^x$ such that $|I \cap V^x| \geqslant \frac{C'}{12t}$. Call these I's *good* for V^x and recall that each I is a subset of an independent set $I_j \in \{I_1, \ldots, I_{n_i}\}$ on C vertices in A_i.

For each I^a which is good for V^a and I^b which is good for V^b with $I^a \subset I_j, I^b \subset I_h$, if $j \neq h$, the probability that in A_i there is a complete bipartite graph between I_j and I_h is precisely p'. Consider some I^a, I^b, which are good for V^a, V^b, respectively, with a complete bipartite graph between their respective supersets I_j and I_h in A_i. We know that there is no blue $K_{s,s}$ between $I^a \cap V^a$ and $I^b \cap V^b$, so by Theorem 3.14, as $\frac{C'}{12t} \gg s$, at least half of the edges between I^a and I^b are red.

Let X be the random variable counting the edges between all pairs I^a, I^b with $I^a \in W^a$ and $I^b \in W^b$ which are good for V^a and V^b, respectively and have distinct supersets $I_j \supset I^a$ and $I_h \supset I^b$ in

A_i, then

$$\mathbb{E}[X] = \sum_{I^a, I^b} |I^a \cap V^a||I^b \cap V^b| p' \geqslant \frac{1}{2}\left(\frac{\beta m}{2}\right)^2 p' = \Theta(n \log n),$$

where the inequality comes from the fact that at least $\beta m/2$ vertices in each V^a and V^b are in good I^a and I^b, respectively; furthermore, for each good I^a, there is at most a constant number of vertices in V^b that belong to the same superset $I_j \supset I^a$, and this is accounted for by the factor of $\frac{1}{2}$, which gives a generous lower bound. Note that at least $X/2$ of these X edges are red. Letting $X_{j,h}$ denote the number of edges between all good pairs $I^a \in W^a, I^b \in W^b$ with $I^a \subset I_j$ and $I^b \subset I_h$, note that $X = \sum_{j \neq h} X_{j,h}$, and we can think of the sum as only going over the pairs j, h such that there exists at least one good pair I^a, I^b with $I^a \subset I_j$ and $I^b \subset I_h$. As with probability p',

$$X_{j,h} = \sum_{\text{good } I^a \subset I_j, I^b \subset I_h} |I^a \cap V^a||I^b \cap V^b|,$$

in which case $1 \leqslant X_{j,h} \leqslant |I_j \cap V^a||I_h \cap V^b| \leqslant C^2$, and $X_{j,h} = 0$ otherwise, we can apply Lemma 3.3. We get

$$\Pr[X < \mathbb{E}[X]/2] \leqslant e^{-\Theta(\mathbb{E}[X])} = e^{-\Theta(n \log n)}.$$

By a union bound over all at most 2^n relevant subsets R of $V(\Gamma)$, all 2^{2n} subsets V^a and V^b and all at most 2^C subsets I^a of each $I_j \cap V^a$ and subsets I^b of each $I_h \cap V^b$, we have that with probability $1 - e^{-\Theta(n \log n)}$ for a fixed $i \in [\lambda z/2]$ we have $X \geqslant \frac{\beta^2 m^2 p'}{16}$, so the number of red edges between V^a and V^b in A_i is at least $\frac{\beta^2 m^2 p'}{32}$. Thus, $d_{A_i^r}(V^a, V^b) \geqslant \frac{\beta^2 p'}{32}$, where A_i^r denotes the red subgraph of A_i. Moreover, as the probability of failure is sufficiently small, by a union bound we get that $d_{A_i^r}(V^a, V^b) \geqslant \frac{\beta^2 p'}{32}$ for each $i \in [\lambda z/2]$.

As each edge is in at most 5 A_i's by Lemma 5.20, we have

$$d_{A^r}(V^a, V^b) \geqslant \frac{\lambda z}{2} \frac{\beta^2 p'}{5 \cdot 32} \geqslant \frac{\lambda \gamma^2 p''}{10^5 \alpha^2} = \left(\frac{q\gamma^2}{800\alpha^2 e}\right)^{21} \frac{\gamma^2 p''}{10^5 \alpha^2},$$

where we used Claim 6. □

The following lemma shows that under certain conditions on $A := A_1 \cup \cdots \cup A_z$, regular pairs in A remain regular after subsampling to get $\tilde{A}_1 \cup \cdots \cup \tilde{A}_z$. All of those conditions hold with high probability for an outcome of $A_1 \cup \cdots \cup A_z$, which we make use of in the choice of our host graph in the proof of Theorem 1.1.

Lemma 5.23. *Let $0 < \gamma' \ll \varepsilon \ll \gamma$, and $\mu \gg \gamma'$, and let $K \in \mathbb{N}$. Let A be an outcome of the random graph distribution $A_1 \cup \ldots \cup A_z$ which is (γ', p'')-upper-uniform. Suppose there is a red/blue colouring of A and K disjoint sets of vertices $V_1, \ldots, V_K$ with $|V_i| = \mu n$, such that the red subgraph of $A[V_i, V_j]$ is (ε, p'')-regular with density at least $\gamma p''$, for all $i \neq j$. Furthermore, assume that each A_i has at most $n^2 p'$ edges, and that at most $n^{3/2}$ edges are in more than one A_i and no edge is in at* *least five A_i's. Then with high probability the graph $\tilde{A} := \tilde{A}_1 \cup \cdots \cup \tilde{A}_z$ is such that the red subgraph of $\tilde{A}[V_i, V_j]$ is $(4\varepsilon, \tilde{p}'')$-regular with density at least $\frac{\gamma}{2}\tilde{p}''$.*

Proof. We refer to the red subgraphs of $A, \tilde{A}, A_i, \tilde{A}_i$ for some $i \in [z]$ as $A^r, \tilde{A}^r, A_i^r, \tilde{A}_i^r$, respectively. Let U, W be distinct sets among $V_1, \dots, V_K$. We show that the statement is with high probability satisfied for U, W, which together with a union bound over all such pairs completes the proof.

Claim 7. Let $\nu > 0$, and let $U' \subseteq U$ and $W' \subseteq W$ with $|U'| \geqslant \varepsilon|U|$ and $|W'| \geqslant \varepsilon|W|$. Then with high probability $d_{\tilde{A}^r}(U', W') = (1 \pm \nu)\frac{\tilde{p}''}{p''} d_{A^r}(U', W')$.

Proof. As $A^r[U, W]$ is an (ε, p'')-regular pair, we have $|d_{A^r}(U', W') - d_{A^r}(U, W)| \leqslant \varepsilon p''$, so

$$e_{A^r}(U', W') \geqslant p''(\gamma - \varepsilon)|U'||W'| \geqslant p''(\gamma - \varepsilon)\varepsilon^2\mu^2 n^2 = \Omega(p'' n^2) = \Omega(n^{\frac{3}{2}+\delta} \log n).$$

Now, for each $k \in [z]$, let $\mathcal{D}_k$ denote the set of copies of complete bipartite graphs $K_{C,C}$ formed by two blocks in A_k. For each $D \in \mathcal{D}_k$, denote with X_D the random variable which counts the number of red edges in $\tilde{A}_k$ between U' and W' contained in D. Note that by assumption $|\mathcal{D}_k| \leqslant e(A_k) \leqslant n^2 p'$. Denote $Y_k = e_{\tilde{A}_k^r}(U', W')$ and observe that $Y_k = \sum_{D \in \mathcal{D}_k} X_D$. By Lemma 5.17, we have that the expectation of Y_k satisfies

$$\mathbb{E}[Y_k] = \sum_{e \in E_{A_k^r}(U', W')} \Pr\left[e \in E(\tilde{A}_k)\right] = e_{A_k^r}(U', W')\frac{\tilde{p}'}{p'}.$$

Furthermore, note that Y_k can be viewed as a random variable on a product of probability spaces, where the coordinates are given by X_D, for each $D \in \mathcal{D}_k$. Observe also that changing one coordinate can change Y_k only by at most C^2.

We now want to bound Y_k for each k, and for this we have two cases. In the first case, if $e_{A_k^r}(U', W') \leqslant (\log n)^{-\frac{1}{4}}|U'||W'|p'$, then clearly we have $Y_k \leqslant (\log n)^{-\frac{1}{4}}|U'||W'|p'$. On the other hand, if $e_{A_k^r}(U', W') \geqslant (\log n)^{-\frac{1}{4}}|U'||W'|p'$, by McDiarmid's inequality (Theorem 3.1), and setting $t = \nu\mathbb{E}[Y_k]/10$ we get

$$\Pr[Y_k \notin (1 \pm \nu/10)\mathbb{E}[Y_k]] \leqslant \exp\left(-\frac{\Omega(\mathbb{E}[Y_k]^2)}{|\mathcal{D}_k| C^4}\right) \leqslant \exp\left(-\Omega\left(\frac{(\log n)^{-\frac{1}{2}} n^4 \tilde{p}'^2}{n^2 p'}\right)\right) \leqslant e^{-\Omega(n\sqrt{\log n})}.$$

Now, a union bound over all 2^{2n} possible choices of U', W' and z choices of k shows us that with high probability we have $Y_k \in (1 \pm \nu/10)\mathbb{E}[Y_k]$ for all U', W' and A_k with $e_{A_k^r}(U', W') \geqslant (\log n)^{-\frac{1}{4}}\ |U'|\ |W'|p'$. Hence, having in mind the bounds from both cases, we get with high probability that

$$\begin{aligned} e_{\tilde{A}^r}(U', W') &\leqslant \sum_{k=1}^{z} e_{\tilde{A}_k^r}(U', W') \leqslant \sum_{k=1}^{z} \left(e_{A_k^r}(U', W')\frac{\tilde{p}'}{p'}(1 + \nu/10) + (\log n)^{-\frac{1}{4}} n^2 p'\right) \\ &\leqslant (1 + \nu/10)\left(e_{A^r}(U', W') + 4n^{\frac{3}{2}}\right)\frac{\tilde{p}'}{p'} + o(zn^2 p') \leqslant (1 + \nu)e_{A^r}(U', W')\frac{\tilde{p}''}{p''}, \end{aligned}$$

TABLE 1 Probabilities.

Notation	p	$\tilde{p}$	p'	$\tilde{p}'$	p''	$\tilde{p}''$
Definition	$n^{\delta-1/2}$	$p = 1-(1-\tilde{p})^{\binom{C}{2}}$	$\frac{\log n}{n}$	$p' = 1-(1-\tilde{p}')^{C^2}$	$1-(1-p')^z$	$1-(1-\tilde{p}')^z$
Asymptotics	$n^{\delta-1/2}$	$\frac{n^{\delta-1/2}}{\binom{C}{2}}$	$\frac{\log n}{n}$	$\frac{\log n}{nC^2}$	$\frac{n^{\delta-1/2}\log n}{C}$	$\frac{n^{\delta-1/2}\log n}{C^3}$

where we used that the sum of all red edges between U' and W' over all the A_k's overcounts $e_{A^r}(U',W')$ by at most $4n^{3/2}$, and the fact that $\frac{\tilde{p}'}{p'} \approx \frac{\tilde{p}''}{p''}$ by Table 1. For the lower bound, we let $I \subseteq [z]$ be the set of indices k for which $e_{A_k^r}(U',W') \geqslant (\log n)^{-\frac{1}{4}}|U'||W'|p'$ to get

$$\begin{aligned}
e_{\tilde{A}^r}(U',W') &\geqslant \sum_{k=1}^{z} e_{\tilde{A}_k^r}(U',W') - 4n^{\frac{3}{2}} \geqslant \sum_{k\in I}\left((1-\nu/10)e_{A_k^r}(U',W')\frac{\tilde{p}'}{p'}\right) - 4n^{\frac{3}{2}} \\
&\geqslant (1-\nu/10)\frac{\tilde{p}'}{p'}\left(e_{A^r}(U',W') - z(\log n)^{-\frac{1}{4}}|U'||W'|p'\right) - 4n^{\frac{3}{2}} \\
&\geqslant (1-\nu/5)\frac{\tilde{p}'}{p'}e_{A^r}(U',W') \geqslant (1-\nu)\frac{\tilde{p}''}{p''}e_{A^r}(U',W'),
\end{aligned}$$

using the condition on repeated edges across A_i's, and that $\frac{\tilde{p}'}{p'} \approx \frac{\tilde{p}''}{p''}$ again. □

By Claim 7, we get that $d_{\tilde{A}^r}(U,W) \geqslant \gamma\tilde{p}''/2$. Let $U' \subseteq U$ and $W' \subseteq W$, with $|U'| \geqslant \varepsilon|U|$ and $|W'| \geqslant \varepsilon|W|$. We apply Claim 7 to U',W' and U,W with $\nu := \varepsilon$, and use the triangle inequality to get

$$\begin{aligned}
&\left|d_{\tilde{A}^r}(U',W') - d_{\tilde{A}^r}(U,W)\right| \\
&\leqslant \left|d_{\tilde{A}^r}(U',W') - \frac{\tilde{p}''d_{A^r}(U',W')}{p''}\right| + \frac{\tilde{p}''}{p''}\left|d_{A^r}(U',W') - d_{A^r}(U,W)\right| + \left|\frac{\tilde{p}''d_{A^r}(U,W)}{p''} - d_{\tilde{A}^r}(U,W)\right| \\
&\leqslant \varepsilon\frac{\tilde{p}''d_{A^r}(U',W')}{p''} + \varepsilon\tilde{p}'' + \varepsilon\frac{\tilde{p}''d_{A^r}(U,W)}{p''} \leqslant 2\varepsilon(1+\gamma')\tilde{p}'' + \varepsilon\tilde{p}'' \leqslant 4\varepsilon\tilde{p}'',
\end{aligned}$$

where we used that A, and therefore also A^r, is (γ', p'')-upper-uniform (for $\gamma' \ll \varepsilon$) in the penultimate inequality, and the fact that $A^r[U,W]$ is (ε, p'')-regular for bounding the middle term on the second line. □

6 | THE PROOF

After having done a big part of the work in the previous sections, we are ready to put everything together to show Theorem 1.1.

We start by describing the key constants we use. We need the following inequalities to hold:

$$\{\eta^{-1}, c^{-1}\} \gg C \gg T_3 \gg \varepsilon_3^{-1} \gg C' \gg T_2 \gg \varepsilon_2^{-1} \gg T_1 \gg \varepsilon_1^{-1} \gg \ell \gg \delta^{-1}. \tag{6.1}$$

We can think of each T_i as the upper bound on the number of sets we get from an application of the sparse regularity lemma (Lemma 3.7) with ε_i. Recall that δ is the constant that determines how close the number of edges $n^{\frac{3}{2}+2\delta}$ in the host graph Γ is to $n^{3/2}$ (see Lemma 5.15). It is important to choose ℓ to be an integer so that $\delta > \frac{1}{4\ell-6}$, as we want δ to be large enough to be able to embed cycles of length at least ℓ later. The constant C is the size of the blocks in $G^C(n,p)$ (see Section 5), whereas cn is the number of vertices of the cubic graph H, which we are to embed in our n-vertex host graph. Finally, as indicated in Subsection 5.2, ηn is the maximum number of vertices not covered by each block matching, and C' is a parameter of the densifiers that we will find (see Definition 5.21).

Let H be a cubic graph on cn vertices. We first apply Lemma 4.4 to H to obtain a decomposition into induced cycles of length at least ℓ and an induced subgraph J with treewidth bounded by 2ℓ. Furthermore, there is a blow-up $\mathcal{T}$ of a tree $T \boxtimes K_{400\ell}$, where T is of maximum degree 400ℓ, which contains the graph J. We also may assume that $v(T) \leqslant cn$, as we have that $v(H) \leqslant cn$.

Before diving into the proof, in the following subsection we state some results from [26] and corollaries of them used for embedding monochromatic blow-ups of trees in coloured host graphs. In particular, the host graphs which we use to apply those embedding theorems are the graphs $A_i' \cup M_i$ defined in Section 5. Furthermore, the results from [26] imply that if we appropriately choose the host graph, then it either contains the required blow-up of a tree in one colour, or it satisfies a certain local density property in the other colour. For completeness, we include the slightly altered proofs from [26] in the appendix.

6.1 | Monochromatic blow-ups of trees in coloured expanders

To state the necessary results, we will need the following definition.

Definition 6.1. We say that an n-vertex graph F is α-joint if for every pair of disjoint sets $S, T \subseteq V(H)$ with $|S|, |T| \geqslant \alpha n$ we have $e(S,T) > 0$.

The following result, Theorem 6.3, can be shown by only slightly modifying the proof in [26], and for completeness we include its proof in the appendix. It states that every blow-up of a bounded degree tree can be found as a monochromatic copy in a constant blow-up of a third power of an α-joint graph. They state the result slightly differently, for a random D-regular graph in place of an α-joint graph, but this has little effect on the argument. We first need the following definition.

Definition 6.2. Let $\mathcal{T}_{n,d}$ be the set of all trees with n vertices and maximum degree at most d. Furthermore, let $\mathcal{T}_{n,d}(k)$ be the family of all graphs $T \boxtimes K_k$ where $T \in \mathcal{T}_{n,d}$.

Note that $\mathcal{T}$ belongs to $\mathcal{T}_{cn,400\ell}(400\ell)$.

Theorem 6.3 [26, Theorem 3.4]. *Let $\alpha \ll c' \ll r^{-1} \ll \{k^{-1}, d^{-1}\}$. Let $\mathcal{G}$ be an α-joint graph on n vertices. Then any red/blue colouring of $\mathcal{G}^3 \boxtimes K_r$ contains a monochromatic copy of each graph in $\mathcal{T}_{c'n,d}(k)$. Furthermore, all graphs in $\mathcal{T}_{c'n,d}(k)$ can be found in $\mathcal{G}^3 \boxtimes K_r$ in the same colour.*

In the rest of the paper, we (evidently) rely on the various parts of our host graph construction, so we refer the reader to Section 5, and in particular to Definition 5.14.

Given a small linear-sized subset S of vertices of our host graph Γ, the following lemma shows the existence of a copy of $\mathcal{T}$ in either the red or the blue subgraph of $(A_i' \cup M_i)[S]$.

Lemma 6.4. *Let $L = A_i' \cup M_i$ for some i. The following holds with high probability for each red/blue colouring of L. Let S be a subset of $V(L)$ of size $|S| = \gamma n$ for $\gamma \gg C^{-1}$. Then either the red or the blue subgraph of $L[S]$ contains $\mathcal{T}$.*

Proof. By Lemma 3.12 applied to the cliques in M_i and the set S, we get a collection $\mathcal{B}$ of at least $\gamma \frac{n}{2C}$ disjoint cliques B such that $B \subseteq S$ and $|B| = \gamma C/8$, where each B is contained in a distinct clique from M_i; here we also used that M_i covers at least $(1-\eta)n$ vertices where $\eta \ll \gamma$ by (6.1). Let F be the subgraph of G_i induced by the vertices of G_i corresponding to cliques in M_i which contain a set B from $\mathcal{B}$.

As the number of vertices of F is $v(F) = |\mathcal{B}| \geqslant \frac{\gamma}{2} v(G_i)$, we have that F is α-joint for all constants α. Indeed, as G_i is a binomial random graph with expected degree logarithmic in its number of vertices, with high probability every pair of linear-sized subsets of vertices has an edge in between (by a standard Chernoff bound).

Now, look at the copy of the graph $F^3 \boxtimes K_{\gamma C/8}$ in L corresponding to the cliques $\mathcal{B}$. As F is with high probability α-joint for an arbitrarily small α (in particular also for $\alpha \ll cC/\gamma$), we infer by Theorem 6.3, setting $c' = \frac{cn}{v(F)} \leqslant \frac{2cC}{\gamma(1-\eta)}$, $r = \gamma C/8$, $k = d = 400\ell$, that the considered copy of $F^3 \boxtimes K_{\gamma C/8}$ in $L[S]$ either contains a red or a blue copy of $\mathcal{T}$. $\square$

Now we show a proposition which states that for each i, either the blue subgraph of $A_i \cup M_i$ contains $\mathcal{T}$, or the red subgraph of $A_i \cup M_i$ satisfies a certain local density property. To do that, we will need the following definition (which also appears in [26] and other prior work), together with a theorem which is implicit in [26] and whose proof can be found in our appendix.

Definition 6.5. For integers s and m, a graph $\mathcal{G}$ with edge-colouring $\psi : E(\mathcal{G}) \to \{\text{red}, \text{blue}\}$ and a vertex partition $(V_1, V_2, \dots, V_m)$ of $\mathcal{G}$, we define the following auxiliary colouring of K_m. For vertices $i, j \in [m]$ of K_m, the edge ij is coloured blue if the bipartite graph between V_i and V_j in $\mathcal{G}$ contains a blue $K_{s,s}$, and red otherwise. This edge-colouring is referred to as the *$(\mathcal{G}, \psi, s)$-colouring of K_m*.

Theorem 6.6 (Proof of Theorem 3.4 in [26]). *Fix integers n_0, d, k, q. Let $s = (d + d^2)k$ and $m \geqslant 20 n_0 d^2 q$. Let $\mathcal{K} = T \boxtimes K_k$ for an n_0-vertex tree T of maximum degree d. Suppose we are given a graph $\mathcal{G}$, a vertex partition $(V_1, V_2, \dots, V_m)$ of $\mathcal{G}$, and an edge-colouring $\psi : E(\mathcal{G}) \to \{red, blue\}$ such that, for all $i \in [m]$, all the edges of $\mathcal{G}[V_i]$ are present and are blue, and $|V_i| \geqslant s$. If $\mathcal{G}$ does not contain a blue copy of $\mathcal{K}$, then there is a red copy of a complete q-partite graph in the $(\mathcal{G}, \psi, s)$-colouring of K_m, such that every part has size at least $\frac{m}{5d^2 q}$.*

We will also need the following deterministic statement, which does not depend on the outcome of $G_1, \dots, G_z$.

Proposition 6.7. *Let $\gamma = \frac{1}{10^3 T_1}$. For every $\frac{10^{11} c\ell^2 C'}{\gamma} < \rho < \frac{\gamma}{20}$, the following holds. Let S be a subset of vertices of Γ with $|S| = \gamma n$ such that the blue subgraph of $\Gamma[S]$ does not contain $\mathcal{T}$. Then for each i one of the following is true.*

(i) *$A_i[S]$ contains a red $(C', \frac{\gamma\rho}{10^{10}\ell^2}, s, q)$-densifier*[†] *with $s = (400\ell)^2(400\ell + 1)$ and $q = 1000$.*

(ii) *There are at least $\frac{|S|}{8C}$ cliques $B \in M_i$ such that $|B \cap S| \geqslant \gamma C/8$ and with some $B' \subseteq B \cap S$ of size at most ρC s.t. there is no blue $K_{C'}$ contained in $B \cap S - B'$.*

Proof. For each $i \in [z]$, let $L_i = A_i \cup M_i$. As in the proof of Lemma 6.4, by Lemma 3.12 applied to the cliques in M_i and the set S, we get a collection $\mathcal{B}$ of at least $\gamma\frac{n}{2C}$ disjoint cliques B such that $B \in M_i$ and $|B \cap S| \geqslant \gamma C/8$. Let F_i be the subgraph of G_i induced by the vertices of G_i corresponding to blocks in $\mathcal{B}$.

Now for each $i \in [z]$, look at the copy of the graph $F_i \boxtimes K_{\gamma C/8}$ in $L_i[S]$, corresponding to the cliques $\mathcal{B}$ (technically, the cliques in the blow-up are of size at least $\gamma C/8$ and not necessarily precisely $\gamma C/8$, so we abuse notation slightly here). For each $v \in F_i$, we refer to the copy of K_t with $t \geqslant \gamma C/8$ that corresponds to v in $F_i \boxtimes K_{\gamma C/8}$ as $B_i(v)$. Then one of the following occurs for each i.

(a) At least half of the vertices $v \in F_i$ are such that $B_i(v)$ contains at least $\rho\frac{C}{C'}$ many vertex-disjoint blue copies of $K_{C'}$.

(b) At least half of the vertices $v \in F_i$ are such that $B_i(v)$ contains some $B_i'(v) \subseteq B_i(v)$ of size at most ρC so that $B_i(v) - B_i'(v)$ has no blue $K_{C'}$.

Indeed, one can remove blue copies of $K_{C'}$ from each $B_i(v)$ repeatedly until either at least ρC vertices are covered or there are no blue copies of $K_{C'}$ remaining.

If (b) holds, there are at least $\frac{v(F_i)}{2} \geqslant \frac{\gamma n}{8C}$ cliques $B_i(v)$, each a subset of a distinct block $B \in M_i$ which fulfills the requirements of (ii) with $B_i'(v)$ as B', so we are done.

Otherwise, if (a) holds, we consider the subgraph F_i' of F_i induced by the vertices $v \in F_i$ such that $B_i(v)$ contains at least $\rho\frac{C}{C'}$ many vertex-disjoint blue copies of $K_{C'}$. For each such v, let $B_i'(v) \subseteq B_i(v)$ be the subset of size at least ρC covered by copies of blue $K_{C'}$. Then consider the graph $F_i' \boxtimes K_{\rho C} \subseteq F_i \boxtimes K_{\gamma C/8}$ with the copies of $K_{\rho C}$ corresponding to the $B_i'(v)$'s. We apply Theorem 6.6 to $F_i' \boxtimes K_{\rho C}$ with the blue copies of $K_{C'}$ as a vertex partition $(V_1, \dots, V_m)$ with

$$m := \frac{v(F_i)}{2}\frac{\rho C}{C'} \geqslant \frac{\gamma\rho n}{8C'}.$$

and with $q = 1000$, $d = k = 400\ell$, $n_0 = cn$ and $s = (d^2 + d)k$. We can do this as $m \geqslant \frac{\gamma\rho n}{8C'} \geqslant 20cn(400\ell)^2 \cdot 1000$ and $C' \gg s$, by (6.1). As the blue subgraph of $F_i' \boxtimes K_{\rho C}$ does not contain $\mathcal{T}$ by assumption, there is a red copy of a complete q-partite graph in the $(F_i' \boxtimes K_{\rho C}, \psi, s)$-colouring of K_m, such that every part has size at least $\frac{m}{5d^2q}$, where ψ is the considered colouring of Γ restricted to $F_i' \boxtimes K_{\rho C}$. This means that there are q collections $W_1, \dots, W_q$ of blue copies of $K_{C'}$, which are subsets of the copies $B_i'(v)$ of $K_{\rho C}$ and are also pairwise vertex-disjoint (even across different W_i's). Furthermore, the collections $W_1, \dots, W_q$ have the property that $|W_i| = \frac{m}{5d^2q}$ and for every pair X, Y with $X \in W_i$ and $Y \in W_j$ where $i \neq j$, there is no blue $K_{s,s}$ in the complete bipartite graph between X and Y. This corresponds precisely to a red $(C', \frac{\gamma\rho}{10^{10}\ell^2}, s, q)$-densifier of $A_i[S]$, as

$$|W_i| \geqslant \frac{m}{5d^2q} \geqslant \frac{\gamma\rho n}{40C'(400\ell)^2q} \geqslant \frac{\gamma\rho}{10^{10}\ell^2}\frac{n}{C'}.$$ □

[†] See Definition 5.21.

6.2 | Embedding the cubic graph

Proof of Theorem 1.1. We show that for every $\delta > 0$ there exists a $c > 0$, such that for every n large enough and every cubic graph H on cn vertices, there is an n-vertex host graph with at most $n^{\frac{3}{2}+2\delta}$ edges that is Ramsey for H.

We first give a summary of the proof. The host graph is as described in Section 5. We distinguish between two cases, depending on whether each relatively large induced subgraph of the host graph contains a copy of the bounded treewidth part $\mathcal{T}$ of H in each colour. If that is the case, we have a lot of flexibility and we can afford to embed the rest of H in either colour. We make use of standard techniques to find disjoint vertex sets $V_1, \dots, V_{21}$, all pairs of which are regular and dense in one colour, say red. We use $V_1, \dots, V_{20}$ to embed the long induced cycles, following the strategy from [14]. Finally, we embed $\mathcal{T}$ in V_{21} in the same colour as the cycles, which is possible by assumption. In the second case, there is a relatively large vertex set in the host graph with no copy of $\mathcal{T}$ in say blue. The challenge then is that we have to embed all of H in red, as $\mathcal{T}$ may not exist in blue at all (it exists in red due to Lemma 6.4). Thus, we need to guarantee some density in red. We do that with the help of Proposition 6.7, which allows us to deduce, as there is no blue $\mathcal{T}$, that there is either a red densifier, or many cliques that do not contain too many blue copies of $K_{C'}$. These are the respective set-ups for Lemmas 5.22 and 5.9, which give us the desired sets $V_1, \dots, V_{21}$ that are pairwise regular and dense in red. From here we can proceed as in the first case.

We start by formally describing our host graph. Let Γ be an outcome of $G \cup A'_1 \cup \dots \cup A'_z$ as defined in Section 5, where the graphs $G_1, \dots, G_z$, which give rise to $A'_1, \dots, A'_z$, satisfy the following conditions: the conclusions of Lemmas 5.15, 5.19, 5.20, 5.22 and 6.4 hold[†], and additionally the following property holds. If we now take $\tilde{A}_1, \dots, \tilde{A}_z$ to be subsampled from $A_1, \dots, A_z$ as in Definition 5.16, then with probability at least 0.9, the conclusions of both Lemma 3.9 and Theorem 3.11 hold for $\tilde{A}_1 \cup \dots \cup \tilde{A}_z$. Note that such $G_1, \dots, G_z$ exist by the same argument as the one used for fixing an outcome G of $G^C(n, p)$ in the paragraph before Definition 5.14, except we now use that by Lemma 5.18, $\tilde{A}_1 \cup \dots \cup \tilde{A}_z$ can be coupled as a subgraph of $G(n, \tilde{p}'')$.

Let U be the largest subset of $V(\Gamma)$ such that either the red or the blue subgraph of $\Gamma[U]$ does not contain $\mathcal{T}$. We distinguish two cases, depending on whether U has size at most $\frac{n}{10^3 T_1}$ or not. For example, if Γ is completely red, then we would be in the latter case.

Case I

In the former case, we can embed $H \setminus \mathcal{T}$ in either colour, as we have sufficiently many and well-distributed copies of $\mathcal{T}$ in both colours. We apply the sparse regularity lemma (Theorem 3.7) to $\tilde{G}$ with ε_1 and a large enough t_1. Note that we can do this as we assumed G is such that with probability at least 0.9, $\tilde{G}$ is $(\zeta, \tilde{p})$-upper-uniform for all $\zeta > 0$ (recall the paragraph above Definition 5.14). Using a standard argument invoking Turán's theorem and Ramsey's theorem (see, for example, the proof of [29, Lemma 19]), we obtain 21 sets $V'_1, \dots, V'_{21}$ which are pairwise $(\varepsilon_1, \tilde{p})$-regular and where each pair has density at least $d := \gamma\tilde{p}$ in say red, where $\gamma = 1/4$. We remove the bad subset for each pair V'_i, V'_j via Lemma 3.9 applied with some ε'_1 such that $\varepsilon_1 \ll \varepsilon'_1 \ll \{\ell^{-1}, \gamma\}$ (which is possible by (6.1)). Note that the lemma also guarantees that the density of the pairs does not drop to less than $(1 - \varepsilon'_1)d$. We then take the subsets given by Lemma 3.8 to get $V_1, \dots, V_{21}$, each of size at least $\tilde{n} \geqslant \frac{n}{2T_1}$, such that the red subgraph F of $\tilde{G}$ on those vertices satisfies the following.

- $d_F(v, V_j) \geqslant \tilde{n}d(1 - \varepsilon'_1)/2$ for each $v \in V_i$ and $i \neq j$.

[†] As in footnote , we specify the constants in the usage of these lemmas later.

- for all distinct i, j, h, g (but possibly $h = g$), for each $v \in V_h, w \in V_g$, and any $N_1 \subseteq N_F(v, V_i)$ and $N_2 \subseteq N_F(w, V_j)$ of size $|N_1| = |N_2| = \frac{\tilde{n}d}{20}$, (N_1, N_2) and (N_1, V_j) are (ε', p)-regular of density at least $\frac{d}{2}$ in F.

The first step of our embedding procedure is to embed the graph $\mathcal{T}$ in the red subgraph of $\Gamma[V_{21}]$, which we can do as V_{21} is of size at least $\frac{n}{2T_1}$, so $\Gamma[V_{21}]$ contains both a red and a blue copy of $\mathcal{T}$ by the assumption of Case I. What is left is to use the remaining vertices in $V_1, \dots, V_{20}$ to embed the induced cycles from the decomposition. By definition, those induced cycles are such that every vertex in each of them has at most one neighbour in the previously embedded part of H. So for each such vertex, the 'candidate set' (i.e. the set where this vertex can be embedded) in each V_i is of size at least $\tilde{n}d(1 - \varepsilon_1')/2$. Now we use the technique developed in [14] to embed those graphs in $V_1, \dots, V_{20}$. The only difference in our case is that we start the embedding process with some candidate sets of vertices which have a neighbour in a graph from $\mathcal{T}$, which is precisely the set-up for using Theorem 3.11. Recall that by our choice of G in the paragraph above Definition 5.14, Theorem 3.11 is applicable to $\tilde{G}$ with probability at least 0.9. We apply it with $\varepsilon := \varepsilon_1'$, noting that $\{\ell^{-1}, \gamma\} \gg \varepsilon_1' \gg \varepsilon_1$, which finishes the proof in this case.

Case II

In the latter case, there is some U of size $|U| > \frac{n}{10^3 T_1}$ such that either the red subgraph or the blue subgraph of $\Gamma[U]$ does not contain $\mathcal{T}$. Suppose without loss of generality that this holds for the blue subgraph of $\Gamma[U]$. Note that by Lemma 6.4, any subset U' of U of size ϕn with $\phi \gg C^{-1}$ is such that the red subgraph of $\Gamma[U']$ contains $\mathcal{T}$.

We now restrict ourselves to a subset V' of U that has size precisely ϕn where $\phi = \frac{1}{10^3 T_1}$. As the blue subgraph of $\Gamma[V']$ does not contain $\mathcal{T}$, we can apply Proposition 6.7. We do so with $S := V'$ and ρ such that $\varepsilon_2 \ll \rho \ll T_1^{-1}$. Note that the conditions for ρ can be satisfied by (6.1). We get that for each i, one of the following holds.

(a) $A_i[V']$ contains a red $(C', \frac{\rho}{10^{13}\ell^2 T_1}, s, q)$-densifier with $s = (400\ell)^2(400\ell + 1)$ and $q = 1000$.
(b) There are at least $|V'|/8C$ cliques $B \in M_i$ such that $|B \cap V'| \geqslant \frac{C}{8 \cdot 10^3 T_1}$ and with some $B' \subseteq B \cap V'$ of size at most ρC s.t. there is no blue $K_{C'}$ in $B \cap V' - B'$.

We again distinguish between two cases, depending on whether at least $z/2$ of indices $i \in [z]$ satisfy (a) or (b) (recall that z is the number of A_i's).

If (a) is more common, we apply Lemma 5.22 with $R := V'$, $\alpha := \frac{1}{10^3 T_1}$, $\gamma := \frac{\rho}{10^{13}\ell^2 T_1}$, $q := 1000$, $\mu := T_2^{-1}$, $s := (400\ell)^2(400\ell + 1)$ and $\varepsilon := \varepsilon_2$, which we can do as $C' \gg T_2$ and $T_2^{-1} \ll \varepsilon_2 \ll T_1^{-1}$. We get $V_1', \dots, V_{21}'$ of size at least μn, such that all pairs are (ε_2, p'')-regular in the red subgraph of $A_1 \cup \dots \cup A_z$ with density at least $\tau p''$ where $\tau = \tau(T_1, \ell) \gg \varepsilon_2$. As the outcome of Lemma 5.19 holds, the graph $A_1 \cup \dots \cup A_z$ is (γ', p'')-upper-uniform for every constant $\gamma' > 0$, and furthermore each A_i has at most $n^2 p'$ edges by standard concentration bounds (as the expected number of its edges is $\binom{n_i}{2} p' C^2 < n^2 p'/2$). Hence, we can apply Lemma 5.23 to $V_1', \dots, V_{21}'$ and get that all pairs are $(4\varepsilon_2, \tilde{p}'')$-regular with density at least $\tau \tilde{p}''/2$ in the red subgraph of $\tilde{A}_1 \cup \dots \cup \tilde{A}_z$.

If (b) is more common, we apply Lemma 5.9 (possibly taking supersets of the sets B') with $R := V'$, $\alpha := \frac{1}{10^3 T_1}$, $\mu := T_3^{-1}$, $\varepsilon := \varepsilon_3$, which we can do as $\{T_1, C'\} \ll \varepsilon_3^{-1} \ll \{C, T_3\}$. We get $V_1', \dots, V_{21}'$ of size μn, such that all pairs are (ε_3, p)-regular in the red subgraph of G with density at least τp, where $\tau = \tau(C', T_1) \gg \varepsilon_3$. Similarly to the previous paragraph, as G is (γ', p)-upper-uniform for fixed $\gamma' > 0$ by Lemma 5.3, we can now apply Lemma 5.10 to get that all pairs are $(4\varepsilon_3, \tilde{p})$-regular with density at least $\tau \tilde{p}/2$ in the red subgraph of $\tilde{G}$.

In both cases, we can now proceed as in Case I, substituting $\tilde{p}''$ for $\tilde{p}$ if (a) is more common. Furthermore, the density of the regular pairs in red is now $\tau\tilde{p}''/2$ if (a) is more common or $\tau\tilde{p}/2$ if (b) is more common. This density is much larger than $\varepsilon_2\tilde{p}''$ and $\varepsilon_3\tilde{p}$, respectively, which enables us to use the same embedding strategy as in Case I. For completeness, we provide the details below. Let $\varepsilon := \varepsilon_2$ and $\pi := \tilde{p}''$ if (a) is more common and $\varepsilon := \varepsilon_3$ and $\pi := \tilde{p}$ if (b) is more common, and consider the appropriate τ and μ, where $d := \tau\pi$ is the density of the regular pairs and μn is a lower bound on the size of the sets V_i. Finally, let $G_0 = (\tilde{A}_1 \cup \cdots \cup \tilde{A}_z)[V']$ if (a) is more common and $G_0 = \tilde{G}[V']$ if (b) is more common.

We remove the bad subset for each pair V_i', V_j' via Lemma 3.9 applied with ε' such that $\varepsilon \ll \varepsilon' \ll \{\ell^{-1}, \tau\}$, and take the subsets given by Lemma 3.8 to get $V_1, \ldots, V_{21}$, each of size at least $\tilde{n} := \frac{\mu n}{2}$, such that the red subgraph F of G_0 on those vertices satisfies the following.

- $d_F(v, V_j) \geqslant \tilde{n}d(1-\varepsilon')/2$ for each $v \in V_i$ with $i \neq j$.
- for all distinct i, j, h, g (but possibly $h = g$), for each $v \in V_h, w \in V_g$, and any $N_1 \subseteq N_F(v, V_i)$ and $N_2 \subseteq N_F(w, V_j)$ of size $|N_1| = |N_2| = \frac{\tilde{n}d}{20}$, (N_1, N_2) and (N_1, V_j) are (ε', π)-regular of density at least $\frac{d}{2}$ in F.

Recall that any subset $U' \subseteq V'$ of size $\phi' n$ with $\phi' \gg C^{-1}$ is such that the red subgraph of $\Gamma[U']$ contains $\mathcal{T}$. As $|V_{21}| \geqslant \frac{\mu n}{2} \gg C^{-1}n$, we can embed $\mathcal{T}$ in the red subgraph of $\Gamma[V_{21}]$. What is left is to use the remaining vertices in $V_1, \ldots, V_{20}$ to embed the induced cycles from the decomposition. For each vertex, the 'candidate set' (i.e. the set where this vertex can be embedded) in each V_i is of size at least $\tilde{n}d(1-\varepsilon')/2$. Due to our assumptions on G and $G_1, \ldots, G_z$ from the paragraph above Definition 5.14 and the second paragraph of the current proof, with probability at least 0.9, the graph G_0 and its red subgraph F are such that we can apply Theorem 3.11 to them. We do so with $\varepsilon := \varepsilon'$, noting that $\{\ell^{-1}, \tau\} \gg \varepsilon' \gg \varepsilon$, which finishes the proof. □

7 | CONCLUDING REMARKS

In this paper, we have shown that the size-Ramsey number of n-vertex cubic graphs is of order $O\left(n^{3/2+o(1)}\right)$. In fact, our proof gives a stronger universality result, for any 2-colouring of the $n^{3/2+o(1)}$ edges of our host graph, there is a colour class which contains all cubic graphs on n vertices. On the other hand, it is known that any graph which contains all n-vertex cubic graphs must have $\Omega(n^{4/3})$ edges, even without colouring (see [1]). Hence, the optimal partition universal graph for the class of all n-vertex cubic graphs has at least $\Omega(n^{4/3})$ and at most $n^{3/2+o(1)}$ edges, and it is not clear to us where the truth lies.

Going back to size-Ramsey numbers of n-vertex cubic graphs H, it might be true that in general $\hat{r}(H) = o(n^{4/3})$, but in that case an upper bound proof would require several distinct host graph constructions for different cubic graphs H. But it is even not completely clear that a general upper bound of $o(n^{3/2})$ should hold. Our proof technique reaches certain hard natural barriers, the most significant one being that at density $p = o(n^{-1/2})$, at least in 'uniformly' dense graphs, regularity inheritance between the candidate sets is no longer guaranteed. That was essential to our approach, as we relied on the regularity method and the KŁR conjecture to embed cycles into the host graph. Therefore, if possible, pushing the upper bound below $n^{3/2}$ would certainly require new ideas and a different approach. Let us note here that in the special case of the grid graph, Conlon, Nenadov and Trujić [13] managed to overcome the regularity inheritance barrier, by using a host graph tailored for the grid graph, exploiting its structural properties. In particular, parts of their host graph are locally very dense, so they are able to use regularity inheritance locally, and get away with using much smaller global density. They get the bound of $\hat{r}(H) = O(n^{5/4})$ where H is the grid graph on n vertices. Finally, for the class of cubic graphs, one could alternatively hope to avoid using regularity inheritance, but this would most probably require entirely new embedding techniques.

Recall that Kohayakawa, Rödl, Schacht and Szemerédi [29] showed that $\hat{r}(H) \leqslant n^{2-1/\Delta+o(1)}$ for all n-vertex graphs H with maximum degree Δ, which was improved to $\hat{r}(H) \leqslant n^{2-\frac{1}{\Delta-1/2}+o(1)}$ by Nenadov [34] in the special case when H does not contain a triangle and when $\Delta \geqslant 5$. If one tries to generalise our approach to arbitrary bounded Δ to show a bound of $\hat{r}(H) \leqslant n^{2-1/(\Delta-1)+o(1)}$ with the appropriate modifications, everything goes through, except for the regularity inheritance of the candidate sets. More precisely, the candidate sets are now the common neighbourhoods of tuples of already embedded vertices, and hence it is significantly harder to make sure that those common neighbourhoods behave well in the sense of regularity inheritance, even though they will typically be of large enough size if one assumes edge probability $p = n^{-1/(\Delta-1)+o(1)}$. It is quite possible that by embedding the parts from the decomposition more carefully, one can control the choice of tuples so that regularity is still inherited, but we chose not to pursue this in this paper. It would certainly be interesting to see if this can be done.

Our proof does not extend to more than two colours, primarily due to the reliance on the machinery in [26] for bounded treewidth embeddings, which is similarly restricted to two colours. Utilising tools from [5], along with further technical adjustments, could be a plausible direction for generalising the result to more colours.

MISSING PROOFS A

For completeness, we provide details on some tools we use in the proof of our main theorem. Most of the exposition here follows closely [26], but we show the adjustments necessary for our applications.

We make use of a well-known result by Friedman and Pippenger [21]. For a graph H and $X \subseteq V(H)$, let $\Gamma_H(X)$ be the set of vertices in $V(H)$ adjacent to some vertex in X. We say that a graph H is (s, d)-*expanding* if for every set $X \subseteq V(H)$ with $1 \leqslant |X| \leqslant s$, it holds that $|\Gamma_H(X)| \geqslant d|X|$.

Lemma A.1 [21]. *If H is a non-empty $(2n-2, d+1)$-expanding graph, then it contains every tree with n vertices and maximum degree at most d.*

The following lemma shows that if all sets of certain size s expand well, then one can remove a small number of vertices to obtain a graph where all sets of size at most s expand well.

Lemma A.2 [17, Lemma 3.1.]. *Let G be a graph such that $|\Gamma_G(X)| \geqslant 3Ks$ for every subset $X \subseteq V(G)$ of size $|X| = s$, for some $s \in \mathbb{N}$ and $K \geqslant 1$. Then there exists a subset $B \subseteq V(G)$ of size $|B| < s$ such that $G - B$ is (s, K)-expanding.*

The next lemma shows a connection between α-joint graphs (recall Definition 6.1) and expanding graphs.

Lemma A.3. *Let G be an n-vertex α-joint graph for some $\alpha > 0$ and let $d \geqslant 1$. Then every induced subgraph of G on at least $10\alpha dn$ vertices contains a non-empty $(\alpha n, d)$-expanding subgraph.*

Proof. Recall that as G is α-joint (Definition 6.1), every pair of disjoint vertex sets of size at least αn have an edge between them. Let G' be an induced subgraph of G on at least $10\alpha dn$ vertices. Every subset S of αn vertices of G' has at least $10\alpha dn - 2\alpha n$ neighbours in G', as there can be only be at most αn vertices outside of S in G' without a neighbour in S, as G is α-joint. So, all sets S of size αn have $|\Gamma_{G'}(S)| \geqslant \frac{10\alpha dn - 2\alpha n}{\alpha n}|S| \geqslant 8d|S|$. Now, by Lemma A.2, there is a subgraph of G' on at least $10\alpha dn - \alpha n$ vertices, which is $(\alpha n, d)$-expanding. $\square$

We also need a lemma from [26], which gives a dichotomy in K_N between the containment of all trees in $\mathcal{T}_{n,d}$ on the one hand, and a complete q-partite graph on the other hand, for appropriately chosen N with respect to n, d, q.

Lemma A.4 [26, Lemma 3.1]. *Fix integers n, d, q and let $N \geqslant 20ndq$. In every red/blue-colouring of $E(K_N)$ there is either a blue copy of every tree in $\mathcal{T}_{n,d}$, or a red copy of a complete q-partite graph in which every part has size at least $\frac{N}{5dq}$.*

The following definition is precisely the same as in [26]. For a tree T with root r, define the *truncation* T' of T as the tree obtained from T by removing each vertex v at a positive even distance from r, and for each such v, adding an edge from the parent of v to each child of v in T. Observe that the maximum degree of T' is at most d^2, where $d := \Delta(T)$.

The next lemma shows that, if $G[V_i]$ is a blue clique for all $i \in [m]$, the existence of a blue tree in the (G, ψ, s)-colouring of K_m implies there is a blue blow-up of a related tree in G.

Lemma A.5 [26, Lemma 3.2]. *Fix integers n_0, d, k, m. Let T be a tree in $\mathcal{T}_{n_0,d}$ rooted at x_0, and let T' be the truncation of T. Let $s = (d + d^2)k$. Suppose we are given a graph G, a vertex partition $(V_1, V_2, \dots, V_m)$ of G, and an edge-colouring $\psi : E(G) \to \{red, blue\}$ such that, for all $i \in [m]$, all the edges of $G[V_i]$ are present and are blue, and $|V_i| \geqslant s$. If there exists a blue copy of T' in the (G, ψ, s)-colouring of K_m, then there exists a blue copy of $T \boxtimes K_k$ in G.*

The next theorem, which we use in the proof of Proposition 6.7, combines the two previous lemmas to show that if a colouring of a certain blow-up does not contain a blow-up of some bounded degree tree in one colour, then it is dense in the other colour. Recall the definition of a $(\mathcal{G}, \psi, s)$-colouring (Definition 6.5).

Theorem 6.6 (Proof of Theorem 3.4 in [26]). *Fix integers n_0, d, k, q. Let $s = (d + d^2)k$ and $m \geqslant 20n_0d^2q$. Let $\mathcal{K} = T \boxtimes K_k$ for an n_0-vertex tree T of maximum degree d. Suppose we are given a graph $\mathcal{G}$, a vertex partition $(V_1, V_2, \dots, V_m)$ of $\mathcal{G}$, and an edge-colouring $\psi : E(\mathcal{G}) \to \{red, blue\}$ such that, for all $i \in [m]$, all the edges of $\mathcal{G}[V_i]$ are present and are blue, and $|V_i| \geqslant s$. If $\mathcal{G}$ does not contain a blue copy of $\mathcal{K}$, then there is a red copy of a complete q-partite graph in the $(\mathcal{G}, \psi, s)$-colouring of K_m, such that every part has size at least $\frac{m}{5d^2q}$.*

Proof. Fix an arbitrary root x_0 of T, and let T' be the truncation of T. By Lemma A.5, as there is no blue copy of $T \boxtimes K_k$ in $\mathcal{G}$, there is no blue copy of T' in the $(\mathcal{G}, \psi, s)$-colouring of K_m. Note that T' belongs to $\mathcal{T}_{n_0,d^2}$. Now Lemma A.4 applied to K_m tells us that there is a red copy of a complete q-partite graph in which every part has size at least $\frac{m}{5d^2q}$. $\square$

For a graph F, we denote by $F\{t\}$ the graph obtained from F by replacing each vertex v by an independent set $I(v)$ of size t, and every edge vw by a complete bipartite graph between the sets $I(v)$ and $I(w)$.

Lemma A.6 [26, Lemma 3.3]. *Fix $t \geqslant 1$. Let F be a graph with maximum degree Δ. Let F' be a spanning subgraph of $F\{t\}$ such that for every edge $vw \in E(F)$ there are at least $(1-\frac{1}{8\Delta})t^2$ edges in F' between $I(v)$ and $I(w)$. Then $F \subseteq F'$.*

With all these ingredients at hand, we are now ready to show a modified version of a theorem in [26], which we use in the proof of Lemma 6.4.

Theorem 6.3 [26, Theorem 3.4]. *Let $\alpha \ll c' \ll r^{-1} \ll \{k^{-1}, d^{-1}\}$. Let $\mathcal{G}$ be an α-joint graph on n vertices. Then any red/blue colouring of $\mathcal{G}^3 \boxtimes K_r$ contains a monochromatic copy of each graph in $\mathcal{T}_{c'n,d}(k)$. Furthermore, all graphs in $\mathcal{T}_{c'n,d}(k)$ can be found in $\mathcal{G}^3 \boxtimes K_r$ in the same colour.*

Proof. Let t be a constant such that $\{k, d\} \ll t \ll r$, and let $s = (d^2 + d)k$. Let $A(v)$ be the copy of K_r that corresponds to $v \in V(\mathcal{G})$. Denote $\mathcal{G}^3 \boxtimes K_r$ by G. Fix an edge-colouring $\psi : E(G) \to \{\text{red}, \text{blue}\}$ of G.

As we can assume that r is at least the Ramsey number $r(t)$, for every $v \in V(\mathcal{G})$ we conclude that $A(v)$ contains a monochromatic copy of K_t, which we denote by $B(v)$. Now, let W be the set of all vertices $v \in V(\mathcal{G})$ in which $B(v)$ is blue. By symmetry between blue and red, we can assume that $|W| \geqslant \frac{1}{2}|V(\mathcal{G})|$. Let $N = |W| \geqslant \frac{n}{2}$.

We define $B(W) = \bigcup_{v \in W} B(v)$ and take φ to be the $(G[B(W)], \psi, s)$-colouring of K_N. If the blue subgraph of K_N contains all trees in $\{T' | T \in \mathcal{T}_{c'n,d}\}$, then by Lemma A.5, the blue subgraph of $G[B(W)]$ contains all graphs in $\mathcal{T}_{c'n,d}(k)$.

From now on, we assume the blue subgraph of K_N does not contain all trees in $\{T' | T \in \mathcal{T}_{c'n,d}\}$. As each T' in this family has $\Delta(T') \leqslant d^2$ and $N \geqslant 20c'nd^2(2k+1)$, by Lemma A.4 there is a family of sets $V_0, V_1, \dots, V_{2k} \subseteq V(K_N)$, each of size at least $\frac{N}{5d^2(2k+1)}$, such that for each $i \neq j$, the complete bipartite graph between V_i and V_j in K_N contains only red edges.

Let an *i-matching* in $\mathcal{G}$ be a matching which consists of edges each incident to one vertex in V_0 and to one vertex in V_i, where $i \in [2k]$. In what follows, we construct a set $S \subseteq V_0$ of size $|S| \geqslant 2^{-2k}|V_0|$ and $2k$ many i-matchings $\{M_i\}_{i=1}^{2k}$, each of which covers S. This is done inductively on i, taking $S_0 := V_0$ as the base case with $i = 0$. Suppose for some $j \leqslant 2k - 1$, we have a set $S_j \subseteq V_0$ such that $|S_j| \geqslant 2^{-j}|V_0|$ and j many i-matchings $\{M_i\}_{i=1}^{j}$ such that M_i covers S_j for each $i \in [j]$. Take a maximum matching M_{j+1} between S_j and V_{j+1}. Suppose for contradiction that M_{j+1} has less than $|S_j|/2$ edges. Consider the vertex sets $X \subset S_j$ and $Y \subset V_{j+1}$ consisting of all vertices that are not incident to edges in M_{j+1}. Note that by the maximality of M_{j+1}, there are no edges between X and Y. As $|X|, |Y| \geqslant |S_j|/2 \geqslant 2^{-2k-2}|V_0| > \alpha n$, this contradicts $\mathcal{G}$ being α-joint. Therefore, at least $|S_j|/2 \geqslant |V_0| \cdot 2^{-(j+1)}$ vertices of S_j are covered by M_{j+1}. Setting $S_{j+1} = V(M_{j+1}) \cap S_j$ at each step, we get the set $S := S_{2k}$ after $2k$ steps, which has the desired properties.

Let $v_i \in V_i$ be the only neighbour of v in M_i, where $v \in S$ and $i \in [2k]$. As $|S| \geqslant 2^{-2k}|V_0| > 20c'(d+1)n$, $\mathcal{G}[S]$ contains all trees in $\mathcal{T}_{c'n,d}$ by Lemmas A.3 and A.1, having in mind that $c' \gg \alpha$ which shows that $\mathcal{G}$ is also $2c'$-joint. Let $T \boxtimes K_k$ be a member of $\mathcal{T}_{c'n,d}(k)$. Denote by $\tilde{T}$ the copy of T as described which we can find in $\mathcal{G}[S]$, and denote its vertex set by U. Pick a root $\tilde{r}$ of $\tilde{T}$ arbitrarily.

For each $v \in V(\tilde{T})$, define $S(v) = \{v_1, v_2, \dots, v_k\}$ if the distance between v and $\tilde{r}$ is even and $S(v) = \{v_{k+1}, v_{k+2}, \dots, v_{2k}\}$ if it is odd. As the vertices in $S(v)$ all belong to different partition classes

V_i, each $S(v)$ is a red clique in K_N, and note that it is also disjoint from all other red cliques $S(u)$ with $u \in V(\tilde{T})$. For every edge $uv \in E(\tilde{T})$, each edge of K_N incident to a vertex u' in $S(u)$ and another vertex v' in $S(v)$ is red, as u' and v' cannot be in the same partition class V_i. Therefore, the graph induced by $\bigcup_{v \in U} S(v)$ in the red subgraph of K_N contains a copy of $T \boxtimes K_k$. We now 'transfer' this copy to the red subgraph of $G[B(W)]$ coloured according to ψ. Notice that each edge in this copy of $T \boxtimes K_k$ is also an edge of $\mathcal{G}^3$, because every two vertices $v_i, v_j \in S(v)$ with $v \in V(\tilde{T})$ are at distance at most 2 in $\mathcal{G}$, and every two vertices $u_i \in S(u), v_j \in S(v)$ with $uv \in E(\tilde{T})$ are at distance at most 3 in $\mathcal{G}$.

By definition, for each $uv \in E(\mathcal{G}^3)$ such that $\varphi(uv)$ is red in K_N, all edges between $B(u)$ and $B(v)$ are present in G, comprising a complete bipartite graph G_{uv}, and there is no blue copy of $K_{s,s}$ in G_{uv}. By Lemma 3.14, the number of blue edges in G_{uv} is at most $(s-1)^s t^{2-1/s} + st + 1 \leqslant \frac{t^2}{16dk}$. Let $F := T \boxtimes K_k$ and let $F' \subseteq G$ be the union of all the red edges in G_{uv} for all $uv \in E(F)$. From Lemma A.6, it follows that F' contains a red copy of $T \boxtimes K_k$. Note that our choice of $T \in \mathcal{T}_{c'n,d}$ was arbitrary, so conditioned on the blue subgraph of K_N not containing all trees in $\{T' | T \in \mathcal{T}_{c'n,d}\}$, G contains a red copy of every graph in $\mathcal{T}_{c'n,d}(k)$. $\square$

ACKNOWLEDGEMENTS

We would like to thank Rajko Nenadov and Miloš Trujić for helpful comments and discussions, as well as the anonymous referees for their very useful feedback, which improved the paper considerably. This research was supported by SNSF Project 217926. Part of this research was conducted while Nemanja Draganić was at ETH Zürich, Switzerland, and partially supported by SNSF Grant 200021_196965. This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie Grant Agreement Number: 101034413. Part of this research was conducted while Kalina Petrova was at the Department of Computer Science, ETH Zürich, Switzerland, supported by SNSF Grant CRSII5 173721.

JOURNAL INFORMATION

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

ORCID

Nemanja Draganić https://orcid.org/0000-0002-1102-3449

REFERENCES

1. N. Alon, M. Capalbo, Y. Kohayakawa, V. Rödl, A. Ruciński, and E. Szemerédi, *Universality and tolerance*, Proceedings 41st Annual Symposium on Foundations of Computer Science, IEEE Computer Society, Redondo Beach, CA, USA, 2000, pp. 14–21.
2. N. Alon and J. H. Spencer, *The probabilistic method*, 4th ed., John Wiley & Sons, Hoboken, NJ, 2016.
3. J. Balogh, R. Morris, and W. Samotij, *Independent sets in hypergraphs*, J. Am. Math. Soc. **28** (2015), no. 3, 669–709.
4. J. Beck, *On size Ramsey number of paths, trees, and circuits. I*, J. Graph Theory **7** (1983), 115–129.
5. S. Berger, Y. Kohayakawa, G. S. Maesaka, T. Martins, W. Mendonça, G. O. Mota, and O. Parczyk, *The size-Ramsey number of powers of bounded degree trees*, J. Lond. Math. Soc. **103** (2021), no. 4, 1314–1332.

6. H. L. Bodlaender and D. M. Thilikos, *Treewidth for graphs with small chordality*, Discrete Appl. Math. **79** (1997), no. 1–3, 45–61.
7. V. Chvatál, V. Rödl, E. Szemerédi, and W. T. j. Trotter, *The Ramsey number of a graph with bounded maximum degree*, J. Comb. Theory, Ser. B **34** (1983), 239–243.
8. D. Clemens, M. Jenssen, Y. Kohayakawa, N. Morrison, G. O. Mota, D. Reding, and B. Roberts, *The size-Ramsey number of powers of paths*, J. Graph Theory **91** (2019), no. 3, 290–299.
9. D. Clemens, M. Miralaei, D. Reding, M. Schacht, and A. Taraz, *On the size-Ramsey number of grid graphs*, Combin. Probab. Comput. **30** (2021), no. 5, 670–685.
10. D. Conlon, *A new upper bound for diagonal Ramsey numbers*, Ann. Math. (2) **170** (2009), no. 2, 941–960.
11. D. Conlon, J. Fox, and Y. Wigderson, *Three early problems on size Ramsey numbers*, Combinatorica **43** (2023), no. 3, 743–768.
12. D. Conlon, W. T. Gowers, W. Samotij, and M. Schacht, *On the KŁR conjecture in random graphs*, Israel J. Math. **203** (2014), no. 1, 535–580.
13. D. Conlon, R. Nenadov, and M. Trujić, *On the size-Ramsey number of grids*, Combinatorics, Probability and Computing **32**, (2023), no. 6, 874–880.
14. D. Conlon, R. Nenadov, and M. Trujić, *The size-Ramsey number of cubic graphs*, Bull. Lond. Math. Soc. **54** (2022), no. 6, 2135–2150.
15. G. Ding and B. Oporowski, *Some results on tree decomposition of graphs*, J. Graph Theory **20** (1995), no. 4, 481–499.
16. N. Draganić, M. Krivelevich, and R. Nenadov, *The size-Ramsey number of short subdivisions*, Random Structures Algorithms **59** (2021), no. 1, 68–78.
17. N. Draganić, M. Krivelevich, and R. Nenadov, *Rolling backwards can move you forward: on embedding problems in sparse expanders*, Trans. Am. Math. Soc. **375** (2022), no. 7, 5195–5216.
18. P. Erdős, *Some remarks on the theory of graphs*, Bull. Am. Math. Soc. **53** (1947), 292–294.
19. P. Erdős, R. J. Faudree, C. C. Rousseau, and R. H. Schelp, *The size ramsey number*, Period. Math. Hungar. **9** (1978), no. 1–2, 145–161.
20. P. Erdős and G. Szekeres, *A combinatorial problem in geometry*, Compos. Math. **2** (1935), 463–470.
21. J. Friedman and N. Pippenger, *Expanding graphs contain all small trees*, Combinatorica **7** (1987), 71–76.
22. L. Gerencsér and A. Gyárfás, *On Ramsey-type problems*, Ann. Univ. Sci. Budap. Rolando Eötvös, Sect. Math. **10** (1967), 167–170.
23. J. Han, M. Jenssen, Y. Kohayakawa, G. O. Mota, and B. Roberts, *The multicolour size-Ramsey number of powers of paths*, J. Comb. Theory, Ser. B **145** (2020), 359–375.
24. J. Han, Y. Kohayakawa, S. Letzter, G. Oliveira Mota, and O. Parczyk, *The size-Ramsey number of 3-uniform tight paths*, Adv. Comb. **2021** (2021), 12. Id/No 5.
25. P. E. Haxell, Y. Kohayakawa, and T. Łuczak, *The induced size-Ramsey number of cycles*, Comb. Probab. Comput. **4** (1995), no. 3, 217–239.
26. N. Kamčev, A. Liebenau, D. R. Wood, and L. Yepremyan, *The size Ramsey number of graphs with bounded treewidth*, SIAM J. Discrete Math. **35** (2021), no. 1, 281–293.
27. Y. Kohayakawa, *Szemerédi's regularity lemma for sparse graphs*, Foundations of computational mathematics. Selected papers of a conference held at IMPA in Rio de Janeiro, Brazil, Springer, Berlin, 1997, pp. 216–230.
28. Y. Kohayakawa and V. Rödl, *Szemerédi's regularity lemma and quasi-randomness*, Recent advances in algorithms and combinatorics, Springer, New York, NY, 2003, pp. 289–351.
29. Y. Kohayakawa, V. Rödl, M. Schacht, and E. Szemerédi, *Sparse partition universal graphs for graphs of bounded degree*, Adv. Math. **226** (2011), no. 6, 5041–5065.
30. A. Kosowski, B. Li, N. Nisse, and K. Suchan, *k-chordal graphs: from cops and robber to compact routing via treewidth*, Algorithmica **72** (2015), no. 3, 758–777.
31. T. Kövári, V. T. Sós, and P. Turán, *On a problem of K. Zarankiewicz*, Colloq. Math. **3** (1954), 50–57.
32. S. Letzter, A. Pokrovskiy, and L. Yepremyan, *Size-Ramsey numbers of powers of hypergraph trees and long subdivisions*, arXiv:2103.01942, 2021.
33. C. McDiarmid, *On the method of bounded differences*, Surv. Comb. **141** (1989), no. 1, 148–188.
34. R. Nenadov, *Ramsey and universality properties of random graphs*, Ph.D. thesis, ETH Zürich, 2016.
35. R. Nenadov, *A new proof of the KŁR conjecture*, Adv. Math. **406** (2022), 108518.

36. N. Pippenger and J. Spencer, *Asymptotic behavior of the chromatic index for hypergraphs*, J. Combin. Theory Ser. A **51** (1989), no. 1, 24–42.
37. F. P. Ramsey, *On a problem of formal logic*, Proc. Lond. Math. Soc. (2) **30** (1929), 264–286.
38. V. Rödl and A. Ruciński, *Lower bounds on probability thresholds for Ramsey properties*, Combinatorics, Paul Erdős is eighty, vol. 1, János Bolyai Mathematical Society, Budapest, 1993, pp. 317–346.
39. V. Rödl and A. Ruciński, *Threshold functions for Ramsey properties*, J. Am. Math. Soc. **8** (1995), no. 4, 917–942.
40. V. Rödl and E. Szemerédi, *On size Ramsey numbers of graphs with bounded degree*, Combinatorica **20** (2000), no. 2, 257–262.
41. A. Sah, *Diagonal Ramsey via effective quasirandomness*, Duke Mathematical Journal **172**, (2023), no. 3, 545–567.
42. D. Saxton and A. Thomason, *Hypergraph containers*, Invent. Math. **201** (2015), no. 3, 925–992.
43. J. Spencer, *Ramsey's theorem: a new lower bound*, J. Comb. Theory, Ser. A **18** (1975), 108–115.
44. J. Spencer, *Asymptotic lower bounds for Ramsey functions*, Discrete Math. **20** (1978), 69–76.
45. K. Tikhomirov, *On bounded degree graphs with large size-Ramsey numbers*, Combinatorica **44** (2024), no. 1, 9–14.
46. P. Turán, *On an external problem in graph theory*, Mat. Fiz. Lapok **48** (1941), 436–452.
47. R. M. Wilson, *An existence theory for pairwise balanced designs, III: Proof of the existence conjectures*, J. Combin. Theory Ser. A **18** (1975), no. 1, 71–79.
48. D. R. Wood, *On tree-partition-width*, European J. Combin. **30** (2009), no. 5, 1245–1253.