

THE TOTAL ACQUISITION NUMBER OF RANDOM GRAPHS

DEEPAK BAL, PATRICK BENNETT, ANDRZEJ DUDEK, AND PAWEŁ PRAŁAT

ABSTRACT. Let G be a graph in which each vertex initially has weight 1. In each step, the weight from a vertex u to a neighbouring vertex v can be moved, provided that the weight on v is at least as large as the weight on u . The total acquisition number of G , denoted by $a_t(G)$, is the minimum possible size of the set of vertices with positive weight at the end of the process.

LeSaulnier, Prince, Wenger, West, and Worah asked for the minimum value of $p = p(n)$ such that $a_t(\mathcal{G}(n, p)) = 1$ with high probability, where $\mathcal{G}(n, p)$ is a binomial random graph. We show that $p = \frac{\log_2 n}{n} \approx 1.4427 \frac{\log n}{n}$ is a sharp threshold for this property. We also show that almost all trees T satisfy $a_t(T) = \Theta(n)$, confirming a conjecture of West.

1. INTRODUCTION

Gossiping and broadcasting are two well studied problems involving information dissemination in a group of individuals connected by a communication network [8]. In the gossip problem, each member has a unique piece of information which they would like to pass to everyone else. In the broadcast problem, there is a single piece of information (starting at one member) which must be passed to every other member of the network. These problems have received attention from mathematicians as well as computer scientists due to their applications in distributed computing [2]. Gossip and broadcast are respectively known as “all-to-all” and “one-to-all” communication problems. In this paper, we consider the problem of acquisition, which is a type of “all-to-one” problem. Suppose each vertex of a graph begins with a weight of 1 (this can be thought of as the piece of information starting at that vertex). A **total acquisition move** is a transfer of all the weight from a vertex v onto a vertex u , provided that immediately prior to the move, the weight on u is at least the weight on v . Suppose a number of acquisition moves are made until no legal moves remain. Such a maximal sequence of moves is referred to as an **acquisition protocol** and the vertices which retain positive weight after an acquisition protocol is called a **residual set**. Note that any residual set is necessarily an independent set. Given a graph G , we are interested in the minimum possible size of a residual set and refer to this number as the **total acquisition number of G** , denoted $a_t(G)$. We are mainly concerned with the question, “for which graphs G is $a_t(G) = 1$?” *i.e.* when can one special member of the network acquire all the information subject to the use of total acquisition moves? The restriction to total acquisition moves can be motivated by the so-called “smaller to larger” rule in disjoint

The third author is supported in part by Simons Foundation Grant #244712 and by a grant from the Faculty Research and Creative Activities Award (FRACAA), Western Michigan University.

The fourth author is supported in part by NSERC and Ryerson University.

set data structures. For example, in the UNION-FIND data structure with linked lists, when taking a union, the smaller list should always be appended to the longer list. This heuristic improves the amortized performance over sequences of union operations.

The parameter $a_t(G)$ was introduced by Lampert and Slater [11] and subsequently studied in [15, 12]. In [11], it is shown that $a_t(G) \leq \lfloor \frac{n+1}{3} \rfloor$ for any connected graph G on n vertices and that this bound is tight. Slater and Wang [15], via a reduction to the three-dimension matching problem, show that it is NP-complete to determine whether $a_t(G) = 1$ for general graphs G . In LeSaulnier *et al.* [12], various upper bounds on the acquisition number of trees are shown in terms of the diameter and the number of vertices, n . They also show that $a_t(G) \leq 32 \log n \log \log n$ (throughout the paper, $\log n$ denotes the natural logarithm) for all graphs with diameter 2 and conjecture that the true bound is constant. For work on game variations of the parameter and variations where acquisition moves need not transfer the full weight of vertex, see [17, 14, 16].

Randomness often plays a part in the study of information dissemination problems, usually in the form of a random network or a randomized protocol, see *e.g.* [4, 5, 6]. In this paper we study the total acquisition number of the **Erdős-Rényi-Gilbert random graph** $\mathcal{G}(n, p)$ where potential edges among n vertices are added independently with probability p . We also consider the total acquisition number of random trees. Our main theorem is the following.

Theorem 1.1. *Fix any $\varepsilon > 0$. If $p = p(n) \geq \frac{1+\varepsilon}{\log 2} \cdot \frac{\log n}{n}$, then with high probability, $a_t(\mathcal{G}(n, p)) = 1$.*

In particular, by taking $p = 1/2$, our result implies that while the question “Is $a_t(G) = 1$?” is NP-complete, the answer is “yes” for almost all graphs.

In [12], the authors mention that understanding the behaviour of $a_t(\mathcal{G}(n, p))$ near the connectivity threshold, $p = \frac{\log n}{n}$, would be of particular interest. In the theory of random graphs it is usually the case that some obvious necessary condition is also a sufficient one (for example, the threshold for connectivity coincides with the one for the minimum degree at least 1; the threshold for hamiltonicity is the same as the one for the minimum degree at least 2; etc.). Hence, one could expect that $a_t(\mathcal{G}(n, p)) = 1$ already at the time a random graph becomes connected. However, it turns out that connectivity is the wrong “obvious” condition. Consider the following observation.

Observation 1.2. *If vertex v is to acquire weight w (at any time during the process of moving weight around), then v has degree at least $\log_2 w$.*

Proof. Note that v can only ever acquire $1 + 2 + \dots + 2^{d(v)-1}$, in addition to the 1 it starts with, so that is a total of $2^{d(v)}$. \square

So if $a_t(G) = 1$ then the vertex which eventually acquires all the weight must have degree at least $\log_2 n$. Now it is true that when $p = \frac{\log n}{n}$, there exist vertices of this degree (see [3]). But just one such vertex does not suffice; a path of significant length consisting of high degree vertices is necessary. Such a path does not exist until the expected degree exceeds $\log_2 n$. So if $p < \log_2 n/n$, then $a_t(\mathcal{G}(n, p)) > 1$. In fact we prove the following stronger theorem.

Theorem 1.3. *Suppose that $p = \frac{c+o(1)}{\log 2} \cdot \frac{\log n}{n}$ for some fixed $c \in (0, 1)$. If $0 < \varepsilon < \min\{c, 1 - c\}$, then with high probability, $n^{1-c-\varepsilon} \leq a_t(\mathcal{G}(n, p)) \leq n^{1-c+\varepsilon}$.*

This result implies that at the connectivity threshold ($p = \frac{\log n}{n}$) the total acquisition number is already of polynomial size, namely it is at least, say, $n^{0.3}$. Theorems 1.1 and 1.3 together imply that $p = \frac{\log_2 n}{n}$ is the sharp threshold for the property $a_t(G) = 1$.

Moreover, we prove the following theorem, confirming a conjecture of West [18, 19]. Before we state the result, we need a few more definitions. For $n \in \mathbb{N}$, let \mathcal{T}_n be the family of labelled trees on n vertices. We say that some given property P holds for **almost all trees** if the ratio between the number of trees in \mathcal{T}_n with property P and the total number of trees in \mathcal{T}_n tends to 1 as $n \rightarrow \infty$.

Theorem 1.4. *For almost all trees $T \in \mathcal{T}_n$,*

$$a_t(T) \geq \frac{n}{3e^3}.$$

1.1. Notation and Conventions. The **random graph** $\mathcal{G}(n, p)$ consists of the probability space $(\Omega, \mathcal{F}, \mathbb{P})$, where Ω is the set of all graphs with vertex set $\{1, 2, \dots, n\}$, \mathcal{F} is the family of all subsets of Ω , and for every $G \in \Omega$,

$$\mathbb{P}(G) = p^{|E(G)|} (1-p)^{\binom{n}{2} - |E(G)|}.$$

This space may be viewed as the set of outcomes of $\binom{n}{2}$ independent coin flips, one for each pair (u, v) of vertices, where the probability of success (that is, adding edge uv) is p . Note that $p = p(n)$ may (and usually does) tend to zero as n tends to infinity. All asymptotics throughout are as $n \rightarrow \infty$ (we emphasize that the notations $o(\cdot)$ and $O(\cdot)$ refer to functions of n , not necessarily positive, whose growth is bounded). We say that an event in a probability space holds **with high probability** (or **w.h.p.**) if the probability that it holds tends to 1 as n goes to infinity. We often write $\mathcal{G}(n, p)$ when we mean a graph drawn from the distribution $\mathcal{G}(n, p)$.

For a vertex v in a graph, we write $d(v)$ for the degree of v .

We will use the following Chernoff bound:

Theorem 1.5 (Chernoff Bound). *If X is a binomial random variable with expectation μ , and $0 < \delta < 1$, then*

$$\Pr[X < (1 - \delta)\mu] \leq \exp\left(-\frac{\delta^2 \mu}{2}\right)$$

and if $\delta > 0$,

$$\Pr[X > (1 + \delta)\mu] \leq \exp\left(-\frac{\delta^2 \mu}{2 + \delta}\right).$$

In Section 2, we prepare the reader for the proof of the main result, Theorem 1.1, that can be found in Section 3. Theorem 1.3 is proved in Section 4 and Theorem 1.4 in Section 5. We conclude the paper with some open problems that can be found in Section 6.

2. WARMING UP BEFORE ATTACKING THEOREM 1.1

First note that in order to prove Theorem 1.1 it is enough to do it for $p = \frac{1+\varepsilon}{\log 2} \cdot \frac{\log n}{n}$ for an arbitrarily small $\varepsilon > 0$. This follows from that fact that $a_t(G) = 1$ is an increasing graph property (see for example Lemma 1.10 in [9]).

In this section, in order to prepare for a delicate and technical argument, we show that the result holds for $p = 2\sqrt{\log n/n}$, which is much larger than the value of p in Theorem 1.1. The idea behind both this warm-up exercise and the more difficult proof of Theorem 1.1 is to embed a certain type of tree in the random graph. The embedded tree contains all the edges that will be used for acquisition moves, and essentially specifies the entire acquisition protocol that will be used.

Let i be the largest integer such that

$$2^i \leq j = \left\lceil \sqrt{n/2} \right\rceil.$$

We construct a tree T rooted at vertex v in the following way. Vertex v has $1 + i + j$ children v_0, v_1, \dots, v_{i+j} . Vertex v_0 is a leaf, and for every $1 \leq \ell \leq i$ we have that vertex v_ℓ has $2^\ell - 1$ children; all remaining children of v have j children. Since the number of children of vertices v_1, v_2, \dots, v_i is at most $2^{i+1} = O(\sqrt{n})$, the number of vertices in T satisfies

$$\frac{n}{2} \leq j^2 \leq |V(T)| \leq j^2 + O(\sqrt{n}) = (1 + o(1))\frac{n}{2}.$$

It is straightforward to see that vertices of T can move their weight to the root v . Indeed, all grandchildren of v can move their weight to the corresponding parents, and then vertices v_0, v_1, \dots, v_i can send (one by one) the weight to the root (vertex v_ℓ sends the weight of 2^ℓ , $\ell = 0, 1, \dots, i$). At that point of the process, v has the weight of $2^{i+1} > j$ so the remaining neighbours can move their weight to v .

Let v be any vertex in $\mathcal{G}(n, p)$. First, we will show that w.h.p. there exists a tree T rooted at v that can be embedded in $\mathcal{G}(n, p)$. It follows from Chernoff Bound that w.h.p. the degree of v is $(2 + o(1))\sqrt{n \log n}$. We select (arbitrarily) $1 + i + j$ neighbours of v and label them as v_0, v_1, \dots, v_{i+j} . We continue discovering neighbours of v_ℓ s (one by one) but not every neighbour of v_ℓ will be used for the tree T (there will be more neighbours than children in the corresponding subtree). Since the total number of vertices in T is $(1 + o(1))n/2$, there will always be at least $n/3$ vertices left that are not embedded yet. Hence, the number of neighbours of v_ℓ that are not embedded yet is a random variable that can be lower bounded by the binomial random variable $\text{Bin}(n/3, p)$ with expected value of $p(n/3) = (2/3)\sqrt{n \log n}$. Hence, using Chernoff Bound, with probability $1 - o(n^{-1})$ there will be enough neighbours of v_ℓ to continue the process. It follows from the union bound that T can be embedded w.h.p.

As we already mentioned, vertices of T can move their weight to the root v . It is enough to show that the remaining vertices can do that too. Let S be the set of neighbours of v that are outside of T . Since

$$s = |S| = \deg(v) - (1 + i + j) = (2 + o(1))\sqrt{n \log n} - O(\sqrt{n}) = (2 + o(1))\sqrt{n \log n},$$

there should be enough vertices in S to dominate the rest of the graph and push the remaining weight to v . The important observation is that at this point of the process v

has weight at least $n/2$ so we do not have to control how much weight we push from a vertex of S to v . It remains to show that S dominates the remaining vertices (*i.e.* each remaining vertex is adjacent to a vertex in S) w.h.p. But this is straightforward to see, since for a given vertex we have that the probability it is not dominated is equal to

$$(1 - p)^s = \exp(-(1 + o(1))ps) = \exp(-(4 + o(1)) \log n) = o(n^{-1}),$$

and the claim holds by the union bound.

In order to generalize these ideas to sparse graphs we have to deal with a number of problems, each of which is relatively easy to overcome but addressing all of them requires more careful argument. The spirit of the proof however, remains the same: we will define a special rooted tree (recursively) which has the property that all the weight can be moved to the root and we will show that this tree can be embedded in $\mathcal{G}(n, p)$ w.h.p. This tree will not quite be spanning, but it will contain a set of vertices, B , which dominates the remaining vertices, R , and can shift their weight to the root as well. Finding a matching from R to B which saturates R completes the proof. There have been numerous results on embedding spanning and almost spanning trees in random graphs [1, 10, 13], but most of these results are for embedding bounded degree trees and are not precise enough for our purposes. Let us list the most important issues and briefly describe the way we want to deal with them. Suppose that $p = \frac{1+\varepsilon}{\log 2} \cdot \frac{\log n}{n}$.

Problem 1: Small average degree

By Observation 1.2, in order for a vertex to be able to accumulate a weight of w , it has to have degree at least $\log_2 w = \frac{1}{\log 2} \log w$. Since the average degree is only $\frac{1+\varepsilon}{\log 2} \log n$, it follows that (almost) every time a neighbour of the root v sends its weight to v , the weight is (almost) doubled. In particular, some children of v must be able to send a large weight to v , much more than the number of their children. Hence, we will need to define the tree recursively. As we will see in Corollary 3.3, the tree will reach level $m = (1 + o(1)) \frac{\log n}{\log \log n}$.

Problem 2: Continuing to embed after a linear number of vertices

As we already mentioned, the root and some vertices on top levels must have degrees close to the average degree in the graph. We will require that the number of children for those vertices is (roughly) $\frac{1+\varepsilon/2}{\log 2} \log n$. However, once a positive fraction of all vertices are already embedded in the tree, the number of available neighbours drops substantially, so in order to be able to continue the process, we will have to decrease the required number of children to $\beta \log n$ for some β . The bottom $\alpha \frac{\log n}{\log \log n}$ levels of the tree will have this property. The number of children on level k will be denoted by c_{m-k-1} . (As explained below, it will be more convenient to count levels from the bottom; hence the notation c_{m-k-1} instead of more natural c_k .)

Problem 3: Vertices with degree smaller than average

Even though the average degree is $\frac{1+\varepsilon}{\log 2} \log n$, it is possible that a vertex does not have the required number of children (either $\frac{1+\varepsilon/2}{\log 2} \log n$ or $\beta \log n$). This is not avoidable but rare, and we will show that w.h.p. there are at most σ children of a given vertex that have this undesired property (we will see in the proof of Lemma 3.8 that $\sigma = \Theta(1/\varepsilon^2)$). Nevertheless, we have to take this into account while constructing the tree.

Before defining our tree, we first define a property of rooted trees which (as we will see soon) guarantees that the root can acquire all the weight on the tree, given that each vertex begins with weight 1.

Definition 2.1 (Cut-off Property). *Let T be a tree rooted at r . We say T has the **cut-off property** if the following holds: for each vertex v with children v_1, v_2, \dots, v_k , and denoting by T_i the subtree rooted at v_i , there exists an i' (which may depend on v) so that $|T_i| = 2^{i-1}$ for $i \leq i'$, and $|T_i| \leq 2^{i'}$ for $i > i'$.*

*In this case, the vertices v_i for $i \leq i'$ are called **exact**. A vertex with an exact ancestor is called **tight**, and vertices that are not tight are **loose**.*

Lemma 2.2. *If T is a tree rooted at r which has the cut-off property, then $a_t(T) = 1$. In particular, vertex r can acquire all the weight.*

Proof. We proceed by induction on the depth of T . The base case (depth 0) is trivial. To see the induction step, let r have children v_1, v_2, \dots, v_k and let T_i be the subtree of T rooted at the v_i . Then the T_i inherit the cut-off property, and all have depth strictly less than the depth of T , and so by induction, all the weight from subtree T_i can be loaded onto the root v_i .

Now it is easy to see that by the cut-off property, r may acquire the weight of each child v_i , going in order of increasing index. \square

It is time to define our recursive construction of a tree which will have the cut-off property.

Definition 2.3. *For any $\rho, m, \sigma \in \mathbb{N}$, and positive integer sequence c_{m-k-1} , construct the rooted tree T_ρ by the following process:*


- (i) *Initialize: The **root** vertex $r = \langle \rangle$, the **weight** $w(r) = \rho$, the **level** $k = 0$.*
- (ii) *Iterate: In level k , for a vertex $\langle i_1, i_2, \dots, i_k \rangle$ with weight $w = w(\langle i_1, i_2, \dots, i_k \rangle) > 1$,*
 - (a) *If $1 < w \leq c_{m-k-1}$, then attach $w - 1$ leaves to vertex $\langle i_1, i_2, \dots, i_k \rangle$ each with weight 1.*
 - (b) *If $w > c_{m-k-1}$, then attach $c := c_{m-k-1}$ children to vertex $\langle i_1, i_2, \dots, i_k \rangle$, labelled $\langle i_1, i_2, \dots, i_k, 1 \rangle, \dots, \langle i_1, i_2, \dots, i_k, c \rangle$. Let i' be the minimum integer $i \geq 0$ such that*

$$\frac{w - 2^i - \sigma}{c - i - \sigma} \leq 2^i + \sigma.$$

Assign weights to the children as follows

$$w(\langle i_1, i_2, \dots, i_{k+1} \rangle) = \begin{cases} 2^{i_{k+1}-1} & \text{if } i_{k+1} \leq i' \\ 1 & \text{if } i' < i_{k+1} \leq i' + \sigma \\ \frac{w - 2^{i'} - \sigma}{c - i' - \sigma} & \text{otherwise.} \end{cases} \quad (1)$$

Here, we assume that $\frac{w - 2^{i'} - \sigma}{c - i' - \sigma}$ is an integer and that $i' + \sigma < c$ so that i' is well defined.



Note that the tree T_ρ has the cut-off property. In this definition, $w(v)$ is meant to represent the number of vertices which will end up in the subtree rooted at v . The sequence c provides a sort of threshold for the recursive part of the definition to come into play. So if the weight on v is at most c , then the entire subtree appears in the form of leaves. If the weight on v exceeds c , then v will have exactly c children and the weight is distributed according to (1).

In Definition 2.3 it is convenient to work “from the top down,” starting with all the weight at the root, and then iteratively assigning to each vertex the number of children necessary to distribute its weight so as to achieve the cut-off property. However, the assumption that $\frac{w-2^{i'}-\sigma}{c-i'-\sigma}$ is always an integer may seem worrisome to the reader. We will ultimately choose the parameters of T_ρ so that this assumption holds. Our choice of parameters will be made by working “from the bottom up,” i.e. specifying what the bottom layer (the layer corresponding to the largest value of k in Definition 2.3) ought to look like, and then working our way up layer by layer. See Definition 3.1 and Lemma 3.5.

Problem 4: Divisibility

In the recursive definition of the tree, we distribute the remaining weight equally among some number of vertices. Hence, we will have to make sure that certain divisibility conditions hold. Unfortunately, it is difficult to find the initial weight ρ that does it. Hence, in order to do that, we first define c_k^* to be the desired number of children of vertices on level $k + 1$ (this time, counted from the bottom), and then assign weights starting from the bottom level and doing calculations upwards. This issue is addressed by Definition 3.1.

Problem 5: Finding the right size tree to embed

Our goal will be to construct a tree that consists of $(\frac{8}{5} + o(1))n$ vertices, that is, a tree T_ρ with $\rho = (\frac{8}{5} + o(1))n$. (Of course, $\mathcal{G}(n, p)$ has only n vertices; T_ρ is an abstract tree that will be “trimmed” before embedding it into a random graph.) As we already mentioned, for a fixed sequence of c_k^* ’s and σ , one can easily (recursively) calculate the weight ρ_k of loose vertices on level k (again, counted from the bottom), and the weight of the root $\rho = \rho_m$. However, it is hard to expect that the desired condition holds, namely, that $\rho_m = (\frac{8}{5} + o(1))n$. In order to solve this problem we start with any sequence c_k^* , take m to be the largest integer such that $\rho_m \leq \frac{8}{5}n$, and then modify the sequence slightly to get the desired sequence c_k with $\rho_m = (\frac{8}{5} + o(1))n$. Let us note that

a non-constructive argument is used here that shows only the existence; the sequence c_k is not explicitly defined. Lemmas 3.2 and 3.4 provide useful relationships between the sequences c and ρ which aid in the proof of Lemma 3.5 which proves the existence of the desired sequence c .

Problem 6: Finishing the embedding

It is difficult to expect that a given tree on n vertices can be embedded in a random graph. Hence, we are going to remove a number of leaves in T_ρ to get another tree T'_ρ on $(\frac{4}{5} + o(1))n$ vertices that can be embedded in $\mathcal{G}(n, p)$ w.h.p. The important property will be that parents of removed leaves can not only dominate the remaining $(\frac{1}{5} + o(1))n$ vertices but also can push all the weight to the root. This issue is addressed by Definition 3.6 and Lemma 3.7.

3. PROOF OF THEOREM 1.1

The proof will be similar to the warm-up exercise in Section 2, in that we will embed a certain type of spanning tree. There has been some research in the problem of embedding spanning (or almost spanning) trees into random graphs ([1, 10, 13]), but these results will not quite achieve our goal. We are working with $\mathcal{G}(n, p)$ for $p = \Theta(\frac{\log n}{n})$, and we must embed a tree with maximum degree $\Delta = \Theta(\log n)$. But the results in [1, 10, 13] all require at least $p = \Omega(\frac{\Delta \log \Delta}{n})$.

The main probabilistic tools we will use are Markov's inequality and the Chernoff bound. The difficulty lies in defining our spanning tree in such a manner that we can argue that it embeds into $\mathcal{G}(n, p)$ using only these simple tools. There are two main reasons why our tree is easier to embed than the trees from [1, 10, 13]. The first is that we will embed a deterministic tree that spans some linear number of vertices (but not all vertices), and then we will span the rest of the vertices in a non-deterministic way that depends on the edges of $\mathcal{G}(n, p)$ (see Problem 6). This is easier than [10, 13], which embed a completely deterministic tree. The second reason our tree is easier to embed is that we construct it to have several leaves adjacent to each vertex. In this way we avoid the problem of not finding enough neighbours with high enough degree. If a vertex v has some neighbours whose degree is not high enough, we let those neighbours be the leaves adjacent to v (see Problem 3).

We will now state a definition that will help us choose the parameters of T_ρ so that it has all the desired properties. As we mentioned before, we will work “from the bottom up” here, which is in contrast to Definition 2.3 where we worked “from the top down.” Specifically, in Definition 3.1 the index j is used to count layers of the tree similarly to the index k in Definition 2.3, but $j = 1$ in Definition 3.1 corresponds to the bottom (i.e. the leaves farthest from the root), whereas $k = 0$ in Definition 2.3 corresponds to the top (i.e. the root). Set $d = \frac{1+\varepsilon}{\log 2} \log n$ and let σ, β, α be constants.

Definition 3.1. *Let*

$$c_j^* := \begin{cases} \beta \log n & \text{if } j \leq \alpha \frac{\log n}{\log \log n} \\ \frac{1+\varepsilon}{\log 2} \log n & \text{otherwise.} \end{cases}$$

and let c be a sequence such that $c_j^* \leq c_j \leq c_j^*(1 + o(1))$. Define the function

$$i^*(x) := \begin{cases} 0 & \text{if } x \leq \sigma \\ \lceil \log_2(x - \sigma) \rceil & \text{otherwise.} \end{cases}$$

Finally, define sequences ρ_1, ρ_2, \dots and b_1, b_2, \dots recursively by putting $\rho_1 := 2$ and

$$\rho_{j+1} := \sigma + 2^{i^*(\rho_j)} + (c_j - i^*(\rho_j) - \sigma) \cdot \rho_j,$$

$b_1 := 1$ and

$$b_{j+1} := (c_j - i^*(\rho_j) - \sigma) \cdot b_j.$$

Note that the sequences ρ, b depend on our choice of the sequence c (we assume that constants $\varepsilon, \sigma, \beta, \alpha$ are fixed in advance). As was mentioned before (see Problem 4), the main purpose of this recursive sequence is to calculate (for a given sequence c and depth m) the weight of the root; in fact, ρ_j is the weight of each loose vertex at level j (counted from the bottom) so the weight of the whole tree is ρ_m . Let us also mention that the purpose of i' in Definition 2.3 was to make sure that the total weight of subtrees rooted at exact vertices together with the weight of the root is at least the weight of each subtree rooted at non-exact children. It is straightforward to see that i^* in Definition 3.1 has the same purpose and so these values are always the same. Finally, let us point out that we fix $\rho_1 = 2$ which indicates that every loose vertex at the level directly above the bottom has precisely one leaf. These leaves will play an important role in our argument and b_j counts how many such leaves we have in the tree rooted at loose vertices at level j (as usual, counted from the bottom).

Let ρ^* be the sequence obtained by setting $c = c^*$ in Definition 3.1. Let m be the largest integer so that $\rho_m^* \leq \frac{8}{5}n$. Note that for every $j \geq 2$, $\rho_j \geq \rho_2 = \Omega(\log n)$ and so

$$\rho_{j+1} = (c_j - \log_2 \rho_j + O(1)) \cdot \rho_j \leq c_j \rho_j. \quad (2)$$

It follows that $\rho_m^* = \Omega\left(\frac{n}{\log n}\right)$, since ρ grows by at most a log factor each time.

Henceforth we will keep $m = m(n)$ as defined above (2) and consider the sequences ρ, b only up to the terms ρ_m, b_m . We will consider sequences c with terms that might be larger than those of c^* . As a result, $\rho_j \geq \rho_j^*$ for all j . However, we will only consider sequences c such that the corresponding sequence ρ has $\rho_m = n^{1+o(1)}$.

Lemma 3.2. *Let c be any sequence such that $c_j^* \leq c_j \leq c_j^*(1 + o(1))$ and the ρ -sequence corresponding to c has $\rho_m = n^{1+o(1)}$. If $\alpha < \frac{\beta \log 2}{2}$ then*

$$\rho_j = \exp\{(j-1) \log \log n + O(j)\}$$

for all $2 \leq j \leq m$.

Proof. Clearly, $\rho_2 = \Theta(\log n)$ and it follows immediately from (2) that for every $2 \leq j < m$ we have

$$\frac{\rho_{j+1}}{\rho_j} \leq c_j \leq 2c_j^* \leq 2 \max\left(\beta, \frac{1 + \frac{\varepsilon}{2}}{\log 2}\right) \log n.$$

Hence, $\rho_j \leq \exp\{(j-1)\log\log n + O(j)\}$ for all $j \leq m$ and so the upper bound holds. In particular, as long as $j \leq \alpha \frac{\log n}{\log\log n}$ we have

$$\rho_j \leq \rho_{\alpha \frac{\log n}{\log\log n}} \leq n^{\alpha(1+o(1))}. \quad (3)$$

For the lower bound, we use (2) one more time and (3) to observe that for every j such that $2 \leq j \leq \alpha \frac{\log n}{\log\log n}$ we have

$$\frac{\rho_{j+1}}{\rho_j} \geq \left(\beta(1+o(1)) - \frac{\alpha}{\log 2}(1+o(1)) \right) \log n \geq \frac{\beta}{2} \log n$$

(by our choice of α). Since for every $j \leq m$ we have $\rho_j \leq \rho_m = n^{1+o(1)}$, for every j such that $\alpha \frac{\log n}{\log\log n} < j < m$ we have

$$\frac{\rho_{j+1}}{\rho_j} \geq \left(\frac{1 + \frac{\varepsilon}{2} + o(1)}{\log 2} - \frac{1 + o(1)}{\log 2} \right) \log n \geq \frac{\varepsilon}{4 \log 2} \log n.$$

It follows that $\rho_j \geq \exp\{(j-1)\log\log n + O(j)\}$ for all $j \leq m$ and the proof is finished. \square

Henceforth, we assume that $\alpha < \frac{\beta \log 2}{2}$. We immediately get the following corollary.

Corollary 3.3. $m = (1 + o(1)) \frac{\log n}{\log\log n}$.

In Definition 2.3, it was assumed that i' was well defined, that is, that the condition $i' + \sigma < c$ holds. Because of the relationship between the two definitions, this condition is equivalent to the condition $c_j - i^*(\rho_j) - \sigma \geq 1$ in Definition 3.1. In the next lemma, we show that the same condition for α as in Lemma 3.2 is enough to guarantee that i' is well defined. The following is a useful property which we will use in the next few arguments.

Lemma 3.4. *We have that i' always exists. In fact, the following stronger property holds: for every j such that $1 \leq j \leq m$ we have*

$$c_j - \log_2 \rho_j = \Omega(\log n).$$

Proof. In order to show that i' exists, we will show that the equivalent condition that $c_j - i^*(\rho_j) - \sigma \geq 1$ in Definition 3.1 holds. In fact, we will show something stronger, namely, that $c_j - \log_2 \rho_j = \Omega(\log n)$ for every $1 \leq j \leq m$. For $j \leq \alpha \frac{\log n}{\log\log n}$ we have

$$c_j - \log_2 \rho_j \geq c_j - \log_2 \rho_{\alpha \frac{\log n}{\log\log n}} = (1 + o(1)) \left(\beta - \frac{\alpha}{\log 2} \right) \log n = \Omega(\log n).$$

If $j > \alpha \frac{\log n}{\log\log n}$, then

$$c_j - \log_2 \rho_j \geq c_j - \log_2 \rho_m = (1 + o(1)) \left(\frac{1 + \frac{\varepsilon}{2}}{\log 2} - \frac{1}{\log 2} \right) \log n = \Omega(\log n).$$

The proof is complete. \square

Our next task is to show that one can adjust a sequence c^* slightly to get another sequence c with $\rho_m = \left(\frac{8}{5} + o(1)\right) n$.

Lemma 3.5. *There exists a sequence of integers c_j with $c_j^* \leq c_j \leq c_j^*(1 + o(1))$ and $\rho_m = \left(\frac{8}{5} + o(1)\right)n$.*

Proof. We will start by setting $c_j = c_j^*$ for all j , and then apply a number of operations to the sequence c . Each operation will consist of increasing a single term c_j by 1 and leaving all other terms the same.

Suppose that the sequence \tilde{c} agrees with sequence c except in the j_0 term where we have $\tilde{c}_{j_0} = c_{j_0} + 1$. Let $\tilde{\rho}$ and ρ be the corresponding sequences, which must then agree for all $j \leq j_0$. Then, it follows from (2) and Lemma 3.4 that

$$\begin{aligned} \frac{\tilde{\rho}_{j_0+1}}{\rho_{j_0+1}} &= \frac{\left((c_{j_0} + 1) - \log_2 \rho_{j_0} + O(1)\right) \cdot \rho_{j_0}}{\left(c_{j_0} - \log_2 \rho_{j_0} + O(1)\right) \cdot \rho_{j_0}} \\ &= 1 + O\left(\frac{1}{c_{j_0} - \log_2 \rho_{j_0}}\right) = 1 + O\left(\frac{1}{\log n}\right). \end{aligned}$$

Now, since $\tilde{\rho}_j \geq \rho_j$, for $j \geq j_0 + 1$ we have

$$\begin{aligned} \frac{\tilde{\rho}_{j+1}}{\rho_{j+1}} &= \frac{\left(c_j - \log_2 \tilde{\rho}_j + O(1)\right) \cdot \tilde{\rho}_j}{\left(c_j - \log_2 \rho_j + O(1)\right) \cdot \rho_j} \leq \frac{\left(c_j - \log_2 \rho_j + O(1)\right) \cdot \tilde{\rho}_j}{\left(c_j - \log_2 \rho_j + O(1)\right) \cdot \rho_j} \\ &= \frac{\tilde{\rho}_j}{\rho_j} \cdot \left(1 + O\left(\frac{1}{c_j - \log_2 \rho_j}\right)\right) = \frac{\tilde{\rho}_j}{\rho_j} \cdot \left(1 + O\left(\frac{1}{\log n}\right)\right). \end{aligned}$$

Hence,

$$\frac{\tilde{\rho}_m}{\rho_m} \leq \left(1 + O\left(\frac{1}{\log n}\right)\right)^m = 1 + O\left(\frac{1}{\log \log n}\right).$$

In other words, each time we increment a term of sequence c , the effect on ρ_m is negligible. However we will now show that if we perform this operation on c enough times (while still not changing it too much each time), the effect on ρ_m can be as much as we need it to be.

Suppose now that $\tilde{c}_j = c_j^* + \frac{\log n}{\log \log \log n} = c_j^*(1 + o(1))$ for all j . Our goal is to show that $\tilde{\rho}_m \geq \frac{8}{5}n$. For a contradiction, suppose that it is not the case, that is, $\tilde{\rho}_m < \frac{8}{5}n$. Note that we have

$$\log_2 \tilde{\rho}_j = \log_2 \rho_j^* + O(j)$$

since from Lemma 3.2 it follows that both $\tilde{\rho}_j$ and ρ_j^* are equal to $\exp\{(j-1)\log\log n + O(j)\}$. Using (2) as usual, by Lemma 3.4 we get that

$$\begin{aligned} \frac{\tilde{\rho}_{j+1}}{\rho_{j+1}^*} &= \frac{\left(c_j + \frac{\log n}{\log\log\log n} - \log_2 \tilde{\rho}_j + O(1)\right) \cdot \tilde{\rho}_j}{\left(c_j - \log_2 \rho_j^* + O(1)\right) \cdot \rho_j^*} \\ &= \frac{c_j - \log_2 \rho_j^* + \frac{\log n}{\log\log\log n} + O(j)}{c_j - \log_2 \rho_j^* + O(1)} \cdot \frac{\tilde{\rho}_j}{\rho_j^*} \\ &= \frac{\tilde{\rho}_j}{\rho_j^*} \left(1 + \Theta\left(\frac{\log n}{(c_j - \log_2 \rho_j^*) \log\log\log n}\right)\right) \\ &= \frac{\tilde{\rho}_j}{\rho_j^*} \left(1 + \Theta\left(\frac{1}{\log\log\log n}\right)\right) = \frac{\tilde{\rho}_j}{\rho_j^*} \exp\left\{\Theta\left(\frac{1}{\log\log\log n}\right)\right\} \end{aligned}$$

And so we have

$$\begin{aligned} \tilde{\rho}_m &= \rho_m^* \cdot \exp\left\{\Theta\left(\frac{m}{\log\log\log n}\right)\right\} \\ &= \Omega\left(\frac{n}{\log n}\right) \cdot \exp\left\{\Theta\left(\frac{\log n}{(\log\log n)(\log\log\log n)}\right)\right\} \gg n \end{aligned}$$

which is a contradiction and so the sequence \tilde{c} is such that $\tilde{\rho}_m \geq \frac{8}{5}n$.

Thus we can apply the operation ‘‘increment one term by 1’’ to the sequence $c = c^*$ several times so that each term gets increased by at most $\frac{\log n}{\log\log\log n} = o(\log n)$, and we are able to do so in such a manner that $\rho_m = \left(\frac{8}{5} + o(1)\right)n$. The proof is finished. \square

Definition 3.6. Define the tree T'_{ρ_j} to be the tree T_{ρ_j} with each leaf in the bottom level being removed if it has a loose parent. Call the parents that lose their children **bereft**.

Note that by induction and definition of ρ_j, b_j , and T'_{ρ_j} , we see that T'_{ρ_j} has $\rho_j - b_j$ many vertices, b_j of which are bereft. It is not difficult to see that by construction, T'_{ρ_j} has the cut-off property. Moreover, if we form another tree T''_{ρ_j} by re-attaching at most one leaf to each bereft parent of T'_{ρ_j} , then T''_{ρ_j} still has the cut-off property.

Our next goal is to show that almost all vertices of T'_{ρ_j} are bereft. Since each bereft vertex has exactly one child in T_{ρ_j} , we get that $|T_{\rho_j}| = (2 + o(1))|T'_{\rho_j}|$. In particular, $|T'_{\rho_j}| = \left(\frac{4}{5} + o(1)\right)n$.

Lemma 3.7. For all $1 \leq j \leq m$ we have $\frac{b_j}{\rho_j - b_j} \rightarrow 1$ as $n \rightarrow \infty$.

Proof. Note that

$$\begin{aligned} \frac{\rho_{j+1}}{b_{j+1}} &= \frac{\sigma + 2^{i^*(\rho_j)} + (c_j - i^*(\rho_j) - \sigma) \cdot \rho_j}{(c_j - i^*(\rho_j) - \sigma) \cdot b_j} = \frac{(c_j - \log_2 \rho_j + O(1)) \cdot \rho_j}{(c_j - \log_2 \rho_j + O(1)) \cdot b_j} \\ &= \frac{\rho_j}{b_j} \cdot \left(1 + O\left(\frac{1}{c_j - \log_2 \rho_j}\right)\right) = \frac{\rho_j}{b_j} \cdot \left(1 + O\left(\frac{1}{\log n}\right)\right), \end{aligned}$$

since $c_j - \log_2 \rho_j = \Omega(\log n)$ by Lemma 3.4. Since $\frac{\rho_1}{b_1} = 2$, we have $\frac{\rho_j}{b_j} = 2 \cdot \left(1 + O\left(\frac{j}{\log n}\right)\right)$. Finally, since $j \leq m = o(\log n)$, for all $j \leq m$ we have that $\frac{\rho_j}{b_j} \rightarrow 2$ as $n \rightarrow \infty$, and the result follows. \square

Now, we are ready to show that T'_{ρ_m} can be embedded into the random graph $\mathcal{G}(n, p)$.

Lemma 3.8. *If $\beta < \frac{1}{10 \log 2}$ and $0 < \alpha < \frac{\beta \log 2}{2}$, then w.h.p. $\mathcal{G}(n, p)$ (with $p = \frac{1+\epsilon}{\log 2} \cdot \frac{\log n}{n}$) contains a copy of T'_{ρ_m} .*

Proof. We will embed T'_{ρ_m} in $\mathcal{G}(n, p)$. Select any vertex (arbitrarily) that will serve as a root of the tree. The embedding is done greedily and from the top down, and at each step we reveal the neighbourhood of one vertex. We group vertices in the same level (i.e. distance from the root) consecutively. The embedding will be determined iteratively as we reveal the random graph. We will not put a vertex of $\mathcal{G}(n, p)$ into our partial embedding until we have exposed all of its children.

We say that a vertex in level k is **bad** if its neighbourhood (into the unexposed vertices) contains less than c_{m-k-1} . We will show that w.h.p. the root is not bad and no vertex has more than σ bad children. Any bad children will be put into the partial embedding as leaves, and the other vertices will be arbitrarily assigned (to non-leaves first and then to leaves, if the number of bad children is smaller than σ).

The tree T'_{ρ_m} has at most

$$\left(\frac{1 + \frac{\epsilon}{2}}{\log 2} \log n\right)^{m - \alpha \frac{\log n}{\log \log n}} = n^{(1-\alpha)(1+o(1))} = o(n)$$

vertices total in levels 0 through $m - \alpha \frac{\log n}{\log \log n} - 1$. Thus, the expected degree (into the unexposed vertices) of each vertex exposed in such a level k is

$$(1 + o(1))d = (1 + o(1)) \frac{1 + \epsilon}{\log 2} \log n > c_k + \frac{\epsilon}{3 \log 2} \log n \quad (4)$$

and so it follows from Chernoff Bound that the probability that a fixed vertex is bad is polynomially small, that is, at most $n^{-\Theta(\epsilon^2)}$. For levels k farther to the bottom, note that the number of vertices that are not embedded yet is always at least $\frac{1+o(1)}{5}n$ and so the expected degree of each exposed vertex in layer k is at least

$$\frac{1 + o(1)}{5}d = (1 + o(1)) \frac{1 + \epsilon}{5 \log 2} \log n > c_k + \frac{1}{10 \log 2} \log n,$$

again yielding that the probability that a fixed vertex is bad is polynomially small (this time the exponent is a universal constant, not a function of ϵ).

Therefore if $\sigma = \Theta(1/\epsilon^2)$ is a large enough constant, then w.h.p. each vertex has at most σ bad children. This proves that our embedding procedure is successful w.h.p. and the proof is finished. \square

Let B be the set of bereft vertices in $\mathcal{G}(n, p)$ and R be the set of remaining, unexposed vertices that are not embedded into tree yet. Note that $|B| = b_m = \left(\frac{4}{5} + o(1)\right)n$ and $|R| = n - (1 + o(1))b_m = \left(\frac{1}{5} + o(1)\right)n$. An important property is that no edge between B and R is exposed at this point, so the next Lemma shows that w.h.p. set B dominates

set R but in such a way that at most one vertex of T is assigned to each bereft vertex. This will finish the proof of the main theorem, Theorem 1.1.

Lemma 3.9. *W.h.p. there is a matching from R to B which saturates R .*

Proof. We are going to use Hall's theorem for bipartite graphs. It is enough to show that for every subset $S \subseteq R$, Hall's condition holds, that is, we have that $|N(S) \cap B| \geq |S|$.

We will use the following useful upper bound:

$$\begin{aligned} & \Pr [\nexists \text{ matching saturating } R] \\ & \leq \Pr [\exists v \in R : N(v) \cap B = \emptyset] \\ & + \Pr [\exists S \subseteq R, T \subseteq B : |S| = k \geq 2, |T| = k - 1, N(S) \cap B = T, e(S : T) \geq 2(k - 1)] \end{aligned} \quad (5)$$

where $e(S : T)$ represents the number of edges between S and T . The first term bounds the probability that Hall's condition fails for some set of size one. To see why the condition in the second term is equivalent to the property that Hall's condition fails for some set of cardinality at least 2 is slightly more complicated. Take a smallest size $S \subseteq R$ with $|S| \geq 2$, which violates Hall's condition, i.e. $|S| > |T|$ where $T = N(S) \cap B$. If $|S| = k$, then $|T| = k - 1$, otherwise we could remove some vertex from S to get a smaller set that violates Hall's condition. Every vertex in T must have degree at least 2 into S , because removing a degree 1 vertex from T and its unique neighbour in S gives us a smaller set which violates Hall's condition. So the number of edges between S and T must be at least $2(k - 1)$.

In order to bound the first term in (5), note that

$$\begin{aligned} \Pr [\exists v \in R : N(v) \cap B = \emptyset] & \leq |R|(1 - p)^{|B|} \leq n \exp\left(-\frac{4}{5 \log 2} \log n\right) \\ & \leq n \exp(-1.15 \log n) = o(1). \end{aligned}$$

To bound the second term in (5), let Y count the number of sets S and T satisfying the condition in this term. Then we may bound the expectation of Y from above by

$$\begin{aligned} \mathbb{E}[Y] & \leq \sum_{k=2}^{|R|} \binom{|R|}{k} \binom{|B|}{k-1} \binom{k(k-1)}{2(k-1)} p^{2(k-1)} (1-p)^{k(|B|-(k-1))} \\ & \leq \sum_{k=2}^{|R|} \left(\frac{|R|e}{k}\right)^k \left(\frac{|B|e}{k-1}\right)^{k-1} \left(\frac{kep}{2}\right)^{2(k-1)} \exp(-pk(|B|-(k-1))), \end{aligned}$$

since $\binom{a}{b} \leq (ae/b)^b$. It follows that

$$\begin{aligned} \mathbb{E}[Y] & \leq \sum_{k=2}^{|R|} \exp\left((2k-1) \log\left(\frac{n}{k}\right) + 2(k-1) \log\left(\frac{k \log n}{n}\right)\right. \\ & \quad \left. - \frac{4}{5 \log 2} k \log n \left(1 - \frac{k}{\frac{4}{5}n} + o(1)\right) + O(k)\right) \\ & \leq \sum_{k=2}^{|R|} \exp\left(\log n + 2(k-1) \log \log n - 1.15k \log n \left(1 - \frac{5k}{4n} + o(1)\right) + O(k)\right). \end{aligned}$$

For each value of k such that $2 \leq k \leq |R|$, each term above is $o(n^{-1.1})$. Since we are summing over $|R| = (\frac{1}{5} + o(1))n$ many terms, we get $\mathbb{E}[Y] = o(1)$ and so $\Pr[Y > 0] = o(1)$ by Markov's inequality. It follows that the probability in (5) is $o(1)$ and the proof is finished. \square

4. PROOF OF THEOREM 1.3

First, let us concentrate on the lower bound. In the rest of this section, set $d = p(n - 1) = c \log n$ where $0 < c < \frac{1}{\log 2}$. Let $\varepsilon', \varepsilon'' > 0$ be constants such that $c + \varepsilon' + \varepsilon'' < \frac{1}{\log 2}$. Set $c' = c + \varepsilon'$ and $c'' = c' + \varepsilon''$ so that $c < c' < c'' < \frac{1}{\log 2}$. Also define the constant $\gamma := \lceil 2 \left(\frac{4c + 2\varepsilon'}{\varepsilon'^2} \right) \rceil$.

We will need the following property of a random graph $\mathcal{G}(n, p)$.

Lemma 4.1. *The following properties hold w.h.p.*

- (i) $G(n, p)$ has no vertices of degree at least $4 \log n$.
- (ii) $G(n, p)$ has no paths of length γ consisting of vertices of degree at least $c' \log n$.

Proof. (i) follows easily from Chernoff Bound.

For (ii), we first note that the expected number of paths on γ vertices is $O(n^\gamma p^{\gamma-1}) = O(n \cdot (\log n)^{\gamma-1})$. Given such a path, we are looking for each vertex in the path to have at least $c' \log n - 2$ additional neighbours among the $n - \gamma$ other vertices. By Chernoff Bound, the probability that one vertex in the path has enough neighbours is at most

$$\exp \left(- \frac{(\frac{\varepsilon'}{c})^2 c \log n}{2 + \frac{\varepsilon'}{c}} \cdot (1 - o(1)) \right) \leq \exp \left(- \frac{\varepsilon'^2}{4c + 2\varepsilon'} \log n \right).$$

Hence the expected number of paths on γ vertices consisting of vertices with degree at least $c' \log n$, is at most

$$O \left(n \cdot (\log n)^{\gamma-1} \cdot \exp \left(- \frac{\varepsilon'^2}{4c + 2\varepsilon'} \log n \right)^\gamma \right) = O \left(n^{-1} \cdot (\log n)^{\gamma-1} \right) = o(1).$$

So by Markov's inequality, w.h.p., there are no such paths. \square

Now, we are ready to show the lower bound.

Proof of the lower bound in Theorem 1.3. Suppose that in any graph G of maximum degree at most $4 \log n$, the vertex v can acquire $n^{c'' \log 2}$ weight. Then, it follows from Observation 1.2 that $d(v) \geq \log_2 (n^{c'' \log 2}) = c'' \log n$. Furthermore, by averaging argument, some neighbour u of v must have acquired at least $\frac{n^{c'' \log 2}}{4 \log n}$ weight, and so

$$d(u) \geq \log_2 \left[\frac{n^{c'' \log 2}}{4 \log n} \right] \geq c'' \log n - O(\log \log n).$$

Applying the same reasoning inductively $\beta = O(1)$ times, we find a path of length β of vertices of degree at least $c'' \log n - O(\log \log n) \geq c' \log n$.

By Lemma 4.1, w.h.p. $\mathcal{G}(n, p)$ has max degree at most $4 \log n$ and no long path of high degree vertices, and so w.h.p. no vertex can ever get a weight more than $n^{c'' \log 2}$.

Thus, at least $n^{1-c'' \log 2}$ vertices have nonzero weight after any legal sequence of moves and the lower bound holds. \square

For the upper bound in Theorem 1.3, we must show that all the weight can be pushed to at most $n^{1-c+\varepsilon}$ many vertices. To do this we basically follow the embedding proof from Theorem 1.1 but this time with many roots. We sketch the idea of this process, and the reader may check the details.

Proof sketch of the upper bound in Theorem 1.3. Suppose $p = \frac{c+o(1)}{\log 2} \cdot \frac{\log n}{n}$ for some $c \in (0, 1)$. We would like to show that for any $0 < \varepsilon' < \min\{c, 1-c\}$, w.h.p., $a_t(\mathcal{G}(n, p)) \leq n^{1-c+\varepsilon'}$. We let $\varepsilon = \varepsilon'/2$ and prove that $a_t(\mathcal{G}(n, p)) \leq n^{1-c+\varepsilon+o(1)}$. Let $\alpha, \beta > 0$ be any two constants such that $\beta < \frac{c}{10 \log 2}$, $\alpha < \frac{\beta \log 2}{2}$ and let

$$c_j^* := \begin{cases} \beta \log n & \text{if } j \leq \alpha \frac{\log n}{\log \log n} \\ \frac{c-\varepsilon}{\log 2} \log n & \text{otherwise.} \end{cases}$$

Let ρ be the sequence defined as in Definition 3.1 with respect to this sequence c^* . Let m be the largest integer such that $\rho_m \leq n^{c-\varepsilon}$. Then $\rho_m = \Omega(n^{c-\varepsilon}/\log n)$ and $m = (c-\varepsilon+o(1)) \frac{\log n}{\log \log n}$. At this point in the proof of Theorem 1.1, we would adjust the sequence c^* to get a precise value for ρ_m . However, in this case, this step is unnecessary since we may simply adjust the number of copies of T'_{ρ_m} which we embed. Let $L = L(n)$ be an integer such that $L \cdot |T'_{\rho_m}| = L \cdot (1/2+o(1))\rho_m = (\frac{4}{5}+o(1))n$. Then $L = n^{1-c+\varepsilon+o(1)}$. We would like to grow L vertex disjoint copies of T'_{ρ_m} . To do this we must have L roots. We begin with $2L$ many vertices which are candidate roots. The probability that a fixed vertex has less than $\frac{c-\varepsilon/2}{\log 2} \log n$ neighbours (among the other $n-2L$ vertices) is at most $n^{-\Theta(\varepsilon^2)}$ by Chernoff Bound. So by Markov's inequality, w.h.p., at least L of these $2L$ vertices have at least $\frac{c-\varepsilon/2}{\log 2} \log n$ neighbours, and we take these L vertices to be the roots. The other L vertices whose neighbourhoods were exposed now play no part in the embedding and can retain their weight until the end.

We may now proceed as in Lemma 3.8. We embed the trees from the top down and group vertices in the same level (this time, distance from their respective root) consecutively. In levels 0 through $m - \alpha \frac{\log n}{\log \log n} - 1$, the L trees have at most

$$L \cdot \left(\frac{c-\varepsilon}{\log 2} \log n \right)^{m - \alpha \frac{\log n}{\log \log n}} = n^{1-\alpha+o(1)} = o(n)$$

vertices total. So, as in (4), the probability that a fixed vertex is bad is polynomially small. We also have that the number of unexposed vertices is always at least $\frac{1+o(1)}{5}n$, so the polynomial bound on the probability that a vertex is bad holds for levels further down as well. So by taking σ to be a large enough constant, we successfully embed the L trees w.h.p.

We are now in the situation where we have $(\frac{4}{5} + o(1))n$ bereft vertices B , and $(\frac{1}{5} + o(1))$ unexposed vertices, R . Note that since $p = \frac{c+o(1)}{\log 2} \cdot \frac{\log n}{n}$ and c may be small here, we will not be able to guarantee that there is a matching from R to B which saturates all vertices in R . However, it is sufficient to find a matching which

saturates all but $n^{1-c+\varepsilon}$ vertices in R , since these remaining $n^{1-c+\varepsilon}$ vertices can keep their weight. Indeed, if such a matching is found then we have shown that

$$a_t(\mathcal{G}(n, p)) \leq L + L + n^{1-c+\varepsilon} = n^{1-c+\varepsilon+o(1)}.$$

The first L represents the candidate roots which were discarded, the second L represents the roots of the T'_{ρ_m} which were embedded and which receive all the weight from their trees, and the last term represents the unmatched vertices from R .

To show that such a matching exists, we may use the defect version of Hall's Theorem: If $|N(S) \cap B| \geq |S| - q$ for all $S \subseteq R$, then there is a matching which saturates all but q vertices of R . Emulating the proof of Lemma 3.9 using this version of Hall's Theorem with $q = n^{1-c+\varepsilon}$ proves the existence of the desired matching. \square

5. PROOF OF THEOREM 1.4

Before we move to the proof of this result, let us mention that our goal is to provide a simple proof of the conjecture and the constant can be easily improved with more effort.

Proof of Theorem 1.4. We say that a subgraph $L = \{v-w-x-y\}$ of a tree T is a **long leaf** if L is an induced path of length 3; in particular, $\deg(v) = 1$, $\deg(w) = \deg(x) = 2$ in T . Observe that the acquisition number of every graph is bounded from below by the number of long leaves. Indeed, it is straightforward to see that, regardless of a strategy used, for every long leaf L we have that at least one vertex from $\{v, w, x\}$ has to have non-zero weight at the end of the process.

Consider the probability space Ω of all labelled trees of order n uniformly distributed. Let T be a randomly chosen tree from Ω . Clearly $|\Omega| = n^{n-2}$, due to Cayley's formula, so for every fixed tree T_0 on n vertices we have $\Pr(T = T_0) = 1/n^{n-2}$. Our goal is to show that a.a.s. the number of long leaves in T is at least $n/(3e^3)$.

Let $X_v = X_v(T)$ be an indicator random variable defined as follows:

$$X_v = \begin{cases} 1 & \text{if } v \text{ is a vertex of degree 1 in a long leaf,} \\ 0 & \text{otherwise.} \end{cases}$$

Let $X = X(T)$ be a random variable counting the number of long leaves in T , that is, $X = \sum_{v \in V(T)} X_v$. Note that for every $v \in V(T)$

$$\begin{aligned} E(X_v) &= \Pr(X_v = 1) = \frac{(n-1)(n-2)(n-3)(n-3)^{n-5}}{n^{n-2}} \\ &= (1+o(1)) \left(1 - \frac{3}{n}\right)^n = (1+o(1)) \frac{1}{e^3}, \end{aligned}$$

since there are $(n-1)(n-2)(n-3)$ choices for the vertices of the long leaf and there are $(n-3)^{n-5}$ ways to embed a tree on remaining vertices. Hence,

$$E(X) = \sum_{v \in V(T)} E(X_v) = (1+o(1)) \frac{n}{e^3}. \quad (6)$$

Now we are going to apply Chebyshev's inequality to show that a.a.s. $X \geq \frac{E(X)}{2} \geq \frac{n}{3e^3}$. It follows that

$$\Pr\left(X \leq \frac{E(X)}{2}\right) \leq \Pr\left(|X - E(X)| \geq \frac{E(X)}{2}\right) \leq \frac{\text{Var}(X)}{\frac{1}{4}(E(X))^2} = 4\left(\frac{E(X^2)}{(E(X))^2} - 1\right).$$

Hence, it suffices to show that $\frac{E(X^2)}{(E(X))^2}$ tends to 1 as $n \rightarrow \infty$. Clearly,

$$E(X^2) = \sum_{v,v'} E(X_v X_{v'}) = E(X) + \sum_{v \neq v'} E(X_v X_{v'}) = E(X) + \sum_{v \neq v'} \Pr(X_v = X_{v'} = 1), \quad (7)$$

where the sums are over ordered pairs. Now, for fixed vertices $v \neq v'$,

$$\Pr(X_v = X_{v'} = 1) = \frac{(n-2)(n-3)(n-4)(n-5)(n-6)^2(n-6)^{n-8}}{n^{n-2}} = (1+o(1))\frac{1}{e^6}, \quad (8)$$

since there are $(n-2)(n-3)(n-4)(n-5)$ choices for vertices w, x, w', x' in the two corresponding long leaves $L = \{v-w-x-y\}$ and $L' = \{v'-w'-x'-y'\}$, and $(n-6)^2$ choices for y, y' (note that it might happen that $y = y'$ but other than that the two leaves cannot overlap). Consequently, (6), (7), and (8) imply that

$$\frac{E(X^2)}{(E(X))^2} = \frac{(1+o(1))\frac{n}{e^3} + (1+o(1))n^2\frac{1}{e^6}}{\left((1+o(1))\frac{n}{e^3}\right)^2} = 1 + o(1),$$

as required. The proof of the theorem is finished. \square

6. CONCLUDING REMARKS

In this paper, we showed that $p = \frac{1}{\log 2} \cdot \frac{\log n}{n}$ is the sharp threshold for the property $a_t(\mathcal{G}(n, p)) = 1$. However, precise behaviour of the total acquisition number in the critical window is not determined and it is left as an open problem. For example one could study when the total acquisition number becomes 2, or when it becomes $\omega(n)$. We analyzed sparser graphs showing that for $c \in (0, 1)$, w.h.p.

$$\log_n a_t\left(\mathcal{G}\left(n, \frac{c}{\log 2} \cdot \frac{\log n}{n}\right)\right) \sim 1 - c,$$

so the exponent of the total acquisition number is determined up to $o(1)$ term. It also remains to be analyzed and better understood. We have not bothered to explicitly state the $o(1)$ term that can be obtained from our proof, and have no reason to believe it is optimal.

On the other hand, it is not difficult to see when this graph parameter becomes sub-linear. It was already anticipated by West [18, 19] that $a_t(G)$ is linear for $p = c/n$ for any constant $c > 0$ and sub-linear for $p \gg 1/n$. This is true, since for $p = c/n$ we have $\Omega(n)$ isolated vertices w.h.p. (see, for example, [9]), and so the total acquisition number is linear w.h.p. For $p = \omega/n$, where $\omega \rightarrow \infty$ the domination number is known to be equal to $(1+o(1))n \log \omega/\omega = o(n)$ w.h.p. [7], so the total acquisition number is also sub-linear w.h.p. An interesting related open question would be to study total acquisition on the giant component of $\mathcal{G}(n, p)$ for $p = c/n$.

REFERENCES

- [1] N. Alon, M. Krivelevich, B. Sudakov, *Embedding nearly-spanning bounded degree trees*. *Combinatorica* **27** (2007), no. 6, 629–644.
- [2] J.C. Bermond, L. Gargano, A.A. Rescigno, U. Vaccaro, *Fast gossiping by short messages*. *SIAM J. Comput.* **27** (1998), no. 4, 917–941
- [3] B. Bollobás, *Random Graphs*, Second Edition, Cambridge Studies in Advanced Mathematics, 73, 2001.
- [4] S. Boyd, A. Ghosh, B. Prabhakar, D. Shah, *Randomized gossip algorithms*. *IEEE Trans. Inform. Theory* **52** (2006), no. 6, 2508–2530.
- [5] A. M. Frieze, M. Molloy, *Broadcasting in random graphs*. *Discrete Appl. Math.* **54** (1994), no. 1, 77–79.
- [6] A. V. Gerbessiotis, *Close-to-optimal and near-optimal broadcasting in random graphs*. *Discrete Appl. Math.* **63** (1995), no. 2, 129–150.
- [7] R. Glebov, A. Liebenau and T. Szabo, *On the concentration of the domination number of the random graph*. *SIAM J. Discrete Math.* **29** (2015), no. 3, 1186–1206.
- [8] S. M. Hedetniemi, S. T. Hedetniemi, A. L. Liestman, *A survey of gossiping and broadcasting in communication networks*. *Networks* **18** (1988), 319–349.
- [9] S. Janson, T. Łuczak and A. Ruciński, *Random Graphs*, Wiley-Intersci. Ser. Discrete Math. Optim., 2000.
- [10] M. Krivelevich, *Embedding spanning trees in random graphs*. *SIAM J. Discrete Math.* **24** (2010), no. 4, 1495–1500.
- [11] D. E. Lampert and P. J. Slater, *The acquisition number of a graph*, *Congr. Numer.* **109** (1995), 203–210.
- [12] T. D. LeSaulnier, N. Prince, P. Wenger, D. B. West, and P. Worah, *Total acquisition in graphs*, *SIAM J. Discrete Math.* **27** (2013), no. 4, 1800–1819.
- [13] R. Montgomery, *Embedding bounded degree spanning trees in random graphs*, preprint, arXiv:1405.6559
- [14] N. Prince, P. S. Wenger, and D. B. West, *Unit acquisition number*, preprint.
- [15] P. J. Slater and Y. Wang, *Some results on acquisition numbers*, *J. Combin. Math. Combin. Comput.* **64** (2008), 65–78.
- [16] P. J. Slater and Y. Wang, *The competitive-acquisition numbers of paths*, *Congr. Numer.* **167** (2004), 33–43.
- [17] P. S. Wenger, *Fractional acquisition in graphs*, *Discrete Appl. Math.* **178** (2014), 142–148.
- [18] D. B. West, *Acquisition parameters of graphs* (slides), at Graph Theory Workshop in Jinhua, China, December 2013, <http://www.math.uiuc.edu/~west/pubs/cumber.pdf>.
- [19] D. B. West, personal communication.

DEPARTMENT OF MATHEMATICAL SCIENCES, MONTCLAIR STATE UNIVERSITY, MONTCLAIR, NJ 07043, USA

E-mail address: `deepak.bal@montclair.edu`

DEPARTMENT OF MATHEMATICS, WESTERN MICHIGAN UNIVERSITY, KALAMAZOO, MI 49008, USA

E-mail address: `patrick.bennett@wmich.edu`

DEPARTMENT OF MATHEMATICS, WESTERN MICHIGAN UNIVERSITY, KALAMAZOO, MI 49008, USA

E-mail address: `andrzej.dudek@wmich.edu`

DEPARTMENT OF MATHEMATICS, RYERSON UNIVERSITY, TORONTO, ON, CANADA, M5B 2K3

E-mail address: `pralat@ryerson.ca`