

Arboricity and spanning-tree packing in random graphs*

Pu Gao[†]

Xavier Pérez-Giménez[‡]

University of Toronto
pu.gao@utoronto.ca

University of Waterloo
xperez@uwaterloo.ca

Cristiane M. Sato

University of Waterloo
and
Center of Mathematics, Computing and Cognition
Universidade Federal do ABC
c.sato@ufabc.edu.br

Abstract

We study the arboricity A and the maximum number T of edge-disjoint spanning trees of the classical random graph $\mathcal{G}(n, p)$. For all $p(n) \in [0, 1]$, we show that, with high probability, T is precisely the minimum between δ and $\lfloor m/(n-1) \rfloor$, where δ is the smallest degree of the graph and m denotes the number of edges. Moreover, we explicitly determine a sharp threshold value for p such that: above this threshold, T equals $\lfloor m/(n-1) \rfloor$ and A equals $\lceil m/(n-1) \rceil$; and below this threshold, T equals δ , and we give a two-value concentration result for the arboricity A in that range. Finally, we include a stronger version of these results in the context of the random graph process where the edges are sequentially added one by one. A direct application of our result gives a sharp threshold for the maximum load being at most k in the two-choice load balancing problem, where $k \rightarrow \infty$.

1 Introduction

STP number and arboricity: The spanning-tree packing (STP) number of a graph is the maximum number of edge-disjoint spanning trees it contains. Computing this parameter is a

*An extended abstract of this article has appeared in the Proceedings of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'14). See [20].

[†]This author has been supported by NSERC PDF.

[‡]Partially supported by FORMALISM TIN2007-66523.

very classical problem in combinatorial optimization. One of the earliest results on the STP number is a min-max relation proved by Tutte [42] and Nash-Williams [34]: the STP number of a graph is the minimum value, ranging over all partitions \mathcal{P} of the vertex set, of the ratio (rounded down) between the number of edges across \mathcal{P} (i.e. edges with ends lying in different classes of \mathcal{P}) and $|\mathcal{P}| - 1$. This characterisation has important consequences in computer science, where the STP number has been used as a measure of network vulnerability in case of attack or edge failure (see [23, 10]). Intuitively speaking, it provides information about the number of edges that must be destroyed in a connected network in order to create a given number of new components. In addition, finding edge-disjoint spanning trees in a graph is relevant to the design of efficient and robust communication protocols (see e.g. a seminal article by Itai and Rodeh [25]). There are two obvious upper bounds on the STP number of a graph with n vertices: the minimum degree, since each spanning tree would need to use at least one edge incident to each vertex; and the number of edges divided by $n - 1$, since each spanning tree has exactly $n - 1$ edges. For further information on the STP number, we refer the reader to a survey by Palmer [36] on this topic.

Another closely related graph parameter that has been widely studied is the arboricity of a graph — i.e. the minimum number of subforests needed to cover all of its edges. A trivial lower bound on the arboricity of a graph with n vertices is the number of edges divided by $n - 1$, since we cannot do better than covering all the edges with a set of edge-disjoint spanning trees. Nash-Williams [35] also provided a min-max relation for the arboricity of a graph, which yields a natural interpretation of arboricity as a measure of density of the subgraphs of a graph. This makes arboricity a useful notion in computer science, since the problem of determining the existence of dense subgraphs in large graphs is relevant to many applications in real world domains like social networking or internet computing. In fact, finding such dense subgraphs and other related problems can often be efficiently solved in linear time for any class of graphs with bounded arboricity (see [8, 22]). This includes important families of graphs such as all minor-closed classes (e.g. planar graphs and graphs with bounded treewidth) and random graphs generated by the preferential attachment model. Another relevant feature of arboricity is its intimate connection with the k -orientability of a graph and its application to certain load-balancing problems (see the discussion below on this topic).

Finding the STP number and the arboricity of a given graph are among the most successful applications of matroids in combinatorial optimization. Both problems can be formulated as matroid union problems and thus can be solved in polynomial time. For more details, see [40, Chapter 51].

Arboricity, k -orientability and load balancing: The k -orientability problem is that of determining whether a graph admits an orientation of its edges so that each vertex has indegree at most k . As established by Hakimi [24], the k -orientability of a graph is determined by the density of the densest subgraph, which brings a strong connection between k -orientability and arboricity. On the other hand, the k -orientability problem of a (random) graph on n vertices and m edges is equivalent to determining whether the maximum load is at most k in the following load-balancing scenario: m balls (jobs) are assigned to n bins (machines) in a way that each ball must pick between two randomly chosen bins, and we

wish to minimise the load of the bins by allowing at most k balls in each bin. This subject has received a lot of attention, since the seminal result by Azar, Broder, Karlin and Upfal [1] on the multiple-choice paradigm of load balancing. They consider a related online setting in which n balls are sequentially thrown into n bins, and each ball is allowed to choose between h given random bins; and prove that the maximum load among all bins can be significantly reduced if each ball has $h \geq 2$ choices rather than one. Research on the power of choice in load balancing has since been very fruitful. We refer readers to the survey by Mitzenmacher, Richa and Sitaraman [31] for a detailed account of the topic. Cain, Sanders and Wormald [5] and Fernholz and Ramachandran [14] simultaneously determined a threshold for a random graph (corresponding to $h = 2$) to be k -orientable. The random hypergraph case ($h \geq 2$) was studied by a few authors due to its applications in Cuckoo hashing [12] and disk scheduling [5]. See [12, 18, 16] for $h \geq 2$ and $k = 1$, and [15] for $h \geq 2$ and general k . A more general case where each ball can take $1 \leq w \leq h$ copies and be assigned to w distinct bins was studied by the first author and Wormald [21] for sufficiently large but constant k . Lelarge closed the gap for small values of k in [29]. For $m/n \geq C \log n$, where $C > 0$ is sufficiently large, Czumaj, Riley and Scheideler [11] proved that $\mathcal{G}(n, m)$ (the random graph on n vertices and m edges with uniform distribution) is a.a.s. $\lceil m/n \rceil$ -orientable. The only case left open for the k -orientability problem is $m/n = O(\log n)$ with $m/n \rightarrow \infty$. Sanders, Egner and Korst [39] proved that $\mathcal{G}(n, m)$ is $(\lceil m/n \rceil + 1)$ -orientable (and of course not $\lfloor (m-1)/n \rfloor$ -orientable), implying that the k -orientability threshold for $\mathcal{G}(n, m)$, for $m/n \rightarrow \infty$, is kn with a window of size $O(1/k) = o(1)$ inside which we cannot determine whether it is k -orientable or not. In this paper, we will squeeze the window of uncertainty to a much smaller size. One conceivable difference between this last case and the one covered in [11] is the following. If $m/n \geq C \log n$ for sufficiently large C , all degrees (including the minimum and maximum degree) of $\mathcal{G}(n, m)$ are very close to the average degree. However, if m/n grows slowly as n grows, the minimum degree of $\mathcal{G}(n, m)$ is likely to be 0 while its average degree is large. Intuitively, load should balance “better” for a growing large parameter k rather than a bounded one, but the previous proofs in [5, 14] dealing with constant k do not generalise to $k \rightarrow \infty$. An online version of load balancing was studied by Berenbrink, Czumaj, Steger and Vöcking [2], where the case $m = \omega(n)$ was particularly investigated.

STP number and arboricity of a random graph: It is then relevant to study the behaviour of the STP number and the arboricity for the classical random graph $\mathcal{G}(n, p)$, in which the vertex set is $[n]$ and each of the $\binom{n}{2}$ possible edges is included independently with probability p (where $p = p(n)$ is a function of n). It is a well-known fact that, for $p = (\log n - \omega(1))/n$, the random graph $\mathcal{G}(n, p)$ is a.a.s.¹ disconnected (see e.g. Theorem 7.3 in [3]), and hence the STP number is zero. Palmer and Spencer [37] showed that a.a.s. the STP number of $\mathcal{G}(n, p)$ equals the minimum degree whenever this has constant value, which typically happens when p is around $(\log n + O(\log \log n))/n$. In fact, they proved a stronger hitting-time result in the context of the evolution of $\mathcal{G}(n, p)$ when p moves gradually from 0 to 1, and showed that a.a.s. the precise time when the minimum degree first becomes k (for constant k) coincides with the time when k edge-disjoint spanning trees first appear.

¹We say that a sequence of events H_n holds *asymptotically almost surely* (a.a.s.) if $\lim_{n \rightarrow \infty} \Pr(H_n) = 1$.

Moreover, Catlin, Chen and Palmer [6] studied the denser case of $p = C(\log n/n)^{1/3}$, where $C > 0$ is a sufficiently large constant, and determined the STP number and the arboricity of $\mathcal{G}(n, p)$ to be a.a.s. equal to $\lfloor m/(n-1) \rfloor$ and $\lceil m/(n-1) \rceil$, respectively, where m denotes the number of edges. In a recent unpublished manuscript, Chen, Li and Lian [7] proved that, for any $(\log n + \omega(1))/n \leq p \leq 1.1 \log n/n$, a.a.s. the STP number of $\mathcal{G}(n, p)$ equals the minimum degree. They also observed that this property a.a.s. does not hold for $p \geq 51 \log n/n$, and posed the question of what is the smallest value of p such that the STP number of $\mathcal{G}(n, p)$ differs from the minimum degree.

Outline of our contribution: In this paper we strengthen the previous work, and characterise the STP number and the arboricity of $\mathcal{G}(n, p)$. A direct application of our results gives a sharp threshold for the maximum load being at most k in the two-choice load balancing problem, where $k \rightarrow \infty$. Our methods rely on a successful combination of some combinatorial optimisation tools together with several probabilistic techniques, and may be hopefully extended in future research to address other relevant related problems in the area. We first prove that for all $p(n) \in [0, 1]$, the STP number is a.a.s. the minimum between δ and $\lfloor m/(n-1) \rfloor$, where δ and m respectively denote the minimum degree and the number of edges of $\mathcal{G}(n, p)$ (see Theorem 1). Note that the quantities δ and $\lfloor m/(n-1) \rfloor$ above correspond to the two trivial upper bounds observed earlier for arbitrary graphs, so this implies that we can a.a.s. find a best-possible number of edge-disjoint spanning trees in $\mathcal{G}(n, p)$. Our argument uses several properties of $\mathcal{G}(n, p)$ in order to bound the number of crossing edges between subsets of vertices with certain restrictions, and then applies the characterisation of the STP number by Tutte and Nash-Williams restated as Theorem 8. Moreover, we determine the ranges of p for which the STP number takes each of these two values: δ and $\lfloor m/(n-1) \rfloor$. In spite of the fact that the property $\{\delta \leq \lfloor m/(n-1) \rfloor\}$ is not necessarily monotonic with respect to p , we show that it has a sharp threshold at $p \sim \beta \log n/n$, where $\beta \approx 6.51778$ is a constant defined in Theorem 2. Below this threshold, the STP number of $\mathcal{G}(n, p)$ is a.a.s. equal to δ ; and above the threshold it is a.a.s. $\lfloor m/(n-1) \rfloor$. In particular, this settles the question raised by Chen, Li and Lian [7]. We also include a stronger version of these results in the context of the random graph process in which p gradually grows from 0 to 1 (or, similarly, the edges are added one by one). This provides a full characterisation of the STP number that holds a.a.s. simultaneously during the whole random graph process (see Theorem 3). The argument combines a more accurate version of the same ideas used in the analysis of the STP number of $\mathcal{G}(n, p)$ together with multiple couplings of $\mathcal{G}(n, p)$ at different values of p . In addition, the article contains several results about the arboricity of $\mathcal{G}(n, p)$. As an almost direct application of our result on the STP number, for p above the threshold $\beta \log n/n$, we determine the arboricity of $\mathcal{G}(n, p)$ to be a.a.s. equal to $\lceil m/(n-1) \rceil$. This significantly extends the range of p in the result by Catlin, Chen and Palmer [6]. We further prove that for all other values of p , the arboricity of $\mathcal{G}(n, p)$ is concentrated on at most two values (see Theorem 4). In order to prove this for the case $pn \rightarrow \infty$, we add $o(n)$ edges to $\mathcal{G}(n, p)$ in a convenient way that guarantees a full decomposition of the resulting graph into edge-disjoint spanning trees. This construction builds upon some of the ideas that we previously use to study the STP number. The case $pn = O(1)$ uses different proof techniques which rely on the structure of the k -core of $\mathcal{G}(n, p)$ together with the

Nash-Williams characterisation of arboricity restated as Theorem 9. Finally, some of the aforementioned results on the arboricity are also given in the more precise context of the random graph process (see Theorem 5 and Corollary 6), similarly as we did for the STP number. As a direct corollary of our result on arboricity, we determine a sharp threshold for the k -orientability of $\mathcal{G}(n, m)$ where $k \rightarrow \infty$ (see Theorem 7). This successfully settles the load-balancing problem in the scenario where $m = \omega(n)$ balls are allocated into n bins, and each ball has two choices, uniformly at random chosen from $[n]$. We prove that in this case, balls can be allocated so that most bins receive an almost even load. This extends the result by Cain, Sanders and Wormald [5] and Fernholz and Ramachandran [14] to the case of $k \rightarrow \infty$.

Related work: Recently, Swamy, Bhashyam, Sundaresan and Viswanath [41] studied several versions of network flow problems on random graphs. In particular, they obtained an almost sure convergence result which implies that the STP of $\mathcal{G}(n, p)$ is a.a.s. $(1 + o(1))pn/2$ as long as $p = \omega(\sqrt{\log n/n})$. The behaviour of the STP number and the arboricity has been also studied in other models of random graphs. Frieze and Łuczak [19] considered the random directed graph in which each vertex chooses k out-neighbours uniformly at random, with fixed k . This graph has k disjoint spanning trees with probability going to 1 (where the orientation of the arcs is ignored). Cooper [9] upper-bounded the arboricity of $\mathcal{G}(n, p)$ for $p = \Omega(1/n)$ and used that to obtain an upper bound of the thickness of $\mathcal{G}(n, p)$: the minimum number of edge-disjoint planar graphs which cover all edges in $\mathcal{G}(n, p)$.

Some variants of arboricity have also been studied. The linear arboricity of a graph is the minimum number of forests consisting only of paths needed to cover all edges of the graph. This parameter was studied by McDiarmid and Reed [30] for random regular graphs.

2 Main results

Let $\mathcal{G}(n, p)$ denote the random graph with vertex set $[n]$ such that each possible edge in $\{\{u, v\} : u, v \in [n], u \neq v\}$ is included independently with probability p . In this article, we regard p as a function of n , and consider asymptotic statements as $n \rightarrow \infty$. Given a sequence of events $(H_n)_{n \in \mathbb{N}}$, we say that H_n happens *asymptotically almost surely* (a.a.s.) if $\Pr(H_n) \rightarrow 1$ as $n \rightarrow \infty$. Given real sequences a_n and b_n (possibly taking negative values), we write: $a_n = O(b_n)$ if there is a constant $C > 0$ such that $|a_n| \leq C|b_n|$ for all n ; $a_n = o(b_n)$ if eventually $b_n \neq 0$ and $\lim_{n \rightarrow \infty} a_n/b_n = 0$; $a_n = \Omega(b_n)$ if eventually $a_n \geq 0$ and $b_n = O(a_n)$; $a_n = \omega(b_n)$ if eventually $a_n \geq 0$ and $b_n = o(a_n)$; $a_n = \Theta(b_n)$ if eventually $a_n \geq 0$, $a_n = O(b_n)$ and $a_n = \Omega(b_n)$; and finally $a_n \sim b_n$ if $a_n = (1 + o(1))b_n$. In particular, in this paper, all constants involved in these notations do not depend on the p under discussion. For instance, if we have $a_n = \Omega(b_n)$, where b_n may be an expression involving $p = p(n)$, then it means that there are constants $C > 0$ and n_0 (both independent of p), such that $a_n \geq C|b_n|$ uniformly for all $n \geq n_0$ and for all p in the range under discussion. We also assume in this paper that n is greater than some absolute constant (e.g. in a few places we require $n \geq 2$), which does not depend on any variables under discussion. We often omit this assumption in the statements of lemmas and theorems.

Given a graph G , let $m(G)$ be the number of edges of G and let $\delta(G)$ denote the minimum

degree of G . Let $T(G)$ be the STP number of G — i.e. the maximum number of edge-disjoint spanning trees in G (possibly 0 if G is disconnected).

Theorem 1. *For every $p = p(n) \in [0, 1]$, we have that a.a.s.*

$$T(\mathcal{G}(n, p)) = \min \left\{ \delta(\mathcal{G}(n, p)), \left\lfloor \frac{m(\mathcal{G}(n, p))}{n-1} \right\rfloor \right\}.$$

Theorem 2. *Let $\beta = 2/\log(e/2) \approx 6.51778$. Then*

- (i) *if $p = \frac{\beta(\log n - \log \log n/2) - \omega(1)}{n-1}$, then a.a.s. $\delta(\mathcal{G}(n, p)) \leq \left\lfloor \frac{m(\mathcal{G}(n, p))}{n-1} \right\rfloor$ and so $T(\mathcal{G}(n, p)) = \delta(\mathcal{G}(n, p))$;*
- (ii) *if $p = \frac{\beta(\log n - \log \log n/2) + \omega(1)}{n-1}$, then a.a.s. $\delta(\mathcal{G}(n, p)) > \left\lfloor \frac{m(\mathcal{G}(n, p))}{n-1} \right\rfloor$ and so $T(\mathcal{G}(n, p)) = \left\lfloor \frac{m(\mathcal{G}(n, p))}{n-1} \right\rfloor$.*

Consider the random graph process $G_0, G_1, \dots, G_{\binom{n}{2}}$ defined as follows: for each $m = 0, \dots, \binom{n}{2}$, G_m is a graph with vertex set $[n]$; the graph G_0 has no edges; and, for each $1 \leq m \leq \binom{n}{2}$, the graph G_m is obtained by adding one new edge to G_{m-1} chosen uniformly at random among the edges not present in G_{m-1} . Equivalently, we can choose uniformly at random a permutation $(e_1, \dots, e_{\binom{n}{2}})$ of the edges of the complete graph with vertex set $[n]$, and define each G_m to be the graph on vertex set $[n]$ and edges e_1, \dots, e_m .

The following theorem is a strengthening of Theorems 1 and 2 in the context of the random graph process just described. Note that the a.a.s. statements in Theorem 3 refer to events that hold *simultaneously* for all m in a certain specified range, as we add edges one by one.

Theorem 3. *Let $\beta = 2/\log(e/2) \approx 6.51778$. The following holds in the random graph process $G_0, G_1, \dots, G_{\binom{n}{2}}$.*

- (i) *A.a.s. $T(G_m) = \min \{ \delta(G_m), \lfloor m/(n-1) \rfloor \}$ for every $0 \leq m \leq \binom{n}{2}$.*
- (ii) *Moreover, for any constant $\epsilon > 0$, a.a.s.*
 - $\delta(G_m) \leq \lfloor m/(n-1) \rfloor$ for every $0 \leq m \leq \frac{(1-\epsilon)\beta}{2} n \log n$, and
 - $\delta(G_m) > \lfloor m/(n-1) \rfloor$ for every $\frac{(1+\epsilon)\beta}{2} n \log n \leq m \leq \binom{n}{2}$.

Remark. We will actually prove a stronger result than Theorem 3 (ii). See Theorem 27 in Section 7.

For any graph G , let $A(G)$ denote the arboricity of G , i.e. the minimum number of subforests of G which cover the whole edge set of G . Let $\mathbb{R}_{\geq 1}$ denote the set of positive real numbers that are greater than one.

Theorem 4. *Let $\beta = 2/\log(e/2) \approx 6.51778$.*

- (i) *For all $p = \frac{\beta(\log n - \log \log n/2) + \omega(1)}{n-1}$, a.a.s. $A(\mathcal{G}(n, p)) = \left\lceil \frac{m(\mathcal{G}(n, p))}{n-1} \right\rceil$; for all $p = \omega(1/n)$, a.a.s. $A(\mathcal{G}(n, p)) \in \left\{ \left\lceil \frac{m(\mathcal{G}(n, p))}{n-1} \right\rceil, \left\lceil \frac{m(\mathcal{G}(n, p))}{n-1} \right\rceil + 1 \right\}$;*

(ii) For all $p = \Theta(1/n)$, there exist a constant $C > 0$ and $k = \Theta(1)$ (both depending on p) such that a.a.s. $A(\mathcal{G}(n, p)) \geq (1 + C)pn/2$ and $A(\mathcal{G}(n, p)) \in \{k, k + 1\}$. Moreover, there is a countable set of real numbers \mathcal{J} such that, for all $c \in \mathbb{R}_{\geq 1} \setminus \mathcal{J}$, there is an integer $k = k(c)$ such that a.a.s. $A(\mathcal{G}(n, p)) = k$, where $p \sim c/n$.

(iii) If $p = o(1/n)$, then a.a.s. $A(\mathcal{G}(n, p)) \leq 1$.

Remark. (a) Note that Theorem 4 (i) is a simple corollary of Theorem 5 below. Indeed, for most values of $p = \omega(1/n)$, we can do even better than the two-value concentration result stated above and a.a.s. determine the exact value of the arboricity (cf. the remark that follows Theorem 5).

(b) It follows from Theorem 4 that, for all $p = \omega(1/n)$, the arboricity of $\mathcal{G}(n, p)$ is asymptotic to $pn/2$, whereas this property fails for $p = O(1/n)$.

(c) There is no one-point concentration in Theorem 4 (ii) for $p \sim c/n$ where $c < 1$. In fact, in this scenario, it is easy to see that a.a.s. $A(\mathcal{G}(n, p)) = 1 + I$, where I denotes the indicator random variable of the event that $\mathcal{G}(n, p)$ contains a cycle. It is well-known that $\lim_{n \rightarrow \infty} \Pr(I = 1) = \alpha(c)$ for some constant $0 < \alpha(c) < 1$ (with $\lim_{c \rightarrow 0} \alpha(c) = 0$ and $\lim_{c \rightarrow 1} \alpha(c) = 1$). Therefore, if $0 < c < 1$, $A(\mathcal{G}(n, p))$ has positive measure on both 1 and 2 as $n \rightarrow \infty$.

Theorem 5. Let $\beta = 2/\log(e/2) \approx 6.51778$. The following holds in the random graph process $G_0, G_1, \dots, G_{\binom{n}{2}}$.

(i) Let m_0 be any function of n such that $m_0/n \rightarrow \infty$ and let $\epsilon > 0$ be any constant. Then, a.a.s. simultaneously for all $m \geq m_0$ such that $\delta(G_m) \leq m/(n-1)$,

$$\left\lceil \frac{m + \phi_1}{n-1} \right\rceil \leq A(G_m) \leq \left\lceil \frac{m + \phi_2}{n-1} \right\rceil, \quad (1)$$

where $\phi_1 = n/\exp\left(\frac{(1+\epsilon)2m}{\beta n}\right) = o(n)$ and $\phi_2 = n/\exp\left(\frac{(1-\epsilon)2m}{\beta n}\right) = o(n)$. In particular, a.a.s. $A(G_m) \in \{\lceil \frac{m}{n-1} \rceil, \lceil \frac{m}{n-1} \rceil + 1\}$ for all m in that range.

(ii) Moreover, a.a.s. simultaneously for every m such that $\delta(G_m) \geq m/(n-1)$ we have

$$A(G_m) = \lceil m/(n-1) \rceil.$$

Remark. Given positive integers a and b , let $R(a, b) = a - b\lfloor a/b \rfloor$ denote the remainder of a divided by b . Then (1) in Theorem 5(i) implies $A(G_m) = \lceil \frac{m}{n-1} \rceil$ for those m such that $0 < R(m, n-1) \leq n-1 - \phi_2$ (which is the case for most values of m), and $A(G_m) = \lceil \frac{m}{n-1} \rceil + 1$ if $R(m, n-1) = 0$ or $R(m, n-1) > n-1 - \phi_1$. For those few remaining values of m such that $n-1 - \phi_2 < R(m, n-1) \leq n-1 - \phi_1$ we can only say $A(G_m) \in \{\lceil \frac{m}{n-1} \rceil, \lceil \frac{m}{n-1} \rceil + 1\}$.

Corollary 6. Let $m_{A=i}$ denote the minimum m such that $A(G_m)$ becomes i in the random graph process $G_0, G_1, \dots, G_{\binom{n}{2}}$. Let i_0 be any function of n such that $i_0 \rightarrow \infty$ and $\epsilon > 0$ be a constant. Then a.a.s.

(i) for every $i_0 \leq i \leq (1 - \epsilon)\beta \log n/2$,

$$(i - 1)(n - 1) - \phi_2 < m_{A=i} < (i - 1)(n - 1) - \phi_1,$$

where $\phi_1 = n / \exp\left(\frac{2(1+\epsilon)}{\beta}i\right) = o(n)$ and $\phi_2 = n / \exp\left(\frac{2(1-\epsilon)}{\beta}i\right) = o(n)$; and

(ii) for every $(1 + \epsilon)\beta \log n/2 \leq i \leq n/2$,

$$m_{A=i} = (i - 1)(n - 1) + 1.$$

Remark. Indeed, as shown in the proof in Section 8, $m_{A=i} > (i - 1)(n - 1) - \phi_2$ holds for all $i_0 \leq i \leq (1 + \epsilon)\beta \log n/2$.

Recall that a graph G is k -orientable if all edges of G can be oriented so that the maximum indegree of the oriented graph is at most k . It was shown by Hakimi [24] that G is k -orientable if and only if it contains no subgraph with average degree more than $2k$. Trivially, this implies that no graph on n vertices with more than kn edges can be k -oriented. Moreover, by the previous corollary together with a result by Nash-Williams (see Theorem 9), we obtain the following theorem, which characterises the k -orientability of $\mathcal{G}(n, m)$ (i.e. the uniform random graph on n vertices and m edges) for $k \rightarrow \infty$.

Theorem 7. Let $\beta = 2/\log(e/2) \approx 6.51778$. Let f be any function of n that goes to infinity as $n \rightarrow \infty$. Then for every integer $k \geq f$ and any $\epsilon > 0$, as $n \rightarrow \infty$,

$$\Pr(\mathcal{G}(n, m) \text{ is } k\text{-orientable}) \rightarrow \begin{cases} 1 & \text{if } m \leq k(n - 1) - \phi \\ 0 & \text{if } m \geq kn + 1, \end{cases}$$

where $\phi = 0$ if $k \geq \frac{1+\epsilon}{2}\beta \log n$ and $\phi = n / \exp\left(\frac{2(1-\epsilon)}{\beta}k\right) = o(n)$ if $f \leq k < \frac{1+\epsilon}{2}\beta \log n$.

Remark. In particular, Theorem 7 implies that the property of being k -orientable has a sharp threshold in $\mathcal{G}(n, m)$ at $m \sim kn$, since we are assuming $k \rightarrow \infty$. By a closer inspection of Corollary 6, we may obtain more accurate bounds on the critical m for k -orientability by distinguishing each of the two cases $f \leq k \leq \frac{1-\epsilon}{2}\beta \log n$ and $k \geq \frac{1+\epsilon}{2}\beta \log n$. Moreover, we recall that the k -orientability of $\mathcal{G}(n, m)$ can be interpreted in terms of a load-balancing problem in a setting where $m = \omega(n)$ jobs (edges) are assigned to n machines (vertices), in a way that each job is allocated to one machine, selected from two randomly given ones. In this context, we want to minimise the maximum load (number of jobs) received by any machine. Observe that Theorem 7 implies that there exists a load distribution such that the maximum load is around m/n , and almost all machines receive a load very close to the maximum load.

The paper is organised as follows. We first introduce some basic tools in Section 3, including two classic theorems by Tutte and Nash-Williams that characterise the spanning tree-packing number and the arboricity of a graph. There, we also prove two deterministic results (Propositions 10 and 11) that are central in our argument. They give a list of conditions under which the STP number equals the minimum between δ and $\lfloor m/(n - 1) \rfloor$. In Section 4, we prove several lemmas about properties of $\mathcal{G}(n, p)$. These lemmas will be

used in conjunction with the two aforementioned deterministic propositions to derive most of the main results in the paper. Finally, we prove Theorem 1 in Section 5, Theorem 2 in Section 6, Theorem 3 in Section 7, Theorem 5 in Section 8 and Theorem 4 in Section 9. In Section 10 we discuss the extension of our results to random uniform hypergraphs.

3 Deterministic tools

In this section, we introduce some basic tools that lie in the core of our argument. Given a graph G , let $V(G)$ denote the vertex set of G and let $E(G)$ denote the edge set of G . Recall that $m(G)$ is the number of edges, and $\delta(G)$ is the minimum degree of G . If $|V(G)| \geq 2$, define $\bar{d}(G) := 2m(G)/(|V(G)| - 1)$. Note that $\bar{d}(G)$ differs from the average degree of G by a small factor of $|V(G)|/(|V(G)| - 1)$. Also, let $t(G) = \min\{\delta(G), \bar{d}(G)/2\}$.

We first restate two well-known results by Tutte and Nash-Williams that provide a useful characterisation of the STP number and the arboricity of a graph G . For any partition \mathcal{P} of the vertex set $V(G)$ of a graph G , let $m(\mathcal{P})$ denote the number of edges in G with ends in distinct parts of \mathcal{P} .

Theorem 8 (Tutte [42] and Nash-Williams [34]). *Let G be a graph and t a positive integer. Then G contains t edge-disjoint spanning trees if and only if, for every partition \mathcal{P} of the vertex set of G such that every class is non-empty,*

$$m(\mathcal{P}) \geq t(|\mathcal{P}| - 1). \quad (2)$$

For any $S \subseteq V(G)$, let $E[S]$ denote the set of edges of G with both ends in S .

Theorem 9 (Nash-Williams [35]). *Let G be a graph and t a positive integer. Then the edge set of G can be covered by t forests if and only if, for every nonempty subset S of vertices of G ,*

$$|E[S]| \leq t(|S| - 1). \quad (3)$$

The next two propositions play a central role in this paper. They make use of Theorem 8 to determine the STP number of any well-behaved graph satisfying certain conditions. For $\epsilon > 0$, we say that a vertex of G is ϵ -light if its degree is at most $\delta(G) + \epsilon\bar{d}(G)$.

Proposition 10. *Let $G = G_n$ be a graph on vertex set $[n]$. Let $\delta := \delta(G)$ and let $\bar{d} := \bar{d}(G)$. Suppose that $\bar{d} \rightarrow \infty$ as $n \rightarrow \infty$ and that there exist constants $\epsilon, \zeta, \eta > 0$ such that the following hold, for all sufficiently large n .*

- (a) *The minimum degree δ is at most $(\epsilon/4)\bar{d}$; there is no pair of adjacent ϵ -light vertices; and all vertices of G have at most one ϵ -light neighbour.*
- (b) *No set of size $s < \zeta n$ induces more than $(\epsilon/4)\bar{d}s$ edges.*
- (c) *For all disjoint $S, S' \subseteq [n]$ with $|S| \geq |S'| \geq \zeta n$, we have that $|E(S, S')| \geq \eta\bar{d}n$.*

Then eventually $T(G) = \delta$.

Proposition 11. *Let $G = G_n$ be a graph on $[n]$. Let $\delta := \delta(G)$ and $\bar{d} := \bar{d}(G)$, and suppose that $\bar{d} \rightarrow \infty$ as $n \rightarrow \infty$. Let $t = \min\{\delta, \bar{d}/2\}$. Suppose moreover that there exist constants $0 < \epsilon, \eta, \zeta \leq 1$ such that the following hold, for sufficiently large n .*

- (a') *Either we have that $\delta > \frac{(1+\epsilon)\bar{d}}{2}$; or there are no adjacent ϵ -light vertices and each vertex of G is adjacent to at most one ϵ -light vertex.*
- (b') *For all $S \subseteq V(G)$, with $|S| \geq \zeta n$, we have that $d(S) \geq \bar{d}(1 - o(1))$, where $d(S)$ denotes the sum of degrees of vertices in S divided by $|S|$.*
- (c') *For all disjoint $S, S' \subseteq V(G)$ with $|S| \geq |S'| \geq \zeta n$, we have that $|E(S, S')| \geq \eta \bar{d} |S| |S'| / n$.*
- (d') *For all $\emptyset \subsetneq S \subsetneq V(G)$, we have that $|E(S, \bar{S})| \geq t$.*
- (e') *No set of size $s < \zeta n$ induces more than $(\epsilon/4)ts$ edges.*

Then eventually $T(G) = \lfloor t \rfloor$.

Propositions 10 and 11 will be used to determine the STP number and the arboricity of $\mathcal{G}(n, p)$ (see the arguments leading to the proofs of Theorems 1, 3, 4 and 5). Basically, according to the range of p , we will show that $\mathcal{G}(n, p)$ (or some modification of $\mathcal{G}(n, p)$) satisfies the conditions in Proposition 10 or Proposition 11 with sufficiently high probability. Proposition 10 is applied when the minimum degree is relatively small compared to \bar{d} , whereas otherwise Proposition 11 is used instead. Thus, we need a good estimation of $\delta(\mathcal{G}(n, p))$; together with several graph-expansion-related properties, as required by conditions (b), (c), (c'), (d') and (e'); and also some properties about the ϵ -light vertices addressed in conditions (a) and (a'). In the following section, we derive bounds on the probability that these properties hold in $\mathcal{G}(n, p)$ for some relevant ranges of p .

Proof of Proposition 10. We will show that every partition of the vertices of G satisfies (2) with $t = \delta$, and thus G has δ edge-disjoint spanning trees by Theorem 8.

Let \mathcal{P} be a partition of $V(G)$. Parts of size one are denoted *singletons*, and singletons consisting of one ϵ -light vertex are called *ϵ -light singletons*. We may assume that

$$\text{every part with size at least 2 has one vertex that is not } \epsilon\text{-light.} \quad (4)$$

This is because, given a part of size at least 2 and with only ϵ -light vertices, we can refine the partition by turning each vertex in this part into a singleton, and this increases the number of parts without increasing the number of edges with ends in distinct parts by Condition (a).

Let \mathcal{K}_1 denote the set of ϵ -light singletons, let \mathcal{K}_2 denote the set of parts of size between 2 and ζn together with the singletons that are not ϵ -light, and let \mathcal{K}_3 denote the set of other parts. For $i = 1, 2, 3$, let $k_i = |\mathcal{K}_i|$. Then, $|\mathcal{P}| = k_1 + k_2 + k_3$.

By Condition (a), no ϵ -light vertices are adjacent. Thus, the number of edges incident with a vertex in \mathcal{K}_1 is at least δk_1 . Suppose \mathcal{K}_2 is non-empty and suppose S is a part in \mathcal{K}_2 , and let r be the number of vertices in S that are not ϵ -light. By the assumption in (4) and the definition of \mathcal{K}_2 , we must have $1 \leq r \leq \zeta n$. The number of edges between these

r vertices is at most $(\epsilon/4)\bar{d}r$ by Condition (b). Since these vertices are not ϵ -light, each of them has degree at least $\delta + \epsilon\bar{d}$. By Condition (a), each of these vertices is adjacent to at most one ϵ -light vertex. Thus,

$$|E(S, \bar{S} \setminus \mathcal{K}_1)| \geq r(\delta + \epsilon\bar{d} - 1) - 2(\epsilon/4)\bar{d}r \geq \delta + (\epsilon/4)\bar{d},$$

where the term -1 in the first inequality accounts for a possible ϵ -light neighbour of each one of these r vertices and we use the fact that $d \rightarrow \infty$. Thus, the number of edges in the partition \mathcal{P} is at least

$$m(\mathcal{P}) \geq \delta k_1 + \frac{\delta + (\epsilon/4)\bar{d}}{2} k_2 \geq \delta(k_1 + k_2), \quad (5)$$

as $\delta \leq (\epsilon/4)\bar{d}$ by Condition (a). If $k_3 \leq 1$, this already shows that $m(\mathcal{P}) \geq \delta(|\mathcal{P}| - 1)$. Otherwise, we have $2 \leq k_3 \leq 1/\zeta$, and the number of edges between any two parts of \mathcal{K}_3 is at least $\eta\bar{d}n$ by Condition (c). We can add these additional edges to (5) and obtain

$$m(\mathcal{P}) \geq \delta(k_1 + k_2) + \eta\bar{d}n \geq \delta(k_1 + k_2) + (\epsilon/4)\bar{d}k_3 \geq \delta|\mathcal{P}|, \quad (6)$$

since eventually $\eta n \geq (\epsilon/4)/\zeta \geq (\epsilon/4)k_3$, and $\delta \leq (\epsilon/4)\bar{d}$. \square

Proof of Proposition 11. We will show that every partition of the vertices of G satisfies (2), and thus G has $\lfloor t \rfloor$ edge-disjoint spanning trees by Theorem 8.

We say that a set $S \subseteq V$ is *large* if $|S| \geq \zeta n$. We say that a partition of V is *simple* if each class either is large or a singleton (that is, it consists of a single vertex). Recall that $m(\mathcal{P})$ denotes the number of edges with ends in distinct parts of \mathcal{P} .

Claim 1. If \mathcal{P} is a simple partition, then $m(\mathcal{P}) \geq t(|\mathcal{P}| - 1)$.

Assume Claim 1 holds, and suppose for a contradiction that there is a partition \mathcal{P} of V such that $m(\mathcal{P}) < t(|\mathcal{P}| - 1)$. By Claim 1, \mathcal{P} is not a simple partition. Given a set S and a vertex $v \in S$, let $d_S(v)$ denote the number of neighbours of v inside S . Since \mathcal{P} is not simple, we can find a non-large part S of \mathcal{P} with at least 2 vertices. By Condition (e') and since $\epsilon \leq 1$, S must contain one vertex w with

$$d_S(w) \leq \frac{2|E[S]|}{|S|} \leq \frac{2\epsilon t|S|}{4|S|} \leq t/2. \quad (7)$$

Moreover, condition (d') implies that

$$m(\mathcal{P}) \geq (t/2)(|\mathcal{P}| - 1). \quad (8)$$

Let \mathcal{P}' be obtained from \mathcal{P} by turning w into a singleton. We have $|\mathcal{P}'| = |\mathcal{P}| + 1$ and $m(\mathcal{P}') = m(\mathcal{P}) + d_S(w)$. Combining this facts together with (7) and (8), we obtain

$$\frac{m(\mathcal{P}')}{|\mathcal{P}'| - 1} = \frac{m(\mathcal{P}) + d_S(w)}{|\mathcal{P}|} \leq \frac{m(\mathcal{P}) + t/2}{|\mathcal{P}|} \leq \frac{m(\mathcal{P}) + \frac{m(\mathcal{P})}{|\mathcal{P}| - 1}}{|\mathcal{P}|} = \frac{m(\mathcal{P})}{|\mathcal{P}| - 1}. \quad (9)$$

Repeat this procedure of turning vertices into singletons until no parts of size between 2 and ζn remain, and therefore obtain a simple partition \mathcal{P}'' . Since (9) holds in each iteration, we have

$$\frac{m(\mathcal{P}'')}{|\mathcal{P}''| - 1} \leq \frac{m(\mathcal{P})}{|\mathcal{P}| - 1} < t,$$

which contradicts Claim 1.

To complete the argument, we proceed to prove Claim 1. Let \mathcal{P} be a simple partition. If all parts of \mathcal{P} are singletons, then we have $m(\mathcal{P}) = \frac{\bar{d}}{2}(n-1) = \frac{\bar{d}}{2}(|\mathcal{P}|-1)$. Suppose otherwise there is at least one large part. Since \mathcal{P} is simple, each large part has at least ζn vertices and so there are at most $\ell := 1/\zeta = O(1)$ large parts. Let k be the number of singletons in \mathcal{P} . Note that $k \leq (1-\zeta)n$ since any large part has at least ζn vertices.

Suppose first that $\zeta n \leq k \leq (1-\zeta)n$. Then the average degree of the singletons is at least $\bar{d}(1-o(1))$ by Condition (b'). Since there is at least one large part, the number of edges between the k singletons and this large part is at least $\eta\zeta\bar{d}k$ by Condition (c'). Hence, $m(\mathcal{P})$ is at least the number of edges incident with a singleton, which is at least

$$\frac{\bar{d}(1-o(1))k + \eta\zeta\bar{d}k}{2} \geq \frac{(1+\eta\zeta/2)k}{k+\ell-1}(\bar{d}/2)(k+\ell-1).$$

This satisfies Equation (2) with $\bar{d}/2 \geq t$ for large enough n , since $k \geq \zeta n$ and $\ell = O(1)$.

Suppose otherwise that $0 \leq k \leq \zeta n$. By Condition (a'), we have that either $\delta > \frac{(1+\epsilon)\bar{d}}{2}$; or there are no adjacent ϵ -light vertices and each vertex is adjacent to at most one ϵ -light vertex. The number of edges between singletons is at most $(\epsilon/4)tk \leq \epsilon tk$ by Condition (e'). In the first case where $\delta > \frac{(1+\epsilon)\bar{d}}{2}$, the total number of edges incident to the singletons is at least

$$\frac{(1+\epsilon)\bar{d}}{2}k - \epsilon tk \geq (1+\epsilon)tk - \epsilon tk \geq tk. \quad (10)$$

Now we consider the second case. Recall that a vertex is ϵ -light if it has degree at most $\delta + \epsilon\bar{d}$. Suppose that there are no adjacent ϵ -light vertices and each vertex is adjacent to at most one ϵ -light vertex. Let K_1 denote the set of singletons that are ϵ -light and K_2 the set of other singletons (singletons that are not ϵ -light). Let $k_i = |K_i|$ for $i = 1, 2$, so $k = k_1 + k_2$ (possibly $k_1, k_2 = 0$). Since there are no adjacent ϵ -light vertices, $|E(K_1, \overline{K_1})| \geq \delta k_1$. Since no two ϵ -light vertices have a common neighbour, we have $d_{[n] \setminus K_1}(v) \geq \delta + \epsilon\bar{d} - 1$, for every $v \in K_2$. Moreover, Condition (e') guarantees that there are at most $\epsilon tk_2 \leq \epsilon\bar{d}k_2/2$ edges inside K_2 , and therefore $|E(K_2, \overline{K_2} \setminus K_1)| \geq (\delta + \epsilon\bar{d} - 1)k_2 - \epsilon\bar{d}k_2/2$. Thus, the total number of edges incident with singletons is at least

$$\delta k_1 + (\delta + \epsilon\bar{d} - 1)k_2 - \epsilon\bar{d}k_2/2 \geq \delta k \geq tk, \quad (11)$$

eventually as $\bar{d} = \omega(1)$ by our assumption. Thus, we have proved that in both cases, the number of edges incident with singletons is at least tk . If the number of large parts is exactly 1, then (2) holds as $|\mathcal{P}| = k + 1$ and $m(\mathcal{P}) \geq tk$ by (10) and (11). Otherwise, if there are at least two large parts, the number of edges between any two of them is at least $\eta\zeta^2\bar{d}n$ by Condition (c'). Thus, for large enough n ,

$$m(\mathcal{P}) \geq tk + \eta\zeta^2\bar{d}n \geq t \left(k + \frac{\eta\zeta^2\bar{d}n}{t} \right) \geq t(k + \ell - 1),$$

since $t \leq \bar{d}/2$ and $\ell = O(1)$. □

4 Properties of $\mathcal{G}(n, p)$

In this section, we always let G denote $\mathcal{G}(n, p)$, and let $\delta := \delta(\mathcal{G}(n, p))$, $m := m(\mathcal{G}(n, p))$ and $\bar{d} := \bar{d}(\mathcal{G}(n, p)) = 2m/(n-1)$. Recall our earlier assumption that $n \geq 2$. So \bar{d} is well defined. For any vertex v , let d_v denote the degree of v in G .

4.1 Typical degrees

Our aim here is to show that m and \bar{d} are a.a.s. concentrated around their expected values, and that most of the vertices of $\mathcal{G}(n, p)$ have degree close to \bar{d} . To do so, we first state a version of the well-know Chernoff's bounds (see e.g. Theorems 4.4 and 4.5 in [32])

Theorem 12 (Chernoff's bounds). *Let X_1, \dots, X_n denote n independent Bernoulli variables. Let $X = \sum_{i=1}^n X_i$ and let $\mu = \mathbf{E}X$. Then for any $0 < \tau < 1$,*

$$\Pr(X \geq (1 + \tau)\mu) \leq \exp(-\tau^2\mu/3), \quad \Pr(X \leq (1 - \tau)\mu) \leq \exp(-\tau^2\mu/2).$$

Lemma 13. *For any function $\tau(n) < 1$, we have that the probability that $|\bar{d} - pn| \leq \tau pn$ and $|m - p\binom{n}{2}| \leq \tau p\binom{n}{2}$ is at least $1 - 2\exp(-A\tau^2pn^2)$ where $A = 1/12$.*

Proof. By the definition of \bar{d} , the events $|\bar{d} - pn| > \tau pn$ and $|m - p\binom{n}{2}| > \tau p\binom{n}{2}$ are equivalent. Then, since the number of edges in $\mathcal{G}(n, p)$ is distributed as $\mathbf{Bin}(\binom{n}{2}, p)$, we apply Chernoff's bound in Theorem 12 and obtain

$$\Pr\left(|m - p\binom{n}{2}| > \tau p\binom{n}{2}\right) \leq 2\exp\left(-\frac{\tau^2 p\binom{n}{2}}{3}\right). \quad \square$$

Lemma 14. *Let $f \geq 0$ be any function of n such that $f \rightarrow \infty$. Then, there exists a constant $C > 0$ such that for every $f/n \leq p \leq 1$ the following holds in $\mathcal{G}(n, p)$ with probability at least $1 - e^{-C(pn)^{1/3}}$. The number of vertices with degree not in $[\bar{d} - (pn)^{2/3}, \bar{d} + (pn)^{2/3}]$ is at most $n/e^{Cf^{1/3}}$.*

Proof. We have that

$$\Pr(|d_v - \bar{d}| > (pn)^{2/3}) \leq \Pr\left(|\bar{d} - pn| > \frac{(pn)^{2/3}}{2}\right) + \Pr\left(|d_v - pn| > \frac{(pn)^{2/3}}{2}\right).$$

By Lemma 13 with $\tau = \frac{1}{2} \cdot (pn)^{-1/3}$,

$$\Pr\left(|\bar{d} - pn| \leq \frac{(pn)^{2/3}}{2}\right) \leq 2\exp(-An \cdot (pn)^{1/3}),$$

where A is a positive constant. By Chernoff's inequality in Theorem 12, for a positive constant B ,

$$\Pr\left(|d_v - pn| > \frac{(pn)^{2/3}}{2}\right) \leq 2\exp\left(-\frac{Bpn}{(pn)^{2/3}}\right) = 2\exp\left(-B(pn)^{1/3}\right).$$

Thus, there is a positive constant C such that,

$$\Pr(|d_v - \bar{d}| > (pn)^{2/3}) \leq \exp\left(-2C(pn)^{1/3}\right).$$

Thus, by Markov's inequality, the probability that the number of vertices with degree outside $[\bar{d} - (pn)^{2/3}, \bar{d} + (pn)^{2/3}]$ is more than $n \exp(-Cf^{1/3})$ is at most

$$\frac{n \exp\left(-2C(pn)^{1/3}\right)}{n \exp\left(-Cf^{1/3}\right)} \leq \exp\left(-C(pn)^{1/3}\right),$$

since $pn \geq f$. □

4.2 Maximum and minimum degree

In this section, we collect several results about the maximum and minimum degree of $\mathcal{G}(n, p)$ relevant to our argument. First, we give an easy upper-bound on the maximum degree.

Lemma 15. *Given any constant $\gamma > 0$, there exist positive constants C and K such that, if $p \leq \gamma \log n / (n - 1)$, then the maximum degree of $\mathcal{G}(n, p)$ is at most $K \log n$ with probability at least $1 - n^{-C}$.*

Proof. Let Δ denote the maximum degree in G and d_v denote the degree of v in G for any vertex v . By union bound and Chernoff's bound in Theorem 12, for any $K > \gamma$,

$$\Pr(\Delta \geq K \log n) \leq n \Pr(d_v \geq K \log n) \leq \exp\left(\frac{-(K \log n - p(n - 1))^2}{3p(n - 1)} + \log n\right).$$

Since $pn \leq \gamma \log n$, it suffices to choose K large enough so that $(K - \gamma)^2 / (3\gamma) - 1 > 0$. □

Our results about the minimum degree require the application of the first and the second moment methods to the number of vertices of low degree. The following lemma gives a lower-tail estimate for a Binomial random variable, and shall be used to bound the expected number of these low-degree vertices; the next lemma after that one will give us a bound on the variance.

Lemma 16. *For every constant $\eta > 0$ there exist positive constants C_1 and C_2 such that the following holds for any function $0 \leq p \leq 1/\sqrt{n}$ and every integer $0 < k \leq (1 - \eta)np$. Let $X \sim \mathbf{Bin}(n, p)$. Then,*

$$\Pr(X \leq k) = C \frac{e^{-pn}}{\sqrt{k}} \left(\frac{epn}{k}\right)^k \quad \text{with } C_1 \leq C \leq C_2.$$

Proof. It follows easily from Stirling's approximation, that there exist two positive constants A_1 and A_2 such that, for every $0 < k < \sqrt{n}$,

$$\frac{A_1}{\sqrt{k}} \left(\frac{en}{k}\right)^k \leq \binom{n}{k} \leq \frac{A_2}{\sqrt{k}} \left(\frac{en}{k}\right)^k.$$

Moreover, there exist positive constants B_1 and B_2 such that, for every $0 \leq p \leq 1/\sqrt{n}$ and every $0 < k < \sqrt{n}$,

$$B_1 e^{-pn} \leq (1-p)^{n-k} \leq B_2 e^{-pn}.$$

Therefore, there exist positive constants C'_1 and C'_2 not depending on p or k such that

$$\Pr(X = k) = \binom{n}{k} p^k (1-p)^{n-k} = C'_1 \frac{e^{-pn}}{\sqrt{k}} \left(\frac{epn}{k}\right)^k \quad \text{with} \quad C'_1 \leq C' \leq C'_2,$$

and the lower bound follows immediately since $\Pr(X \leq k) \geq \Pr(X = k)$. For the upper bound, let $f_i = \Pr(X = i) = \binom{n}{i} p^i (1-p)^{n-i}$, and observe that, for every $i \leq (1-\eta)np$,

$$\frac{f_{i-1}}{f_i} \leq \frac{i}{(n-i)p} \leq 1 - \frac{\eta}{2},$$

since $p \leq 1/\sqrt{n} \leq \frac{\eta}{2-\eta}$ eventually. Hence, there is a constant $D > 0$ only depending on η such that $\Pr(X \leq k) \leq D \Pr(X = k)$. \square

Lemma 17. *Let Y denote the number of vertices of degree at most k in $\mathcal{G}(n, p)$, where $p < 1$. Then $\mathbf{Var}(Y) \leq (\mathbf{E}Y)^2(p/(1-p) + 1/\mathbf{E}Y)$.*

Proof. Recall that d_v has distribution $\mathbf{Bin}(n-1, p)$. Let $q_{\leq}(r, t)$ denote the probability that a random variable with distribution $\mathbf{Bin}(r, p)$ has value at most t and let $q_{=}(r, t)$ denote the probability that it has value exactly t . Then

$$\begin{aligned} \mathbf{E}(Y^2) &= \sum_{u, v \in V} \Pr(d_v \leq k \text{ and } d_u \leq k) \\ &= \mathbf{E}Y + n(n-1) \left(p \cdot q_{\leq}(n-2, k-1)^2 + (1-p) \cdot q_{\leq}(n-2, k)^2 \right). \end{aligned}$$

This holds because for any distinct vertices $u, v \in V$, the number of neighbours of u in $V \setminus \{v\}$ and the number of neighbours of v in $V \setminus \{u\}$ are independent random variables with distribution $\mathbf{Bin}(n-2, p)$. Clearly, $q_{\leq}(r, t) = q_{\leq}(r, t-1) + q_{=}(r, t)$. And so

$$\begin{aligned} &p \cdot q_{\leq}(n-2, k-1)^2 + (1-p) \cdot q_{\leq}(n-2, k)^2 \\ &= q_{\leq}(n-2, k-1)^2 + 2(1-p) \cdot q_{=}(n-2, k) q_{\leq}(n-2, k-1) + (1-p) \cdot q_{=}(n-2, k)^2. \end{aligned} \quad (12)$$

Moreover,

$$q_{\leq}(n-1, k)^2 = \left(q_{\leq}(n-2, k-1) + (1-p) \cdot q_{=}(n-2, k) \right)^2.$$

Thus,

$$\begin{aligned} \mathbf{E}(Y^2) &= \mathbf{E}Y + n(n-1) \left(q_{\leq}(n-1, k)^2 + (1-p) q_{=}(n-2, k)^2 (1 - (1-p)) \right) \\ &\leq \mathbf{E}Y + n^2 q_{\leq}(n-1, k)^2 \left(1 + \frac{p}{1-p} \right) = \mathbf{E}(Y)^2 \left(1 + \frac{p}{1-p} + \frac{1}{\mathbf{E}(Y)} \right). \quad \square \end{aligned}$$

The following lemma bounds the probability that $\delta(\mathcal{G}(n, p))$ deviates significantly from its expected value in the range $0.9 \log n / (n-1) \leq p \leq \gamma \log n / (n-1)$, where $\gamma \geq 0.9$ is a constant. We will apply this lemma when we require more precise probability bounds than those that would follow from the Chernoff's inequalities in Theorem 12.

Lemma 18. *Let $\gamma \geq 0.9$ and $0 < \eta < 1$ be constants. Then there exists a constant $C > 0$ such that, for any functions p and α of n satisfying $0 < \alpha \leq 1 - \eta$ and $0.9 \log n / (n - 1) \leq p \leq \gamma \log n / (n - 1)$, the following holds:*

$$(i) \Pr\left(\delta \leq \alpha p(n-1)\right) \leq C \exp\left(\log n - p(n-1)\left(1 - \alpha \log\left(\frac{e}{\alpha}\right)\right) - \frac{1}{2} \log \log n\right) \text{ and}$$

$$(ii) \Pr\left(\delta > \alpha p(n-1)\right) \leq C \left(\frac{\log n}{n} + \exp\left(p(n-1)\left(1 - \alpha \log\left(\frac{e}{\alpha}\right)\right) - \log n + \frac{1}{2} \log \log n\right)\right).$$

Proof. Given an arbitrary vertex v , let $d_v \sim \mathbf{Bin}(n-1, p)$ be the degree of v . By Lemma 16, there exists a function $C' \in [C_1, C_2]$, where C_1, C_2 are positive constants that depend only on η and γ such that

$$\Pr(d_v \leq \alpha p(n-1)) = C' \exp\left(-p(n-1)\left(1 - \alpha \log\left(\frac{e}{\alpha}\right)\right) - \frac{1}{2} \log \log n\right). \quad (13)$$

Set C to be a constant larger than $C_2 + 2\gamma + 1/C_1 \geq C' + 2\gamma + 1/C'$. From (13), the expected number of vertices with degree at most $\alpha p(n-1)$ is

$$C' \exp\left(\log n - p(n-1)\left(1 - \alpha \log\left(\frac{e}{\alpha}\right)\right) - \frac{1}{2} \log \log n\right),$$

which implies (i), since $C \geq C'$. Finally, the proof of (ii) follows from Lemma 17 and Chebyshev's inequality, since $p/(1-p) \leq 2\gamma \log n/n$ and $C \geq 2\gamma + 1/C'$. \square

It is convenient to state an easy consequence of Lemma 18 as a separate result. We will use the following lemma when p is very close to the threshold function $\beta \log n / (n - 1)$ described in Theorem 2, in order to have a fairly precise bound of the probability that $\delta(\mathcal{G}(n, p))$ deviates slightly from $pn/2$. It is normally applied by choosing ϵ so that $|\epsilon| \log n$ is negligible compared to the other terms in (14) and (15).

Lemma 19. *Let $\gamma > \beta = 2/\log(e/2)$ and $0 < \eta < 1/2$ be constants. Then there exist positive constants C and D such that the following holds. For any functions p and ϵ of n satisfying $|\epsilon| \leq 1 - 2\eta$ and $0.9 \log n / (n - 1) \leq p \leq \gamma \log n / (n - 1)$, we have*

$$\Pr\left(\delta \leq \frac{(1+\epsilon)}{2} p(n-1)\right) \leq C \exp\left(-\frac{h}{\beta} - \frac{1}{2} \log \log n + D|\epsilon| \log n\right), \quad \text{and} \quad (14)$$

$$\Pr\left(\delta > \frac{(1+\epsilon)}{2} p(n-1)\right) \leq C \left(\frac{\log n}{n} + \exp\left(\frac{h}{\beta} + \frac{1}{2} \log \log n + D|\epsilon| \log n\right)\right), \quad (15)$$

where $h = h(n)$ is defined by

$$p = \frac{\beta \log n + h}{n-1}.$$

Proof. By putting $\alpha = (1 + \epsilon)/2$, we have that $\eta \leq \alpha \leq 1 - \eta$, and there is some constant $D' > 0$ depending only on η such that $|\alpha \log(e/\alpha) - 1 + 1/\beta| \leq D'|\epsilon|$. The result follows immediately from Lemma 18 and setting $D = \gamma D'$. \square

At this point, the reader may suspect that the relevant range of p for the study of the evolution of $\delta(\mathcal{G}(n, p))$ corresponds to $p = \Theta(\log n/n)$. Indeed, a careful application of Lemma 18 yields the following: if $p \sim c \log n/n$ for some constant $c > 1$, then a.a.s. $\bar{d} \sim c \log n$ and $\delta \sim g(c)\bar{d}$, where $g : [1, \infty) \rightarrow (0, 1)$ is a strictly increasing continuous function with $\lim_{c \rightarrow 1} g(c) = 0$ and $\lim_{c \rightarrow \infty} g(c) = 1$. We do not prove the above claim, as it is not needed in our argument, but rather collect several related statements together in the following lemma.

Lemma 20.

- (i) For any $p \leq 0.9 \log n/(n-1)$, a.a.s. $\delta(\mathcal{G}(n, p)) = 0$.
- (ii) For any constant $\epsilon > 0$, there exist constants $\gamma > 1$ and $C > 0$ such that, for every $\frac{0.9 \log n}{n-1} \leq p \leq \frac{\gamma \log n}{n-1}$, we have that $\delta(\mathcal{G}(n, p)) \leq \epsilon \bar{d}(\mathcal{G}(n, p))$ with probability at least $1 - n^{-C}$.
- (iii) Let $\gamma > 1$ be a constant. There exist positive constants ϵ and C such that, for $p \geq \gamma \log n/(n-1)$, we have that $\delta > \epsilon p(n-1)$ with probability at least $1 - n^{-C}$.
- (iv) For every constants $0 < \theta < 1$ and $C > 0$, there exists a constant $\gamma > 0$, such that for all $p \geq \gamma \log n/(n-1)$, we have $\Pr(\delta(\mathcal{G}(n, p)) \leq \theta p(n-1)) \leq n^{-C}$.

Proof. Part (i) is a well-known fact (see e.g. [4]). We first prove part (iv). By Theorem 12,

$$\Pr\left(\delta(\mathcal{G}(n, p)) \leq \theta p(n-1)\right) \leq n \exp\left(-\frac{(1-\theta)^2 p(n-1)}{2}\right) \leq n \exp\left(-\frac{(1-\theta)^2 \gamma \log n}{2}\right).$$

Thus, the statement holds by choosing γ sufficiently large so that $(1-\theta)^2 \gamma/2 - 1 > C$.

Next, we prove part (iii). From part (iv), there is a $\rho > 1$ such that, for any $p \geq \rho \log n/(n-1)$, $\delta(\mathcal{G}(n, p)) > (1/2)p(n-1)$ with probability at least $1 - 1/n$. Thus, for $\gamma \geq \rho$, statement (iii) follows immediately by choosing any $\epsilon \leq 1/2$ and any $C \leq 1$. So suppose otherwise that $1 < \gamma < \rho$. By Lemma 18, there is a positive constant C' so that, for every p in the considered range,

$$\Pr(\delta(\mathcal{G}(n, p)) \leq \epsilon p(n-1)) \leq C' \exp\left(\log n - \gamma \log n \left(1 - \epsilon \log\left(\frac{\epsilon}{\epsilon}\right)\right)\right).$$

Since $\gamma > 1$, we can choose $0 < \epsilon \leq 1/2$ small enough so that $\gamma(1 - \epsilon \log(\frac{\epsilon}{\epsilon})) > 1$. Then there is a constant $C'' > 0$ such that the above probability is at most $n^{-C''}$. The statement follows by choosing $C = \min\{1, C''\}$.

Finally, we prove part (ii). We assume without loss of generality that $\epsilon < 1$. We have that

$$\Pr\left(\delta(\mathcal{G}(n, p)) > \epsilon \bar{d}(\mathcal{G}(n, p))\right) \leq \Pr\left(\bar{d}(\mathcal{G}(n, p)) < 0.8 \log n\right) + \Pr\left(\delta(\mathcal{G}(n, p)) > 0.8 \epsilon \log n\right). \quad (16)$$

By Lemma 13, there is a positive constant B such that, for any $p \geq 0.9 \log n/(n-1)$,

$$\Pr\left(\bar{d}(\mathcal{G}(n, p)) < 0.8 \log n\right) \leq \Pr\left(\bar{d}(\mathcal{G}(n, p)) < (8/9)pn\right) \leq \exp(-Bn \log n). \quad (17)$$

Let $1 < \gamma < 8/7$ be a sufficiently small constant which we specify later, and put $\alpha = 0.8/\gamma$. In particular, $0.7 < \alpha < 0.8$. By Lemma 18, there is a constant $D > 0$ such that, for every p in the range $0.9 \log n/(n-1) \leq p \leq \gamma \log n/(n-1)$, we have

$$\begin{aligned} \Pr\left(\delta(\mathcal{G}(n, p)) > 0.8\epsilon \log n\right) &\leq \Pr\left(\delta(\mathcal{G}(n, p)) > \epsilon \alpha p(n-1)\right) \\ &\leq D \left(\frac{\log n}{n} + \exp\left(\gamma \log n \left(1 - \epsilon \alpha \log\left(\frac{e}{\epsilon \alpha}\right)\right) - \log n + \frac{1}{2} \log \log n\right) \right), \end{aligned}$$

where we also used that $1 - \epsilon \alpha \log(\frac{e}{\epsilon \alpha}) > 0$, as $\epsilon \alpha < 0.8$. Moreover, choosing $\gamma > 1$ small enough ensures that

$$B := \gamma \left(1 - \epsilon \alpha \log\left(\frac{e}{\epsilon \alpha}\right)\right) < \gamma \left(1 - 0.7\epsilon \log\left(\frac{e}{0.7\epsilon}\right)\right) < 1,$$

and the above probability is at most

$$D \left(\log n/n + \exp\left(- (1 - B) \log n + \frac{1}{2} \log \log n\right) \right). \quad (18)$$

Combining (16), (17) and (18) yields statement (ii), for C sufficiently small. \square

Finally, we include a result that compares the minimum degree of $\mathcal{G}(n, p)$ and $\mathcal{G}(n, \hat{p})$, when p and \hat{p} are close to one another.

Lemma 21. *For any constants $1 < \gamma_1 < \gamma_2$ and $\epsilon > 0$, there exist positive constants η and C such that, for any functions p and \hat{p} satisfying $\gamma_1 \log n/(n-1) \leq p \leq \hat{p} \leq \gamma_2 \log n/(n-1)$ and $\hat{p}/p - 1 \leq \eta$,*

$$\frac{\delta(\mathcal{G}(n, \hat{p}))}{\delta(\mathcal{G}(n, p))} - 1 \leq \epsilon$$

with probability at least $1 - n^{-C}$.

Proof. Assume without loss of generality that $0 < \epsilon < 1$. Choose constants γ_0 and γ_3 such that $1 < \gamma_0 < \gamma_1$ and $\gamma_2 < \gamma_3$. The function $f(y) = 1 - y \log(e/y)$ is a bijection from $[0, 1]$ to $[0, 1]$ (defining $f(0) = 1$), and is strictly decreasing in that domain. Therefore, we can define the constants $y_i = f^{-1}(1/\gamma_i)$ for $i = 0, 1, 2, 3$, which satisfy $0 < y_0 < y_1 < y_2 < y_3 < 1$. Moreover, there exists a constant $D > 0$ such that $f' \leq -D$ for all y in $[y_0, y_3]$, since this interval is a compact set and $f' < 0$ is continuous there.

Pick two positive constants ζ and ξ sufficiently small so that the following conditions are satisfied: $y_2 + \zeta \leq y_3$; $\zeta \leq \epsilon y_0/3$; $\xi \leq \gamma_1 D \zeta/4$; and $\gamma_1/(1 + \xi) \geq \gamma_0$ (note that the choice of ξ depends on ζ). With all these constants in mind, we choose $\eta > 0$ in the statement sufficiently small so that $\eta \leq \gamma_1 D \zeta/4$ and $\eta \leq \epsilon/3$.

Put $x = p(n-1)/\log n$ and $\hat{x} = \hat{p}(n-1)/\log n$. In general, x and \hat{x} are functions of n with $\gamma_1 \leq x \leq \hat{x} \leq \gamma_2$, and moreover, from the assumption in the statement, $\hat{x} \leq (1 + \eta)x$. Define $\alpha = f^{-1}((1 + \xi)/x)$ and $\hat{\alpha} = \alpha + \zeta$, also functions of n . From the previous conditions $\gamma_1 \leq x \leq \gamma_2$ and $1 + \xi \leq \gamma_1/\gamma_0$, we deduce that $y_0 \leq \alpha \leq y_2$. From this and since $y_2 + \zeta \leq y_3$, we get $\hat{\alpha} \leq y_3$. So in particular $\alpha, \hat{\alpha} \in [y_0, y_3]$. We have

$$x(1 - \alpha \log(e/\alpha)) = x f(\alpha) = 1 + \xi. \quad (19)$$

Moreover, using the bound on f' in $[y_0, y_3]$ and some of the earlier constraints on x , \hat{x} , ξ and η ,

$$\hat{x}(1 - \hat{\alpha} \log(e/\hat{\alpha})) \leq (1 + \eta)x(f(\alpha) - D\zeta) \leq (1 + \eta)(1 + \xi - \gamma_1 D\zeta) \leq 1 - \gamma_1 D\zeta/2. \quad (20)$$

Using Lemma 18 together with (19) and (20), we conclude that

$$\delta(\mathcal{G}(n, p)) > \alpha x \log n \quad \text{and} \quad \delta(\mathcal{G}(n, \hat{p})) \leq \hat{\alpha} \hat{x} \log n$$

with probability at least $1 - n^{-C}$, for any positive constant C satisfying $C < \min\{1, \xi, \gamma_1 D\zeta/2\}$. This last event implies that

$$\delta(\mathcal{G}(n, \hat{p})) \leq (1 + \zeta/\alpha)(1 + \eta)\alpha x \log n \leq (1 + \epsilon/3)^2 \delta(\mathcal{G}(n, p)) \leq (1 + \epsilon)\delta(\mathcal{G}(n, p)),$$

since $\zeta \leq \epsilon y_0/3 \leq \epsilon\alpha/3$, $\eta \leq \epsilon/3$ and $\epsilon < 1$. This completes the proof of the Lemma. \square

4.3 Light vertices

Recall that an ϵ -light vertex was defined to be a vertex of degree at most $\delta + \epsilon\bar{d}$. The following result shows that a.a.s. all ϵ -light vertices of $\mathcal{G}(n, p)$ are at least three steps apart for a certain range of p .

Lemma 22. *Suppose $0.9 \log n/(n-1) \leq p \leq \gamma \log n/(n-1)$ for some constant $\gamma \geq 0.9$. Then there exist constants $\epsilon > 0$ and $C > 0$ such that the following holds in $\mathcal{G}(n, p)$ with probability at least $1 - n^{-C}$. There is no pair of adjacent ϵ -light vertices and no two ϵ -light vertices have a common neighbour.*

Proof. Let $x = p(n-1)/\log n$. For each $x \in [0.9, \gamma]$, define $\alpha = \alpha(x)$ to be the only solution in $(0, 1)$ of

$$x(1 - \alpha \log(e/\alpha)) = 0.8. \quad (21)$$

It is straightforward to verify that $\alpha \in (0, 1)$ is well defined and strictly increasing with respect to $x \in [0.9, \gamma]$. Consider the constant $\hat{\epsilon} = 0.1/(\gamma - 0.8)$, and define $\hat{\alpha} = (1 + \hat{\epsilon})\alpha$. Recall that both α and $\hat{\alpha}$ are functions of $x = p(n-1)/\log n$. Then, using (21) and the fact that $\hat{\epsilon} \leq 0.1/(x - 0.8)$, we obtain

$$x(1 - \hat{\alpha} \log(e/\hat{\alpha})) > x(1 - \hat{\alpha} \log(e/\alpha)) = x - (1 + \hat{\epsilon})(x - 0.8) \geq 0.7. \quad (22)$$

From (21) and by Lemma 18 (ii), we can bound

$$\Pr(\delta > \alpha p(n-1)) \leq Dn^{-0.19}, \quad (23)$$

for a constant $D > 0$ not depending on p . Assume for the rest of the argument that D is sufficiently large. Let S be the set of vertices of degree at most $\hat{\alpha}p(n-1)$. By (13) in the proof of Lemma 18 and (22), the probability that a vertex v belongs to S is

$$\Pr(v \in S) = \Pr(d_v \leq \hat{\alpha}p(n-1)) \leq Dn^{-0.7}. \quad (24)$$

We can upper-bound the probability that a pair of vertices u and v are adjacent and belong to S , by

$$p\Pr(d_u \leq \hat{\alpha}p(n-1))\Pr(d_v \leq \hat{\alpha}p(n-1)) = p(\Pr(v \in S))^2.$$

Multiplying this by the number of possible pairs and using (24), we get that the probability that S contains some adjacent pair of vertices is at most

$$\binom{n}{2}p(\Pr(v \in S))^2 \leq D\gamma n^{-0.4} \log n. \quad (25)$$

By a similar argument, the probability that S contains a pair of vertices with a common neighbour is at most

$$\binom{n}{2}(n-2)p^2(\Pr(v \in S))^2 \leq D\gamma^2 n^{-0.4} \log^2 n. \quad (26)$$

Finally, we define $\epsilon = \alpha(0.9)\hat{\epsilon}/2$. Recall that α is increasing in $[0.9, \gamma]$, and then $\epsilon \leq \hat{\alpha}\hat{\epsilon}/2$. It follows from Lemma 13 that $\bar{d} = \bar{d}(\mathcal{G}(n, p))$ is at most $2p(n-1)$ with probability at least $1 - D/n$, assuming that D is large enough. If this event and the one in (23) hold together, then

$$\delta + \epsilon\bar{d} \leq (\alpha + 2\epsilon)p(n-1) \leq \hat{\alpha}p(n-1),$$

and therefore all ϵ -light vertices are contained in S . Putting everything together, the statement holds with probability at least $1 - n^{-C}$, for some small enough constant $C > 0$. \square

We include an extension of the previous lemma in terms of two random graphs $G_1 \sim \mathcal{G}(n, p_1)$ and $G_2 \sim \mathcal{G}(n, p_2)$, with $0 \leq p_1 \leq p_2 < 1$, which are coupled together so that $G_1 \subseteq G_2$. This standard coupling can be achieved in the following way. Let G_1 distributed as $\mathcal{G}(n, p_1)$ and let G_2 the supergraph of G_1 obtained by adding each edge not in G_1 independently with probability $(p_2 - p_1)/(1 - p_1)$. Then $G_1 \subseteq G_2$ and G_2 has the same distribution as $\mathcal{G}(n, p_2)$ (for more details, we refer readers to Section 1.1 in [27]). The following lemma will be used in Section 7, and can be proved in the exact same way as Lemma 22, but replacing p by p' in (25) and (26).

Lemma 23. *Suppose $0.9 \log n/(n-1) \leq p \leq p' \leq \gamma \log n/(n-1)$ for some constant $\gamma \geq 0.9$. Let $G_1 \subseteq G_2$ where $G_1 \sim \mathcal{G}(n, p)$ and $G_2 \sim \mathcal{G}(n, p')$. Then there exist constants $\epsilon > 0$ and $C > 0$ such that the following holds in $\mathcal{G}(n, p)$ and $\mathcal{G}(n, p')$ with probability at least $1 - n^{-C}$. Let S be the set of ϵ -light vertices in G_1 . Then in G_2 , there is no edge induced by S , and no two vertices in S adjacent to a common vertex.*

4.4 Graph expansion

For any sets $S, S' \subseteq [n]$, let $E(S, S')$ be the set of edges in G with one end in S and the other in S' .

Lemma 24. *Let $f \geq 0$ be any function of n such that $f \rightarrow \infty$, and $\zeta > 0$ any fixed constant. Then, there exists a constant $C > 0$ such that for every $f/n \leq p \leq 1$ the following holds in $\mathcal{G}(n, p)$ with probability at least $1 - e^{-Cpm^2}$. For every disjoint sets $S, S' \subseteq [n]$ with $|S|, |S'| \geq \zeta n$ we have $|E(S, S')| \geq (\bar{d}/4)|S||S'|/n$.*

Proof. The variable $|E(S, S')|$ has distribution $\mathbf{Bin}(|S||S'|, p)$. By Lemma 13 with $\tau = 1/4$, we have that

$$\Pr\left(\bar{d} \geq \frac{5}{4}pn\right) \leq 2\exp(-Apn^2)$$

where $A = 1/12$ and $n \geq 2$. By Chernoff's bound in Theorem 12, for a positive constant B ,

$$\Pr\left(\mathbf{Bin}\left(|S||S'|, p\right) < \frac{5p}{16}|S||S'|\right) \leq \exp(-Bp|S||S'|).$$

Hence, the probability that there exist such S and S' is at most

$$\begin{aligned} 2\exp\left(-Apn^2\right) + \sum_{s, s' > \zeta n} \binom{n}{s} \binom{n}{s'} \exp(-Bpss') &\leq \exp\left(-Apn^2\right) + n^2 \cdot 2^n \cdot 2^n \exp(-B\zeta^2pn^2) \\ &\leq 2\exp\left(-Apn^2\right) + \exp(B''n - B\zeta^2pn^2). \end{aligned}$$

for a positive constant B'' and we are done since $pn \geq f \rightarrow \infty$ as $n \rightarrow \infty$. \square

Lemma 25. *Let $f \geq 0$ be any function of n such that $f \rightarrow \infty$, and let $\alpha > 0$ be any fixed constant. Then, there exist constants $\zeta > 0$ and $C > 0$ such that for every $f/n \leq p \leq 1$ the following holds in $\mathcal{G}(n, p)$ with probability at least $1 - Ce^{-(pn)^2}$. For all $s \leq \zeta n$ and every set S of size s , we have that $|E[S]| \leq \alpha pns$.*

Proof. The result is trivial for any set of size $s \leq 2\alpha pn$ since $|E[S]| \leq s^2/2 \leq s(2\alpha pn)/2 = \alpha pns$. Let $\zeta > 0$ be small enough so that $\frac{e\zeta}{2\alpha} < e^{-1/\alpha^2}$. The expected number of sets of size $2\alpha pn \leq s \leq \zeta n$ containing at least αpns edges is at most

$$\begin{aligned} \binom{n}{s} \binom{\binom{s}{2}}{\lceil \alpha pns \rceil} p^{\lceil \alpha pns \rceil} &\leq \left(\frac{en}{s} \left(\frac{es}{2\alpha n}\right)^{\lceil \alpha pns \rceil}\right)^s = \left(\frac{e^2}{2\alpha} \left(\frac{es}{2\alpha n}\right)^{\lceil \alpha pns \rceil - 1}\right)^s \\ &\leq \left(A \left(\frac{e\zeta}{2\alpha}\right)^{\alpha pn}\right)^s < (Ae^{-pn/\alpha})^s, \end{aligned}$$

for some constant $A > 0$ depending only on f and ζ (we used the fact that the exponent $\lceil \alpha pn \rceil - 1 \geq \alpha f - 1$, which eventually becomes positive as $f \rightarrow \infty$).

Summing the expectation above over all $s \geq 2\alpha pn$, we get

$$\sum_{s \geq 2\alpha pn} (Ae^{-pn/\alpha})^s \leq (Ae^{-pn/\alpha})^{2\alpha pn} \frac{1}{1 - Ae^{-f/\alpha}} \leq Ce^{-(pn)^2},$$

for some constant $C > 0$ only depending on f and ζ . \square

For any $S \subseteq [n]$, let \bar{S} denote $[n] \setminus S$.

Lemma 26. *Let $\gamma > 1$ be a fixed constant. There exists a constant $C > 0$ such that for any $p = p(n) \geq \gamma \log n / (n - 1)$, the following holds in $\mathcal{G}(n, p)$ with probability at least $1 - n^{-C}$. For every $S \subsetneq [n]$ with $2 \leq |S| \leq n - 2$, $|E(S, \bar{S})| \geq 1.5\delta$.*

Proof. Without loss of generality, we may assume that $|S| \leq |\bar{S}|$. Since $p \geq \gamma \log n / (n - 1)$ for some $\gamma > 1$, by Lemma 20 (iii), there exist constants $\epsilon > 0$ and $C_1 > 0$ such that with probability at least $1 - n^{-C_1}$, $\delta = \delta(\mathcal{G}(n, p)) \geq \epsilon pn$. Let $\alpha = \epsilon/8$. Then by Lemma 25, there exist constants $\zeta > 0$ and $C_2 > 0$ such that with probability at least $1 - n^{-C_2}$, for all sets S with size at most ζn , $|E(S, S)| \leq \alpha pn|S|$. Then, with probability at least $1 - n^{-C_1} - n^{-C_2}$, for all these S , $|E(S, \bar{S})| \geq \delta|S| - 2\alpha pn|S| \geq (3/4)\delta|S| \geq 1.5\delta$, as $|S| \geq 2$. Now by Lemmas 24 and 13, there exists another constant $C_3 > 0$ such that with probability at least $1 - n^{-C_3}$, for all sets S with size at least ζn , $|E(S, \bar{S})| \geq (\bar{d}/4)\zeta^2 n$, where $\bar{d} = \bar{d}(\mathcal{G}(n, p)) \sim np$. Clearly, $(\bar{d}/4)\zeta^2 n \geq 1.5\delta$ with probability at least $1 - n^{-C_4}$ for some $C_4 > 0$. The lemma follows by choosing $C < \min\{C_i : 1 \leq i \leq 4\}$. \square

5 Proof of Theorem 1

We proceed to prove Theorem 1, as a consequence of Propositions 10 and 11. For the rest of the argument, let $\delta := \delta(\mathcal{G}(n, p))$ and let $\bar{d} := \bar{d}(\mathcal{G}(n, p))$. We split the argument into cases depending on the range of p .

First observe that by Lemma 20 (i) we can assume that $p \geq 0.9 \log n / (n - 1)$, since for $p \leq 0.9 \log n / (n - 1)$ the random graph $\mathcal{G}(n, p)$ is a.a.s. disconnected and has minimum degree zero, so the statement of Theorem 1 holds trivially.

Let γ_2 be a large enough constant so that for $p \geq \gamma_2 \log n / (n - 1)$ we have $\delta > (3/4)\bar{d}$ a.a.s. (see Lemma 20 (iv) and Lemma 13). Let $\epsilon < 1/2$ be the constant given by Lemma 22 with $\gamma = \gamma_2$. Let $\gamma_1 \in (1, \gamma_2)$ be the constant given by Lemma 20 (ii) with $\epsilon/4$.

For $0.9 \log n / (n - 1) \leq p \leq \gamma_1 \log n / (n - 1)$, we only need to show that $\mathcal{G}(n, p)$ a.a.s. satisfies the hypothesis of Proposition 10. First, we note from Lemma 13 that $\bar{d} \sim pn \rightarrow \infty$. Condition (a) holds by our choice of γ_1 . Condition (b) follows from Lemma 25 with any $\alpha < \epsilon/4$, since $\bar{d} \sim pn$. Fix ζ as given by that lemma. Condition (c) with $\eta = \zeta^2/4$ is a consequence of Lemma 24.

Finally, we show that $\mathcal{G}(n, p)$ a.a.s. satisfies the conditions in Proposition 11 for the range $p \geq \gamma_1 \log n / (n - 1)$. First note that $\delta = \Omega(\bar{d})$ by Lemma 20 (iii). Condition (a') is satisfied for $p \geq \gamma_2 \log n / (n - 1)$, since a.a.s. $\delta > \frac{(1+\epsilon)\bar{d}}{2}$ (by our choice of γ_2 and since $\epsilon < 1/2$); and it is also satisfied for $\gamma_1 \log n / (n - 1) \leq p \leq \gamma_2 \log n / (n - 1)$, since a.a.s. no ϵ -light vertices are adjacent nor have a common neighbours (by our choice of ϵ). For condition (e'), note that $\epsilon t / \bar{d}$ is bounded away from 0 since $\delta = \Omega(\bar{d})$. Therefore, the condition follows from Lemma 25 with $\alpha = \epsilon t / (8\bar{d})$ (also using that a.a.s. $\bar{d} \geq pn/2$), and this determines our choice of ζ . Condition (b') holds a.a.s. by Lemma 14. Condition (c') holds a.a.s. by Lemma 24. Condition (d') holds a.a.s. by Lemma 26. \square

6 Proof of Theorem 2

The number of edges in $G \sim \mathcal{G}(n, p)$ is a binomial random variable distributed as $\mathbf{Bin}(\binom{n}{2}, p)$. If $p < 0.9 \log n / n$, then by Lemma 20 (i), a.a.s. $\delta(G) = 0$ and thus a.a.s. $\delta(G) \leq \bar{d}(G)/2$. Assume $p \geq 0.9 \log n / n$. By Lemma 13, a.a.s. $|\bar{d}/2 - pn/2| \leq \omega_n \sqrt{p}$, where $\bar{d} = \bar{d}(G)$. By Lemma 20 (iv), there is a constant $\gamma > 0$, such that for all $p \geq \gamma \log n / n$, a.a.s. $\delta(G) \geq$

$(3/4)pn$. Hence, for p in this range, a.a.s. $\delta(G) > \bar{d}/2$. Now we only consider $0.9 \log n/n \leq p \leq \gamma \log n/n$.

Let ω_n be a positive-valued function of n that goes to infinity arbitrarily slowly as $n \rightarrow \infty$, and let $\epsilon = \omega_n/\sqrt{pn}$. Define $f = f(n)$ by $p = \frac{\beta \log n + f}{n-1}$. To prove statement (ii), we assume $f \geq -\beta \log \log n/2 + \omega_n$. By Lemma 19 (with $h = f$), we have that

$$\Pr\left(\delta \leq \frac{1}{2}(1 + \epsilon)pn\right) = O\left(\exp\left(-\frac{f}{\beta} - \frac{1}{2} \log \log n + O(\epsilon \log n)\right)\right) = o(1), \quad (27)$$

as $-f/\beta - \log \log n/2 \leq -\omega_n$, whereas $\epsilon \log n = O(\omega_n \sqrt{p}) = o(1)$. Moreover, by Lemma 13, $\Pr(\bar{d}/2 \geq (1 + \epsilon)pn/2) = o(1)$. Thus, a.a.s. $\delta > \bar{d}/2$ and so $T(G) = \lfloor \bar{d}/2 \rfloor$ by Theorem 1. This completes the proof of statement (ii). On the other hand, if $f \leq -\beta \log \log n/2 - \omega_n$, then $f/\beta + \log \log n/2 \leq -\omega_n$, and thus by Lemma 19 (with $h = f$, and ϵ replaced by $-\epsilon$),

$$\Pr\left(\delta > \frac{1}{2}(1 - \epsilon)pn\right) = O\left(\frac{\log n}{n} + \exp\left(\frac{f}{\beta} + \frac{1}{2} \log \log n + O(\epsilon \log n)\right)\right) = o(1). \quad (28)$$

Again by Lemma 13, $\Pr(\bar{d}/2 \leq (1 - \epsilon)pn/2) = o(1)$. Thus, a.a.s. $\delta < \bar{d}/2$ and thus $T(G) = \delta(G)$ by Theorem 1, which completes the proof of statement (i). \square

7 Proof of Theorem 3

A standard tool to investigate the random graph process $G_0, \dots, G_m, \dots, G_{\binom{n}{2}}$ is the related continuous random graph process $(\mathcal{G}_p)_{p \in [0,1]}$ defined as follows. For each edge e of the complete graph with vertex set $[n]$, we associate a random variable P_e uniformly distributed in $[0, 1]$ and independent from all others. Then, for any $p \in [0, 1]$, we define \mathcal{G}_p to be the graph with vertex set $[n]$ and precisely those edges e such that $p \geq P_e$. Note that for each p , \mathcal{G}_p is distributed as $\mathcal{G}(n, p)$. This provides us with a useful way of coupling together $\mathcal{G}(n, p)$ for several values of p , since $p \leq p'$ implies $\mathcal{G}_p \subseteq \mathcal{G}_{p'}$. Moreover, let $p(m) = \min\{p \in [0, 1] : \mathcal{G}_p \text{ has at least } m \text{ edges}\}$. Then, $\mathcal{G}_{p(0)}, \dots, \mathcal{G}_{p(m)}, \dots, \mathcal{G}_{p(\binom{n}{2})}$ is distributed as $G_0, \dots, G_m, \dots, G_{\binom{n}{2}}$, since all P_e are different with probability 1. For more details on the connection between $(\mathcal{G}_p)_{p \in [0,1]}$ and $(G_m)_{0 \leq m \leq \binom{n}{2}}$ and further properties, we refer the reader to [27].

In this article, we prove several statements that hold a.a.s. simultaneously for all m in the random graph process $(G_m)_{0 \leq m \leq \binom{n}{2}}$. To do so, it is often convenient to use small bits of the continuous random graph process as follows. Given p_0 and p_1 as functions of n such that $0 \leq p_0 \leq p_1 \leq 1$, we consider $(\mathcal{G}_p)_{p_0 \leq p \leq p_1}$. Let $m_0 = m(\mathcal{G}_{p_0})$ and $m_1 = m(\mathcal{G}_{p_1})$. (Note that m_0 and m_1 are random variables with $m_0 \leq m_1$, since $\mathcal{G}_{p_0} \subset \mathcal{G}_{p_1}$.) We colour all edges of $\mathcal{G}_{p_0} = G_{m_0}$ red and the remaining $m_1 - m_0$ edges in $\mathcal{G}_{p_1} \setminus \mathcal{G}_{p_0}$ blue. Then we can interpret $G_{m_0}, G_{m_0+1}, \dots, G_{m_1}$ as a random graph process in which we sequentially add blue edges to G_{m_0} , so that each G_m has the m_0 red edges of G_{m_0} together with the first $m - m_0$ blue edges we add in the process. This interpretation will be used many times throughout the argument.

We first prove the following result, which is stronger than part (ii) of Theorem 3, and is also used in the argument for part (i).

Theorem 27. *Consider the random graph process $(\mathcal{G}_p)_{0 \leq p \leq 1}$. We have that a.a.s.*

(i) *for all $p \leq \frac{\beta(\log n - \log \log n/2) - \omega(1)}{n-1}$, we have $\delta(\mathcal{G}_p) \leq \bar{d}(\mathcal{G}_p)/2$; and*

(ii) *for all $p \geq \frac{\beta(\log n - \log \log n/2) + \omega(1)}{n-1}$, we have $\delta(\mathcal{G}_p) > \bar{d}(\mathcal{G}_p)/2$.*

Moreover, for every constant $0 < \theta < 1$, there is a constant $\rho > 0$ such that a.a.s.

(iii) *for all $\rho \log n / (n-1) \leq p \leq 1$, we have $\delta(\mathcal{G}_p) > \theta \bar{d}(\mathcal{G}_p)$.*

Proof. First, we prove statement (iii). We will show that for every $0 < \theta < 1$, there exists $\rho > 0$ such that a.a.s.

$$\delta(G_m) > \theta \bar{d}(G_m), \quad \text{for all } m \geq m_0 = (\rho/4)n \log n. \quad (29)$$

Then, let $p_0 = \rho \log n / (n-1)$. By Chernoff's bound in Theorem 12, a.a.s. $m(\mathcal{G}(n, p_0)) > m_0$, i.e. a.a.s. $p(m_0) < p_0$. It follows then that a.a.s. $\delta(\mathcal{G}_p) > \theta \bar{d}(\mathcal{G}_p)$ for all $p \geq p_0$. Now we prove (29). For each m , let $\bar{p} = m / \binom{n}{2}$. Then

$$\Pr\left(\delta(G_m) \leq \theta \cdot 2m / (n-1)\right) = \Pr\left(\delta(\mathcal{G}(n, \bar{p})) \leq \theta \bar{p}n \mid m(\mathcal{G}(n, \bar{p})) = m\right).$$

By the choice of \bar{p} , $h(i) = \Pr(m(\mathcal{G}(n, \bar{p})) = i)$ is maximized at $i = m$. Hence, $\Pr(m(\mathcal{G}(n, \bar{p})) = m) \geq n^{-2}$. Thus,

$$\Pr\left(\delta(G_m) \leq \theta \cdot 2m / (n-1)\right) \leq \frac{\Pr\left(\delta(\mathcal{G}(n, \bar{p})) \leq \theta \bar{p}n\right)}{\Pr(m(\mathcal{G}(n, \bar{p})) = m)} \leq n^2 \Pr\left(\delta(\mathcal{G}(n, \bar{p})) \leq \theta \bar{p}n\right).$$

By Lemma 20 (iv), for every $0 < \theta < 1$, we can choose $\rho > 0$ sufficiently large such that the probability on the right-hand side above is less than $1/n^5$ for every $m \geq (\rho/4)n \log n$ (correspondingly, $\bar{p} \geq (\rho/2) \log n / (n-1)$). Hence, taking a union bound over the $O(n^2)$ possible values of m , we deduce that claim (29) is true with probability at least $1 - O(n^{-1})$.

Next, we prove statements (i) and (ii). Let $f = o(\sqrt{\log n})$ be a function that goes to ∞ arbitrarily slowly, as $n \rightarrow \infty$. Let $p_i = (\beta(\log n - \log \log n/2) - f^2 - if) / (n-1)$ and let $q_i = (\beta(\log n - \log \log n/2) + f^2 + if) / (n-1)$, for each $i \geq 1$. Let T be the smallest integer such that $p_T \leq 0.9 \log n / (n-1)$ and redefine $p_T = 0.9 \log n / (n-1)$. Let ρ be the constant satisfying statement (iii) with $\theta = 3/4$. Let T' be the smallest integer such that $q_{T'} \geq \rho \log n / (n-1)$ and redefine $q_{T'} = \rho \log n / (n-1)$. Obviously, $T, T' = O(\log n)$.

Claim 2. There exists a positive constant C , such that, for every $1 \leq i < T$,

$$\Pr\left(\delta(\mathcal{G}_{p_i}) > \bar{d}(\mathcal{G}_{p_{i+1}})/2\right) \leq C(f^{-i} + \log n/n),$$

and for every $1 \leq i < T'$,

$$\Pr\left(\delta(\mathcal{G}_{q_i}) \leq \bar{d}(\mathcal{G}_{q_{i+1}})/2\right) \leq C(f^{-i} + n^{-1}).$$

By Lemma 20 (i) and from the monotonicity of $\delta(\mathcal{G}_p)$ with respect to p , a.a.s. for all $p < p_T = 0.9 \log n(n-1)$, we have $\delta(\mathcal{G}_p) = 0$ and thus $\delta(\mathcal{G}_p) \leq \bar{d}(\mathcal{G}_p)/2$ holds. By Claim 2, with probability at least

$$1 - \sum_{1 \leq i < T} C(f^{-i} + \log n/n) = 1 - o(1),$$

for all $1 \leq i < T$ and for all $p_{i+1} \leq p \leq p_i$,

$$\delta(\mathcal{G}_p) \leq \delta(\mathcal{G}_{p_i}) \leq \bar{d}(\mathcal{G}_{p_{i+1}})/2 \leq \bar{d}(\mathcal{G}_p)/2.$$

Thus, a.a.s. $\delta(\mathcal{G}_p) \leq \bar{d}(\mathcal{G}_p)/2$ for all $p \leq \frac{\beta(\log n - \log \log n/2) - \omega(1)}{n-1}$, since f is an arbitrary slowly growing function in $\omega(1)$, and statement (i) follows.

By (iii) (with $\theta = 3/4$), we only need to prove that a.a.s. $\delta(\mathcal{G}_p) > \bar{d}(\mathcal{G}_p)/2$ for all p satisfying $q_1 \leq p \leq q_{T'} = \rho \log n/(n-1)$. Similarly as in the previous argument, with probability at least

$$1 - \sum_{1 \leq i < T'} C(f^{-i} + n^{-1}) = 1 - o(1),$$

for all $1 \leq i < T'$ and for every p with $q_i \leq p \leq q_{i+1}$,

$$\delta(\mathcal{G}_p) \geq \delta(\mathcal{G}_{q_i}) > \bar{d}(\mathcal{G}_{q_{i+1}})/2 \geq \bar{d}(\mathcal{G}_p)/2.$$

Thus, a.a.s. $\delta(\mathcal{G}_p) > \bar{d}(\mathcal{G}_p)/2$ for all $p \geq \frac{\beta(\log n - \log \log n/2) + \omega(1)}{n-1}$, as required in statement (ii).

Finally, we prove Claim 2. In this argument the asymptotic statements are uniform for all $p \in [p_T, p_1] \cup [q_1, q_{T'}]$. By Lemma 13, for $\sigma = n^{-1/3}$ and a positive constant A ,

$$\Pr(|\bar{d}(\mathcal{G}_p) - pn| > \sigma pn) \leq \exp(-A\sigma^2 n^2 p) = o(n^{-1}). \quad (30)$$

Note that the event $\delta(\mathcal{G}_{p_i}) > \frac{1-\sigma}{2} p_{i+1} n$ may be written as $\delta(\mathcal{G}_{p_i}) > \frac{1+\epsilon_i}{2} p_i (n-1)$ for some negative $\epsilon_i = -\Theta(f/\log n)$. Hence, using (30) and also Lemma 19 with $\epsilon = \epsilon_i$ and $h(n) = -\beta \log \log n/2 - f^2 - if$, we get that, for every $1 \leq i < T$,

$$\begin{aligned} \Pr\left(\delta(\mathcal{G}_{p_i}) > \frac{\bar{d}(\mathcal{G}_{p_{i+1}})}{2}\right) &\leq \Pr\left(\delta(\mathcal{G}_{p_i}) > \frac{1-\sigma}{2} p_{i+1} n\right) + \Pr\left(\bar{d}(\mathcal{G}_{p_{i+1}}) < (1-\sigma) p_{i+1} n\right) \\ &= O\left(\frac{\log n}{n} + \exp\left(\frac{-f^2 - if}{\beta} + O(f)\right)\right) + o(n^{-1}) = O(f^{-i} + \log n/n). \end{aligned}$$

Similarly, we write $\delta(\mathcal{G}_{q_i}) \leq \frac{1+\sigma}{2} q_{i+1} n$ as $\delta(\mathcal{G}_{q_i}) \leq \frac{1+\epsilon'_i}{2} q_i (n-1)$ for some $\epsilon'_i = \Theta(f/\log n)$. Using again (30) and Lemma 19 with $\epsilon = \epsilon'_i$ and $h(n) = -\beta \log \log n/2 + f^2 + if$, we obtain, for every $1 \leq i < T'$,

$$\begin{aligned} \Pr\left(\delta(\mathcal{G}_{q_i}) \leq \frac{\bar{d}(\mathcal{G}_{q_{i+1}})}{2}\right) &\leq \Pr\left(\delta(\mathcal{G}_{q_i}) \leq \frac{1+\sigma}{2} q_{i+1} n\right) + \Pr\left(\bar{d}(\mathcal{G}_{q_{i+1}}) > (1+\sigma) q_{i+1} n\right) \\ &= O\left(\exp\left(\frac{-f^2 - if}{\beta} + O(f)\right)\right) + o(n^{-1}) = O(f^{-i} + 1/n). \quad \square \end{aligned}$$

Proof of Theorem 3. We first prove statement (ii). Let $p_1 = \frac{(1-\epsilon/2)\beta \log n}{n-1}$ and let $p_2 = \frac{(1+\epsilon/2)\beta \log n}{n-1}$. For $i = 1, 2$, the number of edges in $\mathcal{G}(n, p_i)$ is distributed as $\mathbf{Bin}\left(\binom{n}{2}, p_i\right)$. By Lemma 13, we have that a.a.s. $m(\mathcal{G}(n, p_1)) \geq \frac{1-\epsilon}{1-\epsilon/2} p_1 \binom{n}{2} = (1-\epsilon)\beta n \log n/2$, and $m(\mathcal{G}(n, p_2)) \leq \frac{1+\epsilon}{1+\epsilon/2} p_2 \binom{n}{2} = (1+\epsilon)\beta n \log n/2$. Then, Theorem 3 (ii) follows immediately from Theorem 27.

We now proceed to prove statement (i) of Theorem 3. Recall that for any graph G , $t(G) = \min\{\delta(G), \bar{d}(G)/2\}$. First, define $p_0 = 0.9 \log n/(n-1)$, $p_1 = \gamma_1 \log n/(n-1)$ and $p_2 = \gamma_2 \log n/(n-1)$, for some constants $1 < \gamma_1 < \gamma_2$ that we specify later. We prove the statement separately for $(\mathcal{G}_p)_{p_0 \leq p \leq p_1}$, $(\mathcal{G}_p)_{p_1 \leq p \leq p_2}$ and $(\mathcal{G}_p)_{p_2 \leq p \leq 1}$. For $(\mathcal{G}_p)_{0 \leq p \leq p_0}$ it is trivially true since a.a.s. $\delta(\mathcal{G}_p) = 0$ for all $0 \leq p \leq p_0$, by Lemma 20 (i) and the monotonicity of $\delta(\mathcal{G}_p)$ with respect to p .

Part 1 ($p_0 \leq p \leq p_1$): Let $\epsilon > 0$ be a constant chosen to satisfy Lemma 23 with $\gamma = 1.1$. Pick a sufficiently small constant $1 < \gamma_1 < 1.1$ and recall $p_0 = 0.9 \log n/(n-1)$ and $p_1 = \gamma_1 \log n/(n-1)$. From Lemma 13, a.a.s.

$$\bar{d}(\mathcal{G}_{p_1}) \leq (4/3)\bar{d}(\mathcal{G}_{p_0}). \quad (31)$$

Moreover, in view of Lemma 20 (ii), we assume that γ_1 is small enough so that a.a.s.

$$\delta(\mathcal{G}_{p_1}) \leq (\epsilon/16)\bar{d}(\mathcal{G}_{p_1}) \leq (\epsilon/12)\bar{d}(\mathcal{G}_{p_0}). \quad (32)$$

Colour edges in \mathcal{G}_{p_1} so that all edges in \mathcal{G}_{p_0} are coloured red and all edges in $\mathcal{G}_{p_1} \setminus \mathcal{G}_{p_0}$ are coloured blue, as described in the beginning of the section. For each vertex $v \in \mathcal{G}_{p_1}$, the red (blue) degree of v is the number of red (blue) edges incident with v . Let S be the set of ϵ -light vertices of \mathcal{G}_{p_0} . Since $\delta(\mathcal{G}_{p_0}) = 0$ a.a.s., the ϵ -light vertices are a.a.s. precisely those vertices with degree at most $\epsilon\bar{d}(\mathcal{G}_{p_0})$ in \mathcal{G}_{p_0} . By the choice of ϵ and Lemma 23, a.a.s. the vertices in S induce no edges and have no common neighbours in the whole process $(\mathcal{G}_p)_{p_0 \leq p \leq p_1}$ as blue edges are added.

For each $p_0 \leq p \leq p_1$, let S_p be the set of $(11\epsilon/16)$ -light vertices of \mathcal{G}_p (i.e. vertices of degree at most $\delta(\mathcal{G}_p) + (11/16)\epsilon\bar{d}(\mathcal{G}_p)$ in \mathcal{G}_p). Note that a.a.s. S contains S_p for all p in this range, since for any $v \in S_p$, its degree in \mathcal{G}_{p_0} (i.e. the red degree of v in \mathcal{G}_p) is at most

$$\delta(\mathcal{G}_p) + (11/16)\epsilon\bar{d}(\mathcal{G}_p) \leq \delta(\mathcal{G}_{p_1}) + (11/16)\epsilon\bar{d}(\mathcal{G}_{p_1}) \leq (\epsilon/12)\bar{d}(\mathcal{G}_{p_0}) + (11/12)\epsilon\bar{d}(\mathcal{G}_{p_0}) = \epsilon\bar{d}(\mathcal{G}_{p_0}),$$

where we used (31) and (32).

We just showed that a.a.s. in $(\mathcal{G}_p)_{p_0 \leq p \leq p_1}$ the set of $(11\epsilon/16)$ -light vertices of \mathcal{G}_p induce no edges and have no common neighbours. Moreover, from (32) and by monotonicity of $\delta(\mathcal{G}_p)$ and $\bar{d}(\mathcal{G}_p)$ with respect to p , we have that a.a.s.

$$\delta(\mathcal{G}_p) \leq \delta(\mathcal{G}_{p_1}) \leq (\epsilon/12)\bar{d}(\mathcal{G}_{p_0}) \leq (\epsilon/12)\bar{d}(\mathcal{G}_p)$$

in the whole process $(\mathcal{G}_p)_{p_0 \leq p \leq p_1}$. Putting all that together, we have that a.a.s. the conditions of Proposition 10 are satisfied in $(\mathcal{G}_p)_{p_0 \leq p \leq p_1}$ (replacing ϵ by $(11/16)\epsilon$), and therefore a.a.s. $T(\mathcal{G}_p) = \delta(\mathcal{G}_p)$ simultaneously for all p in this range.

Part 2 ($p_1 \leq p \leq p_2$): Recall that $p_1 = \gamma_1 \log n / (n - 1)$ and $p_2 = \gamma_2 \log n / (n - 1)$, where γ_1 is as in Part 1, and $\gamma_2 > \gamma_1$ is a sufficiently large constant. In view of Theorem 27 (iii), we assume that γ_2 is large enough so that a.a.s. $\delta(\mathcal{G}_p) \geq (3/4)\bar{d}(\mathcal{G}_p)$ in the whole process $(\mathcal{G}_p)_{p_2 \leq p \leq 1}$.

Define $q_i = (1 + 1/\log n)^i p_1$ for each $i = 0, 1, 2, \dots$, and let T be the smallest integer such that $q_T \geq p_2$. Redefine $q_T = p_2$. We have $T \leq 2 \log(\gamma_2/\gamma_1) \log n = O(\log n)$, since eventually

$$(1 + 1/\log n)^{2 \log(\gamma_2/\gamma_1) \log n} > \gamma_2/\gamma_1.$$

To prove the statement for $(\mathcal{G}_p)_{p_1 \leq p \leq p_2}$, it suffices to see that for every $0 \leq i \leq T - 1$, we have $T(\mathcal{G}_p) = \lfloor t(\mathcal{G}_p) \rfloor$ throughout the process $(\mathcal{G}_p)_{q_i \leq p \leq q_{i+1}}$ with probability at least $1 - 1/\log^2 n$, and then simply take a union bound over all i .

Let ϵ be as in Lemma 23 (putting $\gamma = \gamma_2$), and fix $0 \leq i \leq T - 1$. We verify that with probability at least $1 - 1/\log^2 n$ all conditions (a')–(e') of Proposition 11 are satisfied in $(\mathcal{G}_p)_{q_i \leq p \leq q_{i+1}}$. We colour as before the edges of \mathcal{G}_{q_i} red, and the additional edges in $\mathcal{G}_{q_{i+1}} \setminus \mathcal{G}_{q_i}$ blue.

Let S be the set of vertices that are ϵ -light in \mathcal{G}_{q_i} (they have red degree at most $\delta(\mathcal{G}_{q_i}) + \epsilon \bar{d}(\mathcal{G}_{q_i})$). For each $q_i \leq p \leq q_{i+1}$, define S_p to be the set of vertices that are $\epsilon/2$ -light in \mathcal{G}_p . From Lemma 13 and Lemma 21, we have that

$$\bar{d}(\mathcal{G}_{q_i}) \sim \bar{d}(\mathcal{G}_{q_{i+1}}) \quad \text{and} \quad \delta(\mathcal{G}_{q_{i+1}}) \leq (1 + \epsilon/3)\delta(\mathcal{G}_{q_i}) \quad (33)$$

with probability at least $1 - n^{-C}$, for some small enough constant $C > 0$ not depending on i . These equations imply that $S \supseteq S_p$ for all p in our range, since the red degree of any vertex in S_p is at most

$$\delta(\mathcal{G}_p) + (\epsilon/2)\bar{d}(\mathcal{G}_p) \leq \delta(\mathcal{G}_{q_{i+1}}) + (\epsilon/2)\bar{d}(\mathcal{G}_{q_{i+1}}) \leq (1 + \epsilon/3)\delta(\mathcal{G}_{q_i}) + (\epsilon/2 - o(1))\bar{d}(\mathcal{G}_{q_i}) \leq \delta(\mathcal{G}_{q_i}) + \epsilon \bar{d}(\mathcal{G}_{q_i}),$$

where we also used the trivial fact that $\delta(\mathcal{G}_{q_i}) \leq \bar{d}(\mathcal{G}_{q_i})$. By Lemma 23, with probability at least $1 - n^{-C}$ the vertices in S do not get common neighbours or induced edges as the blue edges are added in $(\mathcal{G}_p)_{q_i \leq p \leq q_{i+1}}$. This implies condition (a') replacing ϵ by $\epsilon/2$. By Lemma 20 (iii), there exists a constant $\sigma > 0$ such that, uniformly for all $p \in [p_1, p_2]$, we have that $\delta(\mathcal{G}_p) \geq \sigma p n$ with probability at least $1 - n^{-C}$. Therefore, $t(\mathcal{G}_{q_i}) \geq \sigma' q_{i+1} n$ with probability at least $1 - n^{-C}$ for a positive constant σ' not depending on i . By applying Lemma 25 to $\mathcal{G}_{q_{i+1}}$ with $\alpha < \epsilon \sigma'/4$, we deduce that condition (e') holds with probability at least $1 - n^{-C}$ during all the process $(\mathcal{G}_p)_{q_i \leq p \leq q_{i+1}}$. This determines our choice of ζ . Next, observe that \mathcal{G}_{q_i} satisfies condition (b') with probability at least $1 - e^{-C(\log n)^{1/3}}$, by Lemma 14, and also condition (c') with probability at least $1 - n^{-C}$, by Lemma 24. Therefore, in view of (33), both Conditions (b') and (c') hold simultaneously in all $(\mathcal{G}_p)_{q_i \leq p \leq q_{i+1}}$ with probability at least $1 - e^{-C(\log n)^{1/3}}$. Condition (d') holds trivially for sets S of size 1. For larger sets, Lemma 26 applied to \mathcal{G}_{q_i} together with (33) imply that this condition holds in all $(\mathcal{G}_p)_{q_i \leq p \leq q_{i+1}}$ with probability $1 - n^{-C}$.

Taking the union bound for all $0 \leq i \leq T - 1 = O(\log n)$, we have that a.a.s. $T(\mathcal{G}_p) = \lfloor t(\mathcal{G}_p) \rfloor$ throughout the process $(\mathcal{G}_p)_{p_1 \leq p \leq p_2}$ by Proposition 11.

Part 3 ($p_2 \leq p \leq 1$): Let γ_2 be as in Part 2, and $p_2 = \gamma_2 \log n / (n - 1)$. Recall from the definition of γ_2 that a.a.s. $\delta(\mathcal{G}_p) \geq (3/4)\bar{d}(\mathcal{G}_p)$ in the whole process $(\mathcal{G}_p)_{p_2 \leq p \leq 1}$, and therefore Condition (a') in Proposition 11 holds. Define $q_i = (1 + 1/\log n)^i p_2$ for each $i = 0, 1, 2, \dots$, and let T be the smallest integer such that $q_T \geq 1$. Redefine $q_T = 1$. Observe that $T \leq 3 \log^2 n$, since eventually $(1 + 1/\log n)^{3 \log^2 n} \geq n^{2.5}$. The same argument as in Part 2 shows that for every $0 \leq i \leq T - 1$, Conditions (b')–(e') in Proposition 11 are satisfied throughout the process $(\mathcal{G}_p)_{q_i \leq p \leq q_{i+1}}$ with probability at least $1 - 1/\log^3 n$. Taking the union bound over all i , we conclude that a.a.s. all condition in Proposition 11 hold and therefore $T(\mathcal{G}_p) = \lfloor \bar{d}(\mathcal{G}_p)/2 \rfloor$, during the whole process $(\mathcal{G}_p)_{p_2 \leq p \leq 1}$. \square

8 Proof of Theorem 5

We first prove statement (ii). In view of Theorem 3, we assume that $T(G_m) = \min\{\delta(G_m), \lfloor m/(n-1) \rfloor\}$ for all $m = 0, 1, \dots, \binom{n}{2}$. Then we pick any m such that $\delta(G_m) \geq \bar{d}(G_m)/2 = m/(n-1)$. If $n-1$ divides m , then $T(G_m) = m/(n-1)$ and thus $A(G_m) = m/(n-1)$. If $n-1$ does not divide m , let m' be the smallest integer $m' > m$ divisible by $n-1$. Since, the minimum degree is always an integer, we have $\delta(G_m) \geq \lceil m/(n-1) \rceil = m'/(n-1)$. Moreover, G_m is a spanning subgraph of $G_{m'}$ and thus $\delta(G_{m'}) \geq \delta(G_m) \geq m'/(n-1)$. Therefore, from our assumption on the random graph process, we have $T(G_{m'}) = m'/(n-1) = \lceil m/(n-1) \rceil$, and these $\lceil m/(n-1) \rceil$ edge-disjoint spanning trees cover all edges of G_m , so $A(G_m) = \lceil m/(n-1) \rceil$. This completes the proof of the statement.

Next we proceed to prove statement (i). Let f be any function of n such that $f \rightarrow \infty$ arbitrarily slowly and $f = o(\log n)$. Define $p_j = (1 + 1/f)^j f/n$ for each $j = 0, 1, 2, \dots$, and let T be the largest integer such that $p_T \leq \frac{\beta \log n}{(1-\epsilon/2)n}$.

Claim 3. for every $0 \leq j < T$ and for every m such that $p_j \leq p(m) \leq p_{j+1}$, the bound $A(\mathcal{G}_{p(m)}) = A(G_m) \leq \lceil \frac{m+\phi_2}{n-1} \rceil$ holds in the random graph process $(\mathcal{G}_p)_{p \in [p_j, p_{j+1}]}$ with probability at least $1 - 1/(p_j n)^2$.

Assuming Claim 3, the probability that $A(G_m) \leq \lceil \frac{m+\phi_2}{n-1} \rceil$ fails somewhere in the random graph process between $\mathcal{G}(n, p_0)$ and $\mathcal{G}(n, p_T)$ is at most

$$\frac{1}{f^2} \sum_{j=0}^{T-1} (1 + 1/f)^{-2j} \leq \frac{1}{f^2} \sum_{j=0}^{\infty} (1 + 1/f)^{-j} = \frac{1 + 1/f}{f} = o(1). \quad (34)$$

Moreover, we have that $p_T \geq \frac{\beta \log n}{(1-\epsilon/4)n}$ eventually (for $n > n_0$ depending only on f). Then there is $\sigma > 0$ such that a.a.s. $m > (1 + \sigma)\beta n \log n$ for every m with $p(m) \geq p_T$. Then, by Theorem 3 (ii), a.a.s. $\delta(G_m) > \bar{d}(G_m)/2$ for all m with $p(m) \geq p_T$. Thus, we only need to restrict our discussion to p_j with $j \leq T$.

Similarly, let T' be the largest integer such that $p_{T'} \leq \frac{\beta \log n}{(1+\epsilon/2)n}$.

Claim 4. for every $0 \leq j < T'$ and for every m such that $p(m) \in [p_j, p_{j+1}]$, the bound $A(G_m) \geq \lceil \frac{m+\phi_1}{n-1} \rceil$ holds in the random graph process $(\mathcal{G}_p)_{p \in [p_j, p_{j+1}]}$ with probability at least $1 - 1/(p_j n)^2$.

By the same argument as in (34), assuming Claim 2, the probability that $A(G_m) \geq \lceil \frac{m+\phi_1}{n-1} \rceil$ fails somewhere in the random graph process $(\mathcal{G}_p)_{p \in [p_0, p_{T'}]}$ is also $o(1)$. Moreover,

note that since $p_{T'} \geq \frac{\beta \log n}{(1+2\epsilon/3)n}$, then a.a.s. for every m with $p(m) \in [p_{T'}, 1]$, we have $2m/n \geq \frac{\beta \log n}{(1+3\epsilon/4)}$ by Lemma 13, and therefore eventually $\phi_1 \leq 1/2$ for all m in this range. Hence, for all m not divisible by $n-1$, we eventually have

$$A(G_m) \geq \left\lceil \frac{m}{n-1} \right\rceil = \left\lceil \frac{m+1/2}{n-1} \right\rceil = \left\lceil \frac{m+\phi_1}{n-1} \right\rceil.$$

Otherwise, for m divisible by $n-1$, the condition $\delta(G_m) < \bar{d}(G_m)/2 = m/(n-1)$ implies $A(G_m) > m/(n-1)$, since we cannot have a full factorisation of G_m into $m/(n-1)$ spanning trees, so then

$$A(G_m) \geq \frac{m}{n-1} + 1 = \left\lceil \frac{m+1/2}{n-1} \right\rceil = \left\lceil \frac{m+\phi_1}{n-1} \right\rceil.$$

Putting everything together, we showed that a.a.s. (1) holds simultaneously for all G_m in the random graph process \mathcal{G}_p for p between f/n and 1. Given any $m_0 = \omega(n)$ as in the statement, we may simply choose $f = m_0/n$. Then a.a.s. $m(\mathcal{G}_{f/n}) = m(\mathcal{G}(n, f/n)) \leq (3/4)fn = (3/4)m_0$, and statement (a) holds for the desired range of m . It only remains to prove Claims 3 and 4.

Proof of Claim 3. In this proof the asymptotic statements are uniform for all p_j and depend only on f . Given any $0 \leq j < T$, we consider the random graph process $(\mathcal{G}_p)_{p \in [p_j, p_{j+1}]}$. Define $\delta_j = \delta(\mathcal{G}_{p_j})$, $\bar{d}_j = \bar{d}(\mathcal{G}_{p_j})$, $m_j = m(\mathcal{G}_{p_j})$, $t_j = \min\{\delta_j, \bar{d}_j/2\}$. By Lemma 13, we have $\bar{d}_j \sim \bar{d}_{j+1} \sim np_j$ with probability at least $1 - Ce^{-p_j n}$ for a positive constant C .

Let $\hat{\epsilon} > 0$ be a sufficiently small constant so that

$$e^{-1} \left(\frac{2e}{1+2\hat{\epsilon}} \right)^{(1+2\hat{\epsilon})/2} < e^{-(1-\epsilon)/\beta} \quad \text{and} \quad e^{-1} \left(\frac{2e}{1-\hat{\epsilon}} \right)^{(1-\hat{\epsilon})/2} > e^{-(1+\epsilon/4)/\beta}. \quad (35)$$

Colour the edges of $G_{m_j} = \mathcal{G}_{p_j}$ red. Let $G_{m_{j+1}} = \mathcal{G}_{p_{j+1}}$. Colour edges in $G_{m_{j+1}} \setminus G_{m_j}$ blue. For any vertex $v \in G_{m_{j+1}}$, define the red (blue) degree of v to be the number of red (blue) edges that are incident with v . Call a vertex light if its red degree is at most $\frac{(1+2\hat{\epsilon})}{2}\bar{d}_{j+1}$. A vertex is called heavy if its red degree is at least $(3/4)\bar{d}_{j+1}$. The vertices that are neither light nor heavy are called medium vertices. We have that for any constant $\alpha > 0$, with probability $1 - e^{-p_j n}$, by Lemma 13 and Lemma 15,

$$\left| m_{j+1} - p_{j+1} \binom{n}{2} \right| < \alpha p_{j+1} n^2, \quad \Delta(G_{m_{j+1}}) = O(\log n), \quad (36)$$

where $\Delta(G)$ denotes the maximum degree of G . By Lemma 16 with $k = \frac{1+2\hat{\epsilon}}{2}\bar{d}_{j+1}$, the expected number of light vertices is at most

$$n \left(\frac{1}{e} \left(\frac{2e}{1+2\hat{\epsilon}} \right)^{\frac{1+2\hat{\epsilon}}{2}} + o(1) \right)^{pn}.$$

Thus by Markov's inequality and (35), the number of light vertices is

$$\ell \leq \frac{n}{\bar{d}_{j+1} \exp\left(\frac{(1-\epsilon)2m_{j+1}}{\beta n}\right)} = o(n), \quad (37)$$

with probability at least $1 - e^{-Dp_j n}$, for a positive constant D . Similarly, by Lemma 16 with $k = \frac{3}{4}\bar{d}_{j+1}$, with probability $1 - e^{-Dp_j n}$ for a positive constant D , the number of heavy vertices is

$$h = n - o(n). \quad (38)$$

For the following construction, assume (36), (37) and (38) hold. We add $g = \bar{d}_{j+1}\ell$ new edges (different from the previous red and blue edges) to $G_{m_{j+1}}$, which we colour green, in such a way that every light vertex is incident with exactly \bar{d}_{j+1} green edges; every heavy vertex is incident to at most one green edge; and no green edge is incident to any medium vertices. (So green edges only connect light and heavy vertices.) This can be done by (37), since the total number of green edges

$$g \leq \frac{n}{\exp\left(\frac{(1-\epsilon)2m_{j+1}}{\beta n}\right)}$$

is much smaller than the number of heavy vertices eventually. Finally, greedily add n yellow edges to $G_{m_{j+1}}$, different to all previous red, blue and green ones, in a way that each yellow edge connects two heavy vertices and each heavy vertex is incident with at most 3 yellow edges (this can be done greedily since we have $h \sim n$ heavy vertices by (38) and the maximum degree (adding red, blue and green degrees together) is $O(\log n)$ by (36)).

We may regard the sequence of graphs $G_{m_j} \subset G_{m_{j+1}}, \dots \subset G_{m_{j+1}}$ as a process in which we sequentially add blue edges to G_{m_j} , so the edges of each G_m are precisely the red ones together with the first $m - m_j$ blue ones. For each m in our range, we define G'_m as $E(G'_m) = G_m \cup E_g \cup E_y$, where E_g is the set of green edges added to $G_{m_{j+1}}$ and E_y is an arbitrary subset of yellow edges added to $G_{m_{j+1}}$ so that the number of edges of the resulting graph G'_m is a multiple of $n - 1$. We now verify that $G'_{m_j}, \dots, G'_{m_{j+1}}$ satisfy all conditions (a')–(e') of Proposition 11, assuming that (36), (37) and (38) and some additional events hold. We give bounds on the probabilities of these events.

First observe that for all $m_j \leq m \leq m_{j+1}$, we have $\bar{d}(G'_m) \leq \frac{2(m_{j+1}+g+n)}{n-1} \leq \bar{d}_{j+1} + 2$ and similarly $\bar{d}(G'_m) \geq \frac{2m_j}{n-1} = \bar{d}_j$, so

$$\bar{d}(G'_m) \sim \bar{d}_{j+1} \sim \bar{d}_j \quad (39)$$

Hence, for all m in the range,

$$\delta(G'_m) \geq \frac{(1 + 2\hat{\epsilon})\bar{d}_{j+1}}{2} \geq \frac{(1 + \hat{\epsilon})\bar{d}(G'_m)}{2}, \quad (40)$$

so (a') holds and $\delta = \Omega(\bar{d}) = \omega(1)$. For any S , let $d_r(S)$ denote the average red degree of S . By Lemma 14, with probability $1 - e^{-C(p_j n)^{1/3}}$ for a positive constant C , for all S with $S \geq \zeta n$, $d_r(S) \geq \bar{d}(G_{m_j})(1 - o(1)) \geq \bar{d}(G'_m)(1 - o(1))$ by (39). Thus, (b') holds by noting that $d(S) \geq d_r(S)$. By applying Lemma 24 to the red edges and using (39), we deduce condition (c') holds with probability at least $1 - e^{-Cp_j n^2}$ for all G'_m .

For (e'), first note that $t(G'_{m_j}) = \Omega(\bar{d}_{j+1})$. Then, Lemma 25 applied to $G_{m_{j+1}}$ (i.e. only red and blue edges) shows that with probability $1 - Ce^{-(p_j n)^2}$ all sets S of size $s < \zeta n$ have at most $(\hat{\epsilon}/8)t(G'_{m_j})s$ red and blue edges inside. Let G'' be obtained by adding all g green edges and all n yellow edges. So $G'_m \subseteq G''$ for all $m_j \leq m \leq m_{j+1}$. We bound the number

of edges induced by S in G'' . Let $s_2 \leq s$ be the number of heavy vertices in S . Each green edge induced by S must be incident to one of the s_2 heavy vertices inside, and the number of yellow edges induced by S is at most $3s_2$, since each heavy vertex is incident to at most 3 of them. Therefore, the number of green and yellow edges induced by S is at most $4s_2 \leq 4s$, and the total number of edges induced by S in G'' (and thus, in all G'_m , $m_j \leq m \leq m_{j+1}$) is at most $(\hat{\epsilon}/8)t(G'_{m_j})s + 4s \leq (\hat{\epsilon}/4)t(G'_{m_j})s$, since $t(G'_{m_j}) \rightarrow \infty$. Thus (e') holds for all $m_j \leq m \leq m_{j+1}$, since $t(G'_{m_j}) \leq t(G'_m)$.

Finally, we prove (d'): Let S be a set with $1 \leq |S| \leq n/2$ (otherwise we take \bar{S}). Suppose first that $1 \leq |S| \leq \zeta n$. From what we proved before for (e') and (40), the number of edges induced by S is at most $(\hat{\epsilon}/4)t(G'_m)s$ and so, for each G'_m ,

$$E(S, \bar{S}) \geq \delta(G'_m)s - (\hat{\epsilon}/4)t(G'_m)s \geq \frac{(1 + \hat{\epsilon})\bar{d}(G'_{m_j})}{2}s - (\hat{\epsilon}/4)t(G'_{m_j})s \geq t(G'_{m_j})s \geq t(G'_m)s.$$

Otherwise, if $\zeta n \leq |S| \leq n/2$, (c') gives us what we need using only red edges.

Hence, in view of Proposition 11, with probability $1 - e^{-C(p_j n)^{1/3}} \geq 1 - 1/(p_j n)^2$ since $p_j n \geq f = \omega(1)$, for all $m_j \leq m \leq m_{j+1}$, we have

$$T(G'_m) = m(G'_m)/(n-1) = \lceil (m+g)/(n-1) \rceil \leq \lceil (m+\phi_2)/(n-1) \rceil,$$

since by construction $m(G'_m)$ is the smallest integer that is at least $m+g$ and is divisible by $n-1$.

This implies the claim since

$$A(G_m) \leq A(G'_m) = m(G'_m)/(n-1) \leq \lceil (m+\phi_2)/(n-1) \rceil. \quad \square$$

Proof of Claim 4. We pick a constant $\hat{\epsilon} > 0$ as in (35). By (36) and the definition of p_j , with probability $1 - e^{-p_j n}$, for all $m_j \leq m \leq m_{j+1}$ we have

$$A(G_m) \geq \lceil m/(n-1) \rceil \geq \bar{d}(G_{m_j})/2 > (1 - \hat{\epsilon})\bar{d}(G_{m_{j+1}})/2 \quad (41)$$

and the maximum degree of G_m is $O(\log n)$. Let us redefine light vertices of $G_{m_{j+1}}$ to be vertices with degree (red degree plus blue degree) at most $(1 - \hat{\epsilon})\bar{d}(G_{m_{j+1}})/2$. By Lemma 16 with $k = \frac{(1-\hat{\epsilon})\bar{d}(G_{m_{j+1}})}{2}$ and (35), the expected number of light vertices in $G_{m_{j+1}}$ is at least

$$n\sqrt{\frac{1}{k}} \left(\frac{1}{e} \left(\frac{2e}{(1-\hat{\epsilon})} \right)^{\frac{1-\hat{\epsilon}}{2}} + o(1) \right)^{pn} \geq ne^{-(1+\epsilon'/4)pn/\beta} \geq 2ne^{-\frac{(1+\epsilon/4)2m_j}{\beta} \frac{2m_j}{n}},$$

for some constant $\epsilon' < \epsilon$ and the inequality holds since $f = o(\log n)$. Note that with probability at least $1 - n^{-C}$, for all $j \leq T'$, $2m_j \leq (1 + \epsilon/8)\beta n \log n$ and so

$$\frac{(1 + \epsilon/4)2m_j}{\beta} \frac{2m_j}{n} \leq \frac{(1 + \epsilon/4)(1 + \epsilon/8)}{1 + \epsilon/2} \log n \leq \sigma \log n,$$

for some $0 < \sigma < 1$, depending only on ϵ . Then, by Lemma 17 and by Chebyshev's inequality, there are

$$\ell' \geq \frac{n}{\exp\left(\frac{(1+\epsilon/4)2m_j}{\beta} \frac{2m_j}{n}\right)}$$

light vertices in $G_{m_{j+1}}$ with probability at least $1 - O(\log n/n) - n^{\sigma-1} \geq 1 - 1/(p_j n)^2$. By (41), with probability at least $1 - e^{-p_j n}$, these light vertices have degree in G_m strictly less than $A(G_m)$ for all m in the range $m_j \leq m \leq m_{j+1}$ since $G_m \subseteq G_{m_{j+1}}$.

For each G_m we construct G'_m as follows. Let \mathcal{F}_m be the set of $A(G_m)$ edge-disjoint forests covering G_m . For every light vertex v and for every forest $F \in \mathcal{F}_m$, if F has no edge incident to v , then we add a new edge connecting v to some non-light vertex, and make this new edge be part of F (this can always be done since both $|\mathcal{F}_m|$ and the maximum degree are $O(\log n)$ and the number of non-light vertices is $n - o(n)$). Observe that G'_m has at least $m + \ell'$ edges since for each light vertex v we added at least one edge as the degree of v is less than $|\mathcal{F}_m| = A(G_m)$. By construction, G'_m and G_m have the same arboricity, so

$$A(G_m) = A(G'_m) \geq \lceil (m + \ell') / (n - 1) \rceil \geq \lceil (m + \phi_1) / (n - 1) \rceil. \quad \square$$

9 Proof of Theorem 4

The main task in this section is to prove part (ii) of Theorem 4. We will need the following lemma bounding the average degree of the subgraphs of $\mathcal{G}(n, p)$ that are not too large.

Lemma 28. *Let $G \sim \mathcal{G}(n, p)$, where $p \sim c/n$ for some constant $c > 0$.*

- (a) *If $c \leq 5$ then for every $\epsilon > 0$, there exists another constant $\alpha = \alpha(\epsilon) > 0$, such that a.a.s. all subgraphs of G with order at most αn have average degree at most $2 + \epsilon$.*
- (b) *There exists an absolute constant $\alpha > 0$ such that for all $c \geq 5$, a.a.s. all subgraphs of G with order at most αn have average degree at most $c/2$.*

Proof. For part (a): let X_s denote the number of sets $S \subseteq [n]$ with $|S| = s$ and $|E[S]| > (1 + \epsilon/2)s$. Let $r = s/n$. Then

$$\begin{aligned} \mathbf{E}X_s &\leq \binom{n}{s} \binom{s^2}{\lceil (1 + \epsilon/2)s \rceil} \left(\frac{c + o(1)}{n} \right)^{(1 + \epsilon/2)s} \\ &\leq \left(\frac{e}{r} \left(\frac{ern}{(1 + \epsilon/2)} \frac{c + o(1)}{n} \right)^{(1 + \epsilon/2)s} \right) \leq ((C + o(1))r^{\epsilon/2})^s, \end{aligned} \quad (42)$$

where $C = e^{2 + \epsilon/2} 5^{1 + \epsilon/2}$ is a constant depending only on ϵ . Thus, by choosing $\alpha = \alpha(\epsilon)$ sufficiently small, we have that for all $r \leq \alpha$, $Cr^{\epsilon/2} < 1/2$. It follows then that

$$\sum_{1 \leq s \leq \alpha n} \mathbf{E}X_s = \sum_{1 \leq s \leq \log n} \mathbf{E}X_s + \sum_{\log n < s \leq \alpha n} \mathbf{E}X_s = O\left(\frac{\log^{\epsilon/2} n}{n^{\epsilon/2}} + 2^{-\log n} \right) = o(1).$$

Part (b) follows by a similar calculation, with $1 + \epsilon/2$ replaced by $c/4$. Then, (42) becomes

$$\mathbf{E}X_s \leq ((1 + o(1))4e^2(4er)^{c/4-1})^s < 2^{-s},$$

by choosing α such that $4e^2(4e\alpha)^{1/4} < 1/3$. Then, $\sum_{1 \leq s \leq \alpha n} \mathbf{E}X_s = o(1)$, by the same calculation as in part (a). \square

Much of the argument that follows relies on the analysis of the structure of the k -core of $\mathcal{G}(n, p)$, which is the largest subgraph of $\mathcal{G}(n, p)$ with minimum degree at least k . We require some definitions. For any integer $k \geq 0$ and real $\mu > 0$, let

$$f_k(\mu) = e^{-\mu} \sum_{i \geq k} \frac{\mu^i}{i!}.$$

For any $k \geq 2$ and $\mu > 0$, define $h_k(\mu) = \frac{\mu}{f_{k-1}(\mu)}$, and let

$$c_k = \inf\{h_k(\mu) : \mu > 0\}. \quad (43)$$

Claim 5.

- (i) For every $k \geq 2$, c_k is a real. Moreover, $1 = c_2 < c_3 < \dots < c_k < c_{k+1} < \dots$ and $\lim_{k \rightarrow \infty} c_k = \infty$.
- (ii) For $k \geq 3$, there is a unique $\mu_k > 0$ where h_k attains a global minimum, so in particular $c_k = h(\mu_k) = \min\{h_k(\mu) : \mu > 0\}$. For $k = 2$, we have $c_2 = 1 = \lim_{\mu \rightarrow 0} h_2(\mu)$, and we define $\mu_2 = 0$ for convenience.
- (iii) For every $k \geq 3$, $h_k(\mu)$ is a strictly increasing continuous function in (μ_k, ∞) , and $\lim_{\mu \rightarrow \infty} h_k(\mu) = \infty$.

In view of this claim, for every $k \geq 2$ and $c > c_k$, we can define $\mu_{c,k} > \mu_k$ to be the unique solution of $h_k(\mu) = c$ in (μ_k, ∞) .

Claim 5 and the following theorem follow from a result about the threshold for the appearance of a giant k -core, first proved by Pittel, Spencer and Wormald [38], and later re-proved by many authors (see [28, 26, 33]).

Theorem 29. *Let $k \geq 2$ be fixed and let c_k be defined as in (43). Then for all $c > c_k$, a.a.s. $\mathcal{G}(n, c/n)$ has a non-empty k -core with $f_k(\mu_{c,k})n + o(n)$ vertices and $\frac{1}{2}\mu_{c,k}f_{k-1}(\mu_{c,k})n + o(n)$ edges. For all $k \geq 3$ and $c < c_k$, a.a.s. $\mathcal{G}(n, c/n)$ has an empty k -core.*

We are particularly concerned about the situation in which the average degree of the k -core of $\mathcal{G}(n, p)$ is around $2k - 2$. So we need to ensure that the k -core has ‘‘at its birth’’ an average degree strictly smaller than $2k - 2$. This is often assumed as true in the literature; we include a proof for completeness in the Appendix.

Lemma 30. *Let $k \geq 3$, and define $\xi_k(\mu) = \mu f_{k-1}(\mu)/f_k(\mu)$ for $\mu > 0$. Then, $\xi_k(\mu)$ is a strictly increasing continuous function with $\lim_{\mu \rightarrow \infty} \xi(\mu) = \infty$. Moreover, $\xi_k(\mu_k) < 2k - 2$.*

Remark. Note that $\xi_k(\mu_k)$ is approximately the average degree of the k -core of $\mathcal{G}(n, p)$ at its birth, by Theorem 29.

In view of Lemma 30, for every $k \geq 3$, we define $\bar{\mu}_k$ to be the only positive solution of $\xi_k(\bar{\mu}_k) = 2k - 2$. We also define $\bar{c}_k = h_k(\bar{\mu}_k)$. Observe that Lemma 30 also implies that $\mu_k < \bar{\mu}_k$ and, since $h_k(\mu)$ is strictly increasing for $\mu > \mu_k$, we have that

$$c_k < \bar{c}_k. \quad (44)$$

Cain, Sanders and Wormald [5] proved that for every $k \geq 2$ and $\epsilon > 0$, a.a.s. we have that if the average degree of the $(k + 1)$ -core of $\mathcal{G}(n, p)$ is at most $2k - \epsilon$, then $\mathcal{G}(n, p)$ is k -orientable: all its edges can be oriented so that no vertex has indegree more than k . This implies that $c \sim \bar{c}_k$ is a sharp threshold for $\mathcal{G}(n, c/n)$ to be $(k - 1)$ -orientable, as it is obvious that if the average degree of the k -core exceeds $2(k - 1)$ then the graph cannot be $(k - 1)$ -oriented. On the other hand, Hakimi's characterisation [24] tells that a graph is $(k - 1)$ -orientable if and only if it contains no subgraph whose average degree is more than $2(k - 1)$. These two results immediately imply the following theorem.

Theorem 31. *Given any positive integer $k \geq 3$, $c \sim \bar{c}_k$ is a sharp threshold of $\mathcal{G}(n, c/n)$ for containing a subgraph with average degree at most $2(k - 1)$.*

Let $d^*(G)$ denote the average degree of a densest subgraph of G . Theorem 31 says that \bar{c}_k is the sharp threshold for $d^*(\mathcal{G}(n, c/n)) \leq 2(k - 1)$. Next we prove that, roughly speaking, $d^*(\mathcal{G}(n, c/n))$ is a continuous function of c .

Lemma 32. *Given $\sigma > 0$, there are constants $\delta_1, \delta_2 > 0$ such that the following holds for any constants $\epsilon > 0$ and $c \geq 1 + \sigma$. Suppose that $p \sim c/n$ and $p' = p + \epsilon/n$. Assume that $\mathcal{G}(n, p) \subseteq \mathcal{G}(n, p')$ are coupled in a way that $\mathcal{G}(n, p')$ is generated by including each edge not in $\mathcal{G}(n, p)$ independently with probability $(p' - p)/(1 - p)$. Then, a.a.s.*

$$d^*(\mathcal{G}(n, p)) + \delta_1\epsilon \leq d^*(\mathcal{G}(n, p')) \leq d^*(\mathcal{G}(n, p)) + \delta_2\epsilon.$$

Proof. By Theorem 29 and since $c \geq 1 + \sigma = c_2 + \sigma$, there is a δ (depending only on σ) such that a.a.s. the 2-core of $\mathcal{G}(n, p)$ has average degree at least $2 + \delta$. It is easy to show that the average degree of $\mathcal{G}(n, p)$ is a.a.s. at least $c/2$ by the Chernoff bound. Hence, a.a.s. the densest subgraph of $\mathcal{G}(n, p)$ has average degree at least $\max\{2 + \delta, c/2\}$. By Lemma 28(a,b), there is a constant $\alpha = \alpha(\sigma) > 0$ (as δ depends only on σ) such that a.a.s. the densest subgraph of $\mathcal{G}(n, p)$ must have order at least αn .

Let H be a densest subgraph of $\mathcal{G}(n, p)$, and let $S = V(H)$. Then, $d^*(\mathcal{G}(n, p))$ is the average degree of H . Clearly, H is the subgraph of $\mathcal{G}(n, p)$ induced by S . Denote by H' the subgraph of $\mathcal{G}(n, p')$ induced by S , and let d' denote the average degree of H' . Let $s = |S|$. We have shown that a.a.s. $s \geq \alpha n$. Conditioning on that, the number of edges in K_S (the complete graph on vertex set S) that are not in H is at least As^2 for some absolute constant $A > 0$. Each of these at least As^2 edges are present in $\mathcal{G}(n, p')$ with probability $(p' - p)/(1 - p) \sim \epsilon/n$, which is at least $\epsilon/2n$ for all large n . Hence, by Theorem 12 (the Chernoff bound), a.a.s. $d' - d^*(\mathcal{G}(n, p)) \geq \delta_1\epsilon$, by setting $\delta_1 = A\alpha/4$. This implies that a.a.s. $d^*(\mathcal{G}(n, p)) + \delta_1\epsilon \leq d^*(\mathcal{G}(n, p'))$, since $d^*(\mathcal{G}(n, p')) \geq d'$.

Similarly, consider a densest subgraph H'_2 of $\mathcal{G}(n, p')$, and let H_2 be the subgraph of $\mathcal{G}(n, p)$ induced by $V(H'_2)$. A similar argument as before yields that a.a.s. $\mathcal{G}(n, p')$ has at most ϵn more edges than $\mathcal{G}(n, p)$. Note that removing a set of at most ϵn edges from $\mathcal{G}(n, p')$ can decrease the average degree of H'_2 by at most $2\epsilon/\alpha$, since $|V(H'_2)| \geq \alpha n$ a.a.s. by Lemma 28. Therefore, $d^*(\mathcal{G}(n, p))$ is at least the average degree of H_2 , which is a.a.s. at least $d^*(\mathcal{G}(n, p')) - \delta_2\epsilon$, by setting $\delta_2 = 2/\alpha$. This implies the second inequality in the lemma. (Note that both δ_1 and δ_2 depend only on α , and thus depend only on σ , as desired.) \square

The following lemma follows directly from Theorem 31, Lemma 32, Claim 5 and (44).

Lemma 33. For every $k \geq 3$, $\bar{c}_k < \bar{c}_{k+1}$. Moreover, $\lim_{k \rightarrow \infty} \bar{c}_k = \infty$.

Corollary 34. Let $\mathcal{J} = \{\bar{c}_k : k \geq 3\}$. For every constant $c \in \mathbb{R}_{\geq 1} \setminus \mathcal{J}$ and every $p \sim c/n$, $A(\mathcal{G}(n, p)) = k$ a.a.s., where $k \geq 2$ is the smallest integer such that $c < \bar{c}_{k+1}$. Moreover, for every $c = \bar{c}_k \in \mathcal{J}$ and every $p \sim c/n$, $A(\mathcal{G}(n, p)) \in \{k-1, k\}$ a.a.s.

Proof. Let $c \geq 1$ be an arbitrary real number not in \mathcal{J} . From Lemma 33 and since $1 < \bar{c}_3 < \bar{c}_3$, we must have that either $1 \leq c < \bar{c}_3$ or there is an integer $k \geq 3$ for which $\bar{c}_k < c < \bar{c}_{k+1}$. Let $\epsilon = (\bar{c}_{k+1} - c)/2$ (in the former case we set $k = 2$) and define $c' = c + \epsilon$. Then $c < c' < \bar{c}_{k+1}$.

Let $p \sim c/n$ and set $p' = c'/n$. By the definition of $c' < \bar{c}_{k+1}$, and by Theorem 31, a.a.s. $d^*(\mathcal{G}(n, p')) \leq 2k$. Then by Lemma 32, a.a.s. $d^*(\mathcal{G}(n, p)) \leq 2k - \delta$ for some constant $\delta > 0$. By Theorem 9, a.a.s. $A(\mathcal{G}(n, p)) \leq k$. On the other hand, we have either $c > \bar{c}_k$ if $k \geq 3$ or $1 \leq c < \bar{c}_3$ if $k = 2$. In the first case, a.a.s. the k -core of $\mathcal{G}(n, p)$ has average degree strictly greater than $2(k-1)$ by the definition of \bar{c}_k and Theorem 29, and hence by Theorem 9, a.a.s. $A(\mathcal{G}(n, p)) \geq k$. It follows then that a.a.s. $A(\mathcal{G}(n, p)) = k$. In the second case, it is well known that a.a.s. $\mathcal{G}(n, p)$ contains a cycle when $c \geq 1$ and thus $A(\mathcal{G}(n, p)) \geq 2 = k$. It follows immediately that a.a.s. $A(\mathcal{G}(n, p)) = 2$ in this case.

The second statement of the corollary concerning $c = \bar{c}_k \in \mathcal{J}$, follows immediately from a standard coupling since the arboricity is a non-decreasing parameter as we add more edges to the graph. \square

Proof of Theorem 4. Part (i) follows as a corollary of Theorem 5. For part (ii), we consider $p = \Theta(1/n)$. Assume $p \leq c/n$ for some constant $c > 0$. Then there exists a constant $0 < \sigma_c < 1$ such that a.a.s. the number of vertices in $G \sim \mathcal{G}(n, p)$ with degree 0 is at least $\sigma_c n$, whereas a.a.s. $m(G) = (1 + o(1))p \binom{n}{2}$. Hence, every forest contained in G has at most $(1 - \sigma_c)n$ edges. It follows then that $A(G) \geq m(G)/(1 - \sigma_c)n \geq \frac{1}{1 - \sigma_c}pn/2$ as required. Next, we will prove that for all $p = \Theta(1/n)$, a.a.s. $A(\mathcal{G}(n, p))$ is concentrated on two bounded values. This directly implies that $A(\mathcal{G}(n, p))$ is bounded and thus, for every $p = \Theta(1/n)$, a.a.s. $A(\mathcal{G}(n, p)) = (1 + \Theta(1))pn/2$.

We proceed to prove that $A(\mathcal{G}(n, p))$ is a.a.s. concentrated on two values. We assume $p \sim c/n$ for some constant $c > 0$, and our claim holds for all $p = \Theta(1/n)$ by the subsubsequence principle (see [27]). We consider two cases. If $c \leq 1$, then all vertices of $\mathcal{G}(n, p)$ with $p \sim c/n$ are contained in isolated trees except a set S of $o(n)$ vertices [13, Theorem 4b]. By Lemma 28, a.a.s. any subset $S' \subseteq S$ contains at most $1.1|S'|$ edges. Hence, a.a.s. there is no subgraph H of $\mathcal{G}(n, p)$, such that $|E(H)|/(|H| - 1) > 2$. It follows then that $A(\mathcal{G}(n, p)) \in \{1, 2\}$ in this case. Otherwise, the a.a.s. two-value concentration of $A(\mathcal{G}(n, p))$ for all $c \geq 1$, and the a.a.s. one-value concentration of $A(\mathcal{G}(n, p))$ for all $c \in \mathbb{R}_{\geq 1} \setminus \mathcal{J}$ follow from Corollary 34.

Finally, when $p = o(1/n)$, we have that $\mathcal{G}(n, p)$ is a.a.s. acyclic by [13]. So a.a.s. $A(\mathcal{G}(n, p)) \leq 1$. \square

10 Extension to random hypergraphs

The Nash-Williams characterisations of the STP number of the arboricity of graphs generalise to hypergraphs. Of course for hypergraphs we need to extend the concepts of connectivity, spanning trees and forests properly. This has been carefully investigated in [17]; in particular,

Theorems 8 and 9 hold also for hypergraphs by using the concepts of hyperforests and spanning hypertrees in [17]. See [17, Theorems 2.8 and 2.10] for details. It is easy to work out that our deterministic propositions 10 and 11 can be generalised (by some slight adjustment of the constants involved in the statements and minor modifications of the proofs) for h -uniform hypergraphs, for any constant $h \geq 2$. Then, all the probabilistic results in Section 4 should extend naturally to random uniform hypergraphs. We believe that a direct extension of our arguments in this paper will yield all results in Section 2 for random uniform hypergraphs.

References

- [1] Y. Azar, A. Z. Broder, A. R. Karlin, and E. Upfal. Balanced allocations. *SIAM journal on computing*, 29(1):180–200, 1999.
- [2] P. Berenbrink, A. Czumaj, A. Steger, and B. Vöcking. Balanced allocations: The heavily loaded case. *SIAM Journal on Computing*, 35(6):1350–1385, 2006.
- [3] B. Bollobás. *Random Graphs*. Cambridge University Press, 2001.
- [4] B. Bollobás and A. Thomason. Random graphs of small order. In *Random graphs*, volume 83, pages 47–97, 1985.
- [5] J. A. Cain, P. Sanders, and N. Wormald. The random graph threshold for k -orientability and a fast algorithm for optimal multiple-choice allocation. In *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 469–476, New York, 2007. ACM.
- [6] P. A. Catlin, Zhi-Hong Chen, and E. M. Palmer. On the edge arboricity of a random graph. *Ars Combin.*, 35(A):129–134, 1993.
- [7] Xiaolin Chen, Xueliang Li, and Huishu Lian. Note on packing of edge-disjoint spanning trees in sparse random graphs. <http://arxiv.org/abs/1301.1097v1>.
- [8] N. Chiba and T. Nishizeki. Arboricity and subgraph listing algorithms. *SIAM Journal on Computing*, 14(1):210–223, 1985.
- [9] C. Cooper. On the thickness of sparse random graphs. *Combinatorics, Probability and Computing*, 1(4):303–309, 1992.
- [10] W. H. Cunningham. Optimal attach and reinforcement of a network. *Journal of the ACM*, 32(3):549–561, 1985.
- [11] A. Czumaj, C. Riley, and C. Scheideler. Perfectly balanced allocation. In *Approximation, Randomization, and Combinatorial Optimization: Algorithms and Techniques, 6th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX) and 7th International Workshop on Randomization and Approximation Techniques in Computer Science (RANDOM), Proceedings*, pages 240–251, 2003.

- [12] M. Dietzfelbinger, A. Goerdt, M. Mitzenmacher, A. Montanari, R. Pagh, and M. Rink. Tight thresholds for cuckoo hashing via xorsat. <http://arxiv.org/abs/0912.0287>.
- [13] P. Erdős and A. Rényi. On the evolution of random graphs. *Bull. Inst. Internat. Statist.*, 38:343–347, 1961.
- [14] D. Fernholz and V. Ramachandran. The k -orientability thresholds for $G_{n,p}$. In *Proc. ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 459 – 468, January 2007.
- [15] N. Fountoulakis, M. Khosla, and K. Panagiotou. The multiple-orientability thresholds for random hypergraphs. In *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1222–1236, 2011.
- [16] N. Fountoulakis and K. Panagiotou. Sharp load thresholds for cuckoo hashing. <http://arxiv.org/abs/0910.5147>.
- [17] A. Frank, T. Király, and M. Kriesell. On decomposing a hypergraph into k connected sub-hypergraphs. *Discrete Appl. Math.*, 131(2):373–383, 2003. Submodularity.
- [18] A. Frieze and P. Melsted. Maximum matchings in random bipartite graphs and the space utilization of cuckoo hashtables. <http://arxiv.org/abs/0910.5535>.
- [19] A. M. Frieze and T. Łuczak. Edge disjoint spanning trees in random graphs. *Period. Math. Hungar.*, 21(1):35–37, 1990.
- [20] P. Gao, X. Pérez-Giménez, and C. M. Sato. Arboricity and spanning-tree packing in random graphs with an application to load balancing. In *Proceedings of ACM-SIAM Symposium on Discrete Algorithms (2014)*, pages 317–326, 2014.
- [21] P. Gao and N. Wormald. Load balancing and orientability thresholds for random hypergraphs. In *STOC*, June 5–8 2010.
- [22] G. Goel and J. Gustedt. Bounded arboricity to determine the local structure of sparse graphs. In *Proceedings of the 32nd international conference on Graph-Theoretic Concepts in Computer Science, WG'06*, pages 159–167, Berlin, Heidelberg, 2006. Springer-Verlag.
- [23] D. Gusfield. Connectivity and edge-disjoint spanning trees. *Information Processing Letters*, 16(2):87–89, 1983.
- [24] S. L. Hakimi. On the degrees of the vertices of a directed graph. *J. Franklin Inst.*, 279:290–308, 1965.
- [25] A. Itai and M. Rodeh. The multi-tree approach to reliability in distributed networks. *Information and Computation*, 79(1):43–59, 1988.
- [26] S. Janson and M. J. Luczak. A simple solution to the k -core problem. *Random Structures Algorithms*, 30(1-2):50–62, 2007.

- [27] S. Janson, T. Łuczak, and A. Ruciński. *Random graphs*. Wiley-Interscience Series in Discrete Mathematics and Optimization. Wiley-Interscience, New York, 2000.
- [28] J. H. Kim. Poisson cloning model for random graphs. In *International Congress of Mathematicians. Vol. III*, pages 873–897. Eur. Math. Soc., Zürich, 2006.
- [29] M. Lelarge. A new approach to the orientation of random hypergraphs. In *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 251–264, 2012.
- [30] C. McDiarmid and B. Reed. Linear arboricity of random regular graphs. *Random Structures Algorithms*, 1(4):443–445, 1990.
- [31] M. Mitzenmacher, A. Richa, and R. Sitaraman. The power of two random choices: a survey of techniques and results. In *Handbook of randomized computing*, volume I, II, pages 255–312. Comb. Optim., 9, Kluwer Acad. Publ., Dordrecht, 2001.
- [32] M. Mitzenmacher and E. Upfal. *Probability and computing: randomized algorithms and probabilistic analysis*. Cambridge University Press, 2005.
- [33] M. Molloy. Cores in random hypergraphs and Boolean formulas. *Random Structures Algorithms*, 27(1):124–135, 2005.
- [34] C. St. J. A. Nash-Williams. Edge-disjoint spanning trees of finite graphs. *J. London Math. Soc.*, 36:445–450, 1961.
- [35] C. St. J. A. Nash-Williams. Decomposition of finite graphs into forests. *J. London Math. Soc.*, 39:12, 1964.
- [36] E. M. Palmer. On the spanning tree packing number of a graph: a survey. *Discrete Math.*, 230(1-3):13–21, 2001. Paul Catlin memorial collection (Kalamazoo, MI, 1996).
- [37] E. M. Palmer and J. J. Spencer. Hitting time for k edge-disjoint spanning trees in a random graph. *Period. Math. Hungar.*, 31(3):235–240, 1995.
- [38] B. Pittel, J. Spencer, and N. Wormald. Sudden emergence of a giant k -core in a random graph. *J. Combin. Theory Ser. B*, 67(1):111–151, 1996.
- [39] Peter Sanders, Sebastian Egner, and Jan Korst. Fast concurrent access to parallel disks. *Algorithmica*, 35(1):21–55, 2003.
- [40] A. Schrijver. *Combinatorial optimization. Polyhedra and efficiency. Vol. B*, volume 24 of *Algorithms and Combinatorics*. Springer-Verlag, Berlin, 2003. Matroids, trees, stable sets, Chapters 39–69.
- [41] V. N. Swamy, S. Bhashyam, R. Sundaresan, and P. Viswanath. An asymptotically optimal push-pull method for multicasting over a random network. *IEEE Transactions on Information Theory*, 59(8):5075–5087, 2013.
- [42] W. T. Tutte. On the problem of decomposing a graph into n connected factors. *J. London Math. Soc.*, 36:221–230, 1961.

Appendix

Proof of Lemma 30. The fact that $\xi_k(\mu)$ is strictly increasing follows from (5.10) in [38]. It is straightforward to verify that $\lim_{\mu \rightarrow \infty} \xi_k(\mu) = \infty$.

Define for $k \geq 3$ and $\mu > 0$

$$g_k(\mu) = \sum_{i \geq 0} \frac{\mu^i}{[k-1+i]_{i+1}} = \frac{(k-2)!}{\mu^{k-1}} \sum_{i \geq k-1} \frac{\mu^i}{i!}.$$

Claim 6. For every $k \geq 3$ and $\mu \geq 1.7(k-1)$, we have $g_k(\mu) > 1$.

Since g_k is strictly increasing in $(0, \infty)$, it suffices to prove the claim for $\mu = 1.7(k-1)$. Cases $3 \leq k \leq 50$ can be verified by direct computations (since we are aiming at a strict inequality, it suffices to compute a finite number of terms in the series defining $g_k(\mu)$). So, let's assume $k \geq 51$, and put $a = \lceil (k-1)/4 \rceil$ and $b = \lceil (k-1)/2 \rceil$. Then

$$\begin{aligned} g_k(\mu) &> \sum_{i=a}^b \frac{\mu^i}{[k-1+i]_{i+1}} > \frac{b-a+1}{k-1} \left(\frac{\mu}{k-1+b} \right)^a \geq \frac{1}{4} \left(\frac{1.7(k-1)}{\lceil 1.5(k-1) \rceil} \right)^{(k-1)/4} \\ &\geq \frac{1}{4} \left(\frac{1.7}{1.52} \right)^{(k-1)/4} > 1, \end{aligned}$$

which completes the proof of the claim.

Recall that for $k \geq 3$, $h_k(\mu)$ has a global minimum at μ_k . In particular, μ_k satisfies $h_k'(\mu_k) = 0$, which is equivalent to

$$(1 + \mu_k) \sum_{i \geq k-1} \frac{\mu_k^i}{i!} = \mu_k \sum_{i \geq k-2} \frac{\mu_k^i}{i!}.$$

Easy manipulations imply that μ_k is determined by $1 = g_k(\mu_k)$. By Claim 6 and since g_k is strictly increasing, we deduce that

$$\mu_k < 1.7(k-1).$$

Since $\xi_k(\mu)$ is strictly increasing, in order to prove the second statement in the lemma it suffices to show that $\xi_k(\mu) < 2k-2$ for $\mu = 1.7(k-1)$. This can be verified by direct computations for $3 \leq k \leq 24$. Suppose otherwise that $k \geq 25$. Then, we have

$$\frac{\mu^{k-1}}{(k-1)!} = \frac{[k+7]_8}{\mu^8} \frac{\mu^{k+7}}{(k+7)!} \leq \frac{(k+7)^8}{\mu^8} \frac{\mu^{k+7}}{(k+7)!} \leq \left(\frac{k+7}{1.7(k-1)} \right)^8 \frac{\mu^{k+7}}{(k+7)!} < 0.15 \frac{\mu^{k+7}}{(k+7)!},$$

and therefore

$$\frac{\mu}{2k-2} \sum_{i \geq k-1} \frac{\mu^i}{i!} = \frac{1.7}{2} \frac{\mu^{k-1}}{(k-1)!} + \frac{1.7}{2} \sum_{i \geq k} \frac{\mu^i}{i!} \leq \frac{1.7}{2} 0.15 \frac{\mu^{k+7}}{(k+7)!} + \frac{1.7}{2} \sum_{i \geq k} \frac{\mu^i}{i!} < \sum_{i \geq k} \frac{\mu^i}{i!}.$$

This implies that $\frac{\mu}{2k-2} f_{k-1}(\mu) < f_k(\mu)$, and hence $\xi_k(\mu) < 2k-2$, as desired. \square