

*Large-Deviations/Thermodynamic Approach to Percolation on the Complete Graph**

Marek Biskup, Lincoln Chayes, S. A. Smith

*Department of Mathematics, University of California at Los Angeles,
Los Angeles, California; e-mail: biskup@math.ucla.edu*

Received 13 June 2005; revised 22 March 2006; accepted 18 April 2006

Published online in Wiley InterScience (www.interscience.wiley.com).

DOI 10.1002/rsa.20169

ABSTRACT: We present a large-deviations/thermodynamic approach to the classic problem of percolation on the complete graph. Specifically, we determine the large-deviation rate function for the probability that the giant component occupies a fixed fraction of the graph while all other components are “small.” One consequence is an immediate derivation of the “cavity” formula for the fraction of vertices in the giant component. As a byproduct of our analysis we compute the large-deviation rate functions for the probability of the event that the random graph is connected, the event that it contains no cycles and the event that it contains only small components. © 2007 Wiley Periodicals, Inc. *Random Struct. Alg.*, 00, 000–000, 2007

Keywords: Erdős-Renyi random graph; percolation; large deviations; random forests

1. INTRODUCTION

For physical systems, mean-field theory often provides a qualitatively correct description of “realistic behavior.” The corresponding analysis usually begins with the derivation of so-called mean-field equations that are self-consistent relations involving the physical quantity of primary interest and the various parameters of the model. This approach may be realized and, to some extent, justified mathematically by considering the model on the complete graph, where each constituent interacts with all others.

Correspondence to: M. Biskup

*Supported by NSF (DMS-0306167, DMS-0301795, DMS-0505356).

© 2007 Wiley Periodicals, Inc.

As an example, let us consider the Ising model on a complete graph K_n of n vertices. Here we have a collection of ± 1 -valued random variables $(\sigma_i)_{i=1}^n$ which are distributed according to the probability measure $\mu_n(\{\sigma\}) = e^{-\beta H_n(\sigma)} / Z_{n,\beta}$, where

$$H_n(\sigma) = -\frac{1}{n} \sum_{i,j=1}^n \sigma_i \sigma_j - h \sum_{i=1}^n \sigma_i \quad (1.1)$$

and where β, h are parameters. The relevant physical quantity is the *empirical magnetization*, $m_n(\sigma) = n^{-1} \sum_{i=1}^n \sigma_i$. In terms of this quantity, $H_n(\sigma) = -\frac{1}{2} n [m_n(\sigma)]^2 - h m_n(\sigma)$ and so

$$\mathbb{E}_n(\sigma_1 | \sigma_j : j \neq 1) = \tanh[\beta(m_n(\sigma) + h)] + O(1/n). \quad (1.2)$$

This permits the following ‘‘cavity argument.’’ Supposing that m_n tends, as $n \rightarrow \infty$, to a value m_\star in probability, we have that $m_\star = \lim_{n \rightarrow \infty} \mathbb{E}_n(\sigma_1)$ obeys

$$m_\star = \tanh[\beta(m_\star + h)]. \quad (1.3)$$

This is the *mean-field equation* for the (empirical) magnetization. Of course, the concentration of the law of m_n still needs to be justified; cf [18] for details.

In the context of percolation [21], the relevant mean-field model goes under the name the Erdős-Renyi Random Graph. Here each edge of K_n is independently occupied with probability α/n , where $0 \leq \alpha < \infty$, and vacant with probability $1 - \alpha/n$. The relevant ‘‘physical’’ quantity is the *giant-component density* ϱ_\star , i.e., the limiting fraction of the vertices that belong to the giant component of the graph. The corresponding mean-field equation,

$$\varrho_\star = 1 - e^{-\alpha \varrho_\star}, \quad (1.4)$$

is also readily derived from heuristic ‘‘cavity’’ considerations. As is well known, $\varrho_\star = 0$ is the only solution for $\alpha \leq \alpha_c = 1$, while for $\alpha > \alpha_c$ there is another, strictly positive solution. This solution tends to zero as $\alpha \downarrow \alpha_c$; hence, we may speak of a continuous transition.

While Eqs. (1.3–1.4) are indeed straightforward to derive, matters at the level of mean-field equations are not always satisfactory; the problem being the existence multiple solutions. As it turns out, for the percolation model (as well as the k -core percolation) the proper choice is always the *maximal* solution, but prescriptions of this sort generically fail, e.g., for the Ising model (1.3) with $h < 0$ and, as often as not, whenever there is a first-order transition. Thus, one is in need of an additional principle which determines which of the solutions is relevant.

The existing mathematical approach to these difficulties—e.g., for percolation [20], see also [2, 11, 23], or the k -core [32]—is to work with sufficient precision until the mean-field conclusions are rigorously established. Another approach—which admits some prospects of extendability beyond the complete graph [5, 6]—is to supplement the picture by the introduction of the *mean-field free-energy function*. For the Ising model, this is a function $m \mapsto \Phi_{\beta,h}(m)$ such that

$$\mu_n(m_n(\sigma) \approx m) = e^{-n\Phi_{\beta,h}(m)+o(n)}, \quad n \rightarrow \infty, \quad (1.5)$$

i.e., $m \mapsto \Phi_{\beta,h}(m)$ is the large-deviation rate function for the probability of observing the event $\{m_n(\sigma) \approx m\}$. This spells the end of the story from the perspective of probability and/or theoretical physics: One seeks the minimum of the free energy function, setting its derivative to zero yields the mean-field equations with the irrelevant solutions corresponding to the local extrema which are not absolute minima; see again [18].

The free-energy approach to mean-field problems has met with success in Ising systems and, to some extent, it has been applied to the Potts and random-cluster models [12, 16, 19, 28]. However, no attempt seems to have been made to extend this technology to “purely geometrical” problems on the complete graph, specifically, ordinary percolation or k -core percolation. The purpose of this note is to derive the large-deviation rate function for the event that the random graph contains a fraction ϱ of vertices in “large” components. As we will see, the function has a unique minimum for all α , which coincides with the “correct” solution of (1.4). We do not necessarily claim that the resultant justification of this equation is easier than that which already exists in the literature. However, the picture presented here provides some additional insights into the model, while the overall approach indeed admits the possibility of generalizations.

2. MAIN RESULTS

Consider the set of vertices $\mathcal{V} = \{1, \dots, n\}$ and let $(\omega_{kl})_{1 \leq k < l \leq n}$ be a collection of i.i.d. random variables taking value one with probability p and zero with probability $1 - p$. Let $\mathcal{E} = \mathcal{E}(\omega)$ be the (random) set $\{(k, l) : 1 \leq k < l \leq n, \omega_{kl} = 1\}$. In accord with the standard notation, cf [2, 11], we will use $\mathcal{G}(n, p)$ to denote the undirected graph with vertices \mathcal{V} and edges \mathcal{E} . Of particular interest are the cases where p decays to zero proportionally to $1/n$. Since these are the only problems we will consider, let us set, for once and all, $p = \alpha/n$ for some fixed $\alpha \in [0, \infty)$. We will denote the requisite probability measure by $P_{n,\alpha}$.

In order to state our main theorems, we need to introduce some notation. First, consider the standard entropy function

$$S(\varrho) = \varrho \log \varrho + (1 - \varrho) \log(1 - \varrho) \tag{2.1}$$

and let

$$\pi_1(\alpha) = 1 - e^{-\alpha}. \tag{2.2}$$

In addition, consider the function

$$\Psi(\alpha) = \left(\log \alpha - \frac{1}{2} \left[\alpha - \frac{1}{\alpha} \right] \right) \wedge 0 \tag{2.3}$$

and note that $\Psi(\alpha) < 0$ if and only if $\alpha > 1$. Finally, let us also define

$$\begin{aligned} \Phi(\varrho, \alpha) &= S(\varrho) - \varrho \log \pi_1(\alpha\varrho) \\ &\quad - (1 - \varrho) \log[1 - \pi_1(\alpha\varrho)] - (1 - \varrho)\Psi(\alpha(1 - \varrho)). \end{aligned} \tag{2.4}$$

Then we have




Fig. 1. The graph of the free energy function $\varrho \mapsto \Phi(\varrho, \alpha)$ for four distinct values of α . For $\alpha \leq 1$, the function is minimized by $\varrho = 0$, while for $\alpha > 1$ the unique minimum occurs at $\varrho > 0$. In any case, the minimizer is the maximal positive solution for ϱ_* from (1.4). The dashed portion of the graphs for $\alpha = 1.6$ and 2.4 marks the part where the background contribution, $\Psi(\alpha(1 - \varrho))$, to $\Phi(\varrho, \alpha)$ in (2.4) is strictly positive. This rules out the zero solution to (1.4) for all $\alpha > 1$.

Theorem 2.1. Consider $\mathcal{G}(n, \alpha/n)$ and let \mathcal{V}_r be the set of vertices that are in connected components of size larger than r . Then for every $\varrho \in [0, 1]$,

$$\lim_{\epsilon \downarrow 0} \lim_{n \rightarrow \infty} P_{n,\alpha}(|\mathcal{V}_{\epsilon n}| = \lfloor \varrho n \rfloor)^{1/n} = e^{-\Phi(\varrho, \alpha)}. \tag{2.5}$$

An inspection of Lemma 6.2 reveals that, conditional on $\{|\mathcal{V}_{\epsilon n}| = \lfloor \varrho n \rfloor\}$, with $\epsilon > 0$, there will be only one large component with probability tending to one as $n \rightarrow \infty$.

Figure 1 shows the graph of Φ for various values of α which is archetypal of free-energy functions in complete graph setting. The figure indicates a unique global minimum; direct, albeit arduous differentiation of (2.4) yields the fact that all local extrema satisfy the mean-field equation (1.4). The extremum at $\varrho = 0$ is ruled out for $\alpha > 1$ by noting that, under these conditions, the last term in (2.4) is strictly positive.

The corresponding conclusion also may be extracted from the following probabilistic argument: Let $m = \lfloor \varrho n \rfloor$ and note that $e^{-nS(\varrho)}$ is then the exponential growth-rate of $\binom{n}{m}$. This allows us to write

$$e^{-n\Phi(\varrho, \alpha)} = e^{o(n)} \binom{n}{m} [\pi_1(\alpha\varrho)]^m [1 - \pi_1(\alpha\varrho)]^{n-m} e^{(n-m)\Psi(\alpha(1-\varrho))}. \tag{2.6}$$

Neglecting the Ψ -term (which provides a lower bound on Φ), one sees a quantity reminiscent of binomial distribution. Well-known results on the latter inform us that the right-hand side is exponentially small unless

$$\pi_1(\alpha\varrho) \approx \frac{m}{n}, \tag{2.7}$$

i.e., unless ϱ satisfies the mean-field equation (1.4). If Ψ is set to zero, there are degenerate minima for $\alpha > 1$; however, the Ψ -function will lift the degeneracy and, in fact, create a local *maximum* at $\varrho = 0$ once $\alpha > 1$. Meanwhile, in the region of the maximal solution, Ψ has vanished and the aforementioned approximation is exact.

Remarks 2.2. (1) A closely-related, but different problem to the one treated earlier has previously been studied using large-deviation techniques. Indeed, in [30], O’Connell derived the large-deviation rate function for the event that the *largest* connected component is of size about κn . Note, however, that this does not restrict the total volume occupied by these component. For κ close to ϱ_* from (1.4)—explicitly, as long as the complement of the large component has effective α less than 1—O’Connell’s rate function coincides with ours. But once κ is sufficiently small, this conditioning will lead to the creation of several large components whose total volume is such that their complement is effectively subcritical. Consequently, O’Connell never needs to address the central issue of our proof; namely, the decay rate of the probability that supercritical percolation has no giant components. (This is what gives rise to the term Ψ in (2.4) and the dashed portion of the graph in Fig. 1.) In fact, this rate function is basically a concatenation of many scaled copies of the undashed portion of the graph in Fig. 1.

(2) While the Ψ -term in (2.4) has a nontrivial effect on the large-deviation questions studied here, it does not play any role for events whose probability is of order unity (or is subexponential in n). This is because Ψ “kicks in” only for ϱ away from the minimizing value. This is not the case for the k -core where the corresponding large-deviation analysis (M. Biskup, L. Chayes, and S. A. Smith, unpublished) suggests that the analogous term kicks in right at the minimizer and may even affect the fluctuation scales. One way to bring Ψ out of the “realm of exponentially-improbable” for percolation would be to give each configuration a weight suppressing large component. However, we will not pursue these matters in the present note.

(3) Our control of the rate function is not sharp enough to provide a detailed description of the critical region, i.e., the situations when $\alpha = 1 + O(n^{-1/3})$. The corresponding analysis of the scaling phenomena inside the “critical window” has been performed in [10, 13, 27, 29, 31]. On the other hand, for $\alpha > 1$ one should be able to sharpen the control of the rate function near its minimum to derive a CLT for the fluctuations of the size of the giant component.

Several ingredients enter our proof of Theorem 2.1 which are of independent interest. We state these as separate theorems. The first one concerns the exponential decay rate for the probability that the random-graph is (completely) connected:

Theorem 2.3. *Let K denote the event that $\mathcal{G}(n, \alpha/n)$ is connected. Then*

$$P_{n,\alpha}(K) = (1 - e^{-\alpha})^n e^{O(\log n)}, \quad n \rightarrow \infty, \tag{2.8}$$

where $O(\log n)$ is bounded by a constant times $\log n$ uniformly on compact sets of $\alpha \in [0, \infty)$.

We remark that Theorem 2.3 holds with $e^{O(\log n)}$ replaced by $C(\alpha) + o(1)$, see [4] for a proof. However, the requisite steps seem far in excess of the derivation in Section 3. Furthermore, various pieces of Theorem 2.3 have been discovered, apparently multiple times, in [3, 25, 26, 34, 35]; cf also the discussion following Lemma 3.3.

Next we present a result concerning the event that $\mathcal{G}(n, \alpha/n)$ contains no cycles. Such problems have been extensively studied under the conditions where this probability is $O(1)$, see e.g. [11]. Our theorem concerns the large-deviation properties of this event:

Theorem 2.4. *Let L be the event that $\mathcal{G}(n, \alpha/n)$ contains no cycles. Then*

$$\lim_{n \rightarrow \infty} P_{n,\alpha}(L)^{1/n} = \begin{cases} \alpha \exp\left(-\frac{\alpha}{2} + \frac{1}{2\alpha}\right), & \text{if } \alpha > 1, \\ 1, & \text{otherwise.} \end{cases} \quad (2.9)$$

Strictly speaking, this result is not needed for the proof of our main theorem; it is actually used to derive the exponential decay for the probability of the event that $\mathcal{G}(n, \alpha/n)$ contains only “small” components. Surprisingly, the decay rates for these two events are exactly the same:

Theorem 2.5. *Let L be the event that $\mathcal{G}(n, \alpha/n)$ contains no cycles and let B_r be the event that there are no components larger than r . Then*

$$\lim_{r \rightarrow \infty} \liminf_{n \rightarrow \infty} P_{n,\alpha}(B_r)^{1/n} = \lim_{\epsilon \downarrow 0} \limsup_{n \rightarrow \infty} P_{n,\alpha}(B_{\epsilon n})^{1/n} = \lim_{n \rightarrow \infty} P_{n,\alpha}(L)^{1/n}. \quad (2.10)$$

Update: In the present paper we prove Theorem 2.4 using enumeration and generating-function techniques. Recently, a probabilistic approach has been developed by which we obtain an expansion of $P_{n,\alpha}(L_n)$ to quantities of order unity. One advantage of the new approach is that it also permits the analysis of the conditional measure $P_{n,\alpha}(\cdot | L_n)$; see (M. Biskup, L. Chayes, and S. A. Smith, unpublished).

To finish the discussion of our results, let us give some reason for the word “thermodynamic” in the title. The motivation comes from an analogy with droplet formation in systems at phase transition. Such situations have been studied extensively in the context of percolation [1, 14] and Ising (and Potts) model [7, 8, 15, 17, 22] under the banner of “Wulff construction,” see [9] for a review of these matters.

One of the principal questions underlying Wulff construction is as follows: Compute the probability—and the characteristics of typical configurations carrying the event—that a given fraction of the system is in one thermodynamic state (e.g., liquid) while the rest is in another state (e.g., gas). It turns out that the typical configurations are such that the two phases separate; a droplet of one phase “floats” in the other phase. The requisite probability is then given by a large-deviation expression whose rate function is composed of three parts: the “surface” energy and entropy of the droplet, the rate function for the probability that the droplet is all in one phase, and the rate function for the probability that the complement of the droplet is in the other phase.

In the case under study, the droplet is exactly the giant component and its weight is just the probability that all vertices in the droplet are connected to each other. The surface energy is (the log of) the probability that no vertex inside is connected to any vertex outside; the entropy is (the log of) the number of ways to choose the corresponding number of sites. The weight of the phase outside simply amounts to the probability that all remaining components are of submacroscopic scale. When the leading-order exponential decay rate of all of these contributions is extracted using Theorems 2.3–2.5, we get a quantity that only depends on the fraction of vertices taken by the droplet. The resulting expression is the one on the right-hand side of (2.6).

3. EVERYBODY CONNECTED

The goal of this section is to prove Theorem 2.3. Our proof is based on showing that the probability in (2.8) is exactly the same probability in a related, directed graph problem.

For a collection of vertices $\mathcal{V}_n = \{1, \dots, n\}$ and a set of edge probabilities $(p_{kl})_{1 \leq k < l \leq n}$, let \mathcal{G} be the inhomogeneous undirected random graph over \mathcal{V}_n . Similarly, let $\vec{\mathcal{G}}$ denote the inhomogeneous *directed* complete random graph with the restriction that the two possible (directed) edges between k and l occur independently, each with probability p_{kl} . To keep our notation distinct from the special case $p_{kl} = \alpha/n$ treated throughout this paper, we will write P instead of $P_{n,\alpha}$.

Definition 3.1. *A labeled directed graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ is said to be grounded at vertex $v \in \mathcal{V}$ if for every $w \in \mathcal{V}$ there exists a (directed) path from w to v in \mathcal{E} .*

The identification of the two problems is now stated as follows:

Lemma 3.2. *Let K be the event that \mathcal{G} is connected and let G be the event that $\vec{\mathcal{G}}$ is grounded at vertex “1.” Then $P(K) = P(G)$.*

Proof. We use induction on the total number of edges incident with vertex “ n .” Indeed, if $p_{kn} = 0$ for all $k = 1, \dots, n - 1$, then $P(K) = P(G)$ because both probabilities are zero. Now let us suppose that $P(K) = P(G)$ when $p_{\ell n} = 0$ for all $\ell = k, \dots, n - 1$ and let us prove that it also for $p_{kn} > 0$. It clearly suffices to show that the partial derivatives of $P(K)$ and $P(G)$ with respect to p_{kn} are equal for all $p_{kn} \in [0, 1]$.

Notice first that both K and G are increasing events. Invoking Russo’s formula, see [33] or [21, Theorem 2.25], we obtain

$$\frac{\partial}{\partial p_{kn}} P(G) = P((n, k) \text{ is pivotal for } G), \tag{3.1}$$

where the event $\{(n, k) \text{ is pivotal for } G\}$ means that if (n, k) is occupied, the event G occurs and if not, it does not. (Note that (n, k) denotes the edge going from “ n ” to “ k .”) The conditions under which this event occurs are straightforward: The set $\mathcal{V}_n = \{1, \dots, n\}$ splits into two disjoint components, one rooted at “1” and the other at n , such that no vertex in the component associated with vertex n has an oriented edge to the other component and k has an oriented path to 1. Similarly, we have

$$\frac{\partial}{\partial p_{kn}} P(K) = P((n, k) \text{ is pivotal for } K). \tag{3.2}$$

Here $\{(n, k) \text{ is pivotal for } K\}$ simply means that, if the edge (n, k) is absent, \mathcal{V}_n consists of two connected components, one containing 1 and the other containing n .

To see the equality of partial derivatives, we split both “pivotal” events according to the component containing the vertex n . If \mathcal{W} is a set of vertices such that $n \in \mathcal{W}$ and $1 \notin \mathcal{W}$, let $\mathcal{G}_{n, \mathcal{W}}$ and $\mathcal{G}_{1, \mathcal{W}}$ be the restrictions of \mathcal{G} to \mathcal{W} , and $\mathcal{V}_n \setminus \mathcal{W}$, respectively. Similarly, let $\vec{\mathcal{G}}_{n, \mathcal{W}}$ and $\vec{\mathcal{G}}_{1, \mathcal{W}}$ be the corresponding “components” of the oriented graph. Let $K_{n, \mathcal{W}}$ and $K_{1, \mathcal{W}}$ be the events that $\mathcal{G}_{n, \mathcal{W}}$ and $\mathcal{G}_{1, \mathcal{W}}$ are connected and let $G_{n, \mathcal{W}}$ and $G_{1, \mathcal{W}}$ be the events that $\vec{\mathcal{G}}_{n, \mathcal{W}}$

is grounded at n and that $\vec{\mathcal{G}}_{1,\mathcal{W}}$ is grounded at 1, respectively. Since these pairs of events are independent, we have

$$P((n, k) \text{ is pivotal for } G) = \sum_{\substack{\mathcal{W}: n \in \mathcal{W} \\ 1, k \notin \mathcal{W}}} P(G_{1,\mathcal{W}})P(G_{n,\mathcal{W}})P(C_{\mathcal{W}}) \Big|_{p_{kn}=0}, \tag{3.3}$$

where $C_{\mathcal{W}}$ is the event that no vertex in \mathcal{W} has a (directed) edge to $\mathcal{V}_n \setminus \mathcal{W}$. But the induction assumption tells us that $P(G_{1,\mathcal{W}}) = P(K_{1,\mathcal{W}})$ and $P(G_{n,\mathcal{W}}) = P(K_{n,\mathcal{W}})$, and the symmetry of edge probabilities for the directed graph tells us that $P(C_{\mathcal{W}})$ is the probability that $\mathcal{G}_{n,\mathcal{W}}$ and $\mathcal{G}_{1,\mathcal{W}}$ are not connected by an edge in \mathcal{G} . Substituting these into (3.3), we get the right-hand side of (3.2). This completes the induction step. ■

From now on, let K and G pertain to the specific random graphs $\mathcal{G}(n, \alpha/n)$ and $\vec{\mathcal{G}}(n, \alpha/n)$. We begin with upper and lower bounds on $P_{n,\alpha}(K)$:

Lemma 3.3. $P_{n,\alpha}(K) \leq (1 - (1 - \alpha/n)^{n-1})^{n-1}$.

Proof. Let E be the event—concerning the graph $\vec{\mathcal{G}}(n, \alpha/n)$ —that every vertex except number 1 has at least one outgoing edge. Then $G \subset E$ and so

$$P_{n,\alpha}(G) \leq P_{n,\alpha}(E) = (1 - (1 - \alpha/n)^{n-1})^{n-1}. \tag{3.4}$$

Invoking Lemma 3.2, this proves the desired upper bound. ■

We remark that the upper bound in Lemma 3.3 has been discovered (and rediscovered) several times in the past. It seems to have appeared in [34] for the first time and later in [25, 35] and also [24]. A generalization to arbitrary connected graphs has been achieved in [26].

Lemma 3.4. $P_{n,\alpha}(K) \geq (1 - (1 - \alpha/n)^{n-1})^{n-1} \frac{1}{n}$.

Proof. Consider the following events for directed random graph $\vec{\mathcal{G}}(n, \alpha/n)$: Let E be the event that every vertex, except vertex number 1, has at least one outgoing edge and let F be the event every such vertex has *exactly* one outgoing edge. Since $G \subset E$, we have

$$P_{n,\alpha}(G) = P_{n,\alpha}(E)P_{n,\alpha}(G|E). \tag{3.5}$$

We claim that

$$P_{n,\alpha}(G|E) \geq P_{n,\alpha}(G|F). \tag{3.6}$$

Indeed, let us pick an outgoing edge for each vertex different from 1, uniformly out of all edges going out of that vertex, and let us color these edges red. Let G' be the event that G occurs using only the red edges. The distribution of red edges conditional on E is the same as conditional on F . Hence $P_{n,\alpha}(G|E) \geq P_{n,\alpha}(G'|E) = P_{n,\alpha}(G'|F)$. But, on F , every available edge is red and so $P_{n,\alpha}(G'|F) = P_{n,\alpha}(G|F)$. Combining these inequalities, (3.6) is proved.

The number of configurations that $\vec{\mathcal{G}}(n, \alpha/n)$ can take on F is exactly $(n - 1)^{n-1}$. On the other hand, the number of configurations which result in $\vec{\mathcal{G}}(n, \alpha/n)$ being grounded is $a_n = n^{n-2}$ —the number of labeled trees with n vertices. Hence

$$P_{n,\alpha}(G|F) \geq \frac{n^{n-2}}{(n - 1)^{n-1}} \geq \frac{1}{n}. \tag{3.7}$$

Using that $P_{n,\alpha}(E) = (1 - (1 - \alpha/n)^{n-1})^{n-1}$ the desired bound follows. ■

Proof of Theorem 2.3. The claim is proved by noting

$$\lim_{n \rightarrow \infty} \frac{(1 - (1 - \alpha/n)^{n-1})^{n-1}}{(1 - e^{-\alpha})^{n-1}} = \exp\left(\left(1 - \alpha/2\right) \frac{\alpha e^{-\alpha}}{1 - e^{-\alpha}}\right) \tag{3.8}$$

and using the results of Lemmas 3.3 and 3.4. ■

4. ONLY TREES

Here we will assemble the necessary ingredients for the proof of Theorem 2.4. The proof is based on somewhat detailed combinatorial estimates and arguments using generating functions.

Recall that L denotes the event that $\mathcal{G}(n, \alpha/n)$ contains no cycles and that B_r denotes the event that all components of $\mathcal{G}(n, \alpha/n)$ have no more than r vertices. We begin by a combinatorial representation of the probability $P_{n,\alpha}(L \cap B_r)$: Let a_ℓ denote the number of labeled trees on ℓ vertices. Then

$$\begin{aligned} P_{n,\alpha}(L \cap B_r) &= \sum_{\substack{\sum m_\ell \ell = n \\ m_\ell = 0 \forall \ell > r}} \frac{n!}{\prod_\ell [m_\ell! (\ell!)^{m_\ell}] } \left(\prod_{\ell \geq 1} \left[a_\ell \left(\frac{\alpha}{n}\right)^{\ell-1} \right]^{m_\ell} \right) \left(1 - \frac{\alpha}{n}\right)^{\binom{n}{2} - n + \sum m_\ell} \\ &= n! \left(\frac{\alpha}{n}\right)^n \left(1 - \frac{\alpha}{n}\right)^{\binom{n}{2} - n} \sum_{k=1}^n \left(\frac{\alpha}{n}\right)^{-k} \left(1 - \frac{\alpha}{n}\right)^k Q_{n,k,r}, \end{aligned} \tag{4.1}$$

where we set $k = \sum_\ell m_\ell \ell$, applied the constraint $\sum_\ell \ell m_\ell = n$ and let $Q_{n,k,r}$ denote the sum

$$Q_{n,k,r} = \sum_{\substack{\sum m_\ell \ell = n \\ \sum m_\ell = k \\ m_\ell = 0 \forall \ell > r}} \prod_{\ell \geq 1} \left(\frac{a_\ell}{\ell!}\right)^{m_\ell} \frac{1}{m_\ell!}. \tag{4.2}$$

We begin by isolating the large- n, k behavior of this quantity:

Proposition 4.1. *Consider the polynomial*

$$F_r(s) = \sum_{\ell=1}^r \frac{s^\ell a_\ell}{\ell!} \tag{4.3}$$

Then for all $n, k, r \geq 1$,

$$Q_{n,k,r} \leq \frac{1}{k!} \inf_{s>0} \frac{F_r(s)^k}{s^n}. \tag{4.4}$$

Moreover, for each $\eta > 0$, there is $n_0 < \infty$ and a sequence $(c_r)_{r \geq 1}$ of positive numbers for which

$$Q_{n,k,r} \geq \frac{c_r}{\sqrt{n}} \frac{1}{k!} \inf_{s>0} \frac{F_r(s)^k}{s^n} \tag{4.5}$$

holds for all $n \geq n_0$, all $k \geq 1$ and all $r \geq 2$ such that $k < (1 - \eta)n$ and $rk > n(1 + \eta)$.

Proof of upper bound. Let us consider the generating function

$$\hat{Q}_r(s, z) = 1 + \sum_{n=1}^{\infty} \sum_{k=1}^n Q_{n,k,r} z^k s^n = \exp\{zF_r(s)\}, \tag{4.6}$$

where we used Fubini-Tonelli to derive the second equality. Since F_r is a polynomial, the Cauchy integral formula yields

$$Q_{n,k,r} = \frac{1}{(2\pi i)^2} \oint ds \oint dz \frac{\exp\{zF_r(s)\}}{s^{n+1} z^{k+1}} = \frac{1}{2\pi i} \frac{1}{k!} \oint ds \frac{F_r(s)^k}{s^{n+1}}, \tag{4.7}$$

where all integrals are over a circle of positive radius centered at the origin of \mathbb{C} . Since all coefficients of F_r are nonnegative, $\theta \mapsto |F_r(s e^{i\theta})|$ for $s > 0$ is maximized at $\theta = 0$. Bounding the integrand by its value at $\theta = 0$, the integral yields a factor 2π ; optimizing over $s > 0$ then gives the upper bound in (4.4). ■

Proof of lower bound. As is common in Tauberian arguments, the lower bound will require somewhat more effort. First let us note that under the conditions $k < (1 - \eta)n$ and $rk > n(1 + \eta)$ the function $s \mapsto F_r(s)^k/s^n$, for $s > 0$, blows up both at 0 and ∞ . Its minimum is thus achieved at an interior point; for the rest of this proof we will fix s to a minimizer of this function. Since $|F_r(s e^{i\theta})| < F_r(s)$ for all $\theta \in (-\pi, \pi] \setminus \{0\}$, the part of the integral in (4.7) corresponding to $|\theta| > \epsilon$ is exponentially small (in n) compared to the infimum in (4.5). We thus need to show the lower bound only for the portion of the integral over θ with $|\theta| \leq \epsilon$, for some fixed $\epsilon > 0$.

Since F_r has positive coefficients, $F_r \neq 0$ in the (complex) ϵ -neighborhood of s . This allows us to define the function

$$g(\theta) = \log \frac{F_r(s e^{i\theta})^\varrho}{s e^{i\theta}}, \quad |\theta| \leq \epsilon, \tag{4.8}$$

where ϱ plays the role of k/n . The function g is analytic in an $O(\epsilon)$ -neighborhood of the origin. The choice of s implies that $g'(0) = 0$ which is equivalent to

$$\frac{sF'_r(s)}{F_r(s)} = \frac{1}{\varrho}. \tag{4.9}$$

For the second derivative we get $g''(0) = -\varrho \text{Var}(X)$, where X is the random variable with law

$$P(X = \ell) = \frac{1}{F_r(s)} \frac{a_\ell s^\ell}{\ell!}, \quad \ell = 1, \dots, r. \tag{4.10}$$

In particular, since our restrictions on ϱ between $\frac{1}{r}(1 + \eta)$ and $1 - \eta$ imply that s is bounded away from zero, this law is nondegenerate and so $g''(0) < 0$.

The analyticity of $\theta \mapsto g(\theta)$ for $\theta = O(\epsilon)$ implies that g''' is bounded in this neighborhood, and so by Taylor's theorem we have

$$g(\theta) = g(0) - A\theta^2 + O(\theta^3), \tag{4.11}$$

where $A = A(r, \varrho)$ is positive uniformly in the allowed range of ϱ s, and $O(\theta^3)$ is a quantity bounded by $|\theta|^3$ times a constant, depending only on r, ϵ , and η . (In particular, we may assume that $O(\theta^3)$ is dominated by $\frac{1}{2}A\theta^2$ for $|\theta| \leq \epsilon$.)

We will split the integral over $\theta \in [-\epsilon, \epsilon]$ into two more parts. Let $\delta > 0$ and note that $ng(0)$ is the logarithm of the infimum in (4.5). Then for θ with $\delta n^{-1/3} \leq |\theta| \leq \epsilon$ we have

$$n \operatorname{Re} g(\theta) \leq ng(0) - \frac{1}{2} A \delta^2 n^{1/3} \tag{4.12}$$

which shows that even this portion of the integral brings a contribution that is negligible compared to the right-hand side of (4.5). But for $|\theta| \leq \delta n^{-1/3}$ we have $nO(\theta^3) = O(\delta)$ and so for $\delta \ll 1$, the Taylor remainder will always have imaginary part between, say, $-\pi/4$ and $\pi/4$. This means that

$$\operatorname{Re} \int_{-\delta n^{-1/3}}^{\delta n^{-1/3}} e^{ng(\theta)} d\theta \geq \frac{1}{2} e^{ng(0)} \int_{-\delta n^{-1/3}}^{\delta n^{-1/3}} e^{-nA\theta^2} d\theta \geq \frac{c}{\sqrt{n}} e^{ng(0)} \tag{4.13}$$

for some constant $c > 0$ which may depend on r and η but not on ϱ and n . Combined with the previous estimates, this proves the lower bound (4.5). ■

In light of the above lemma, the k -th term in the sum on the extreme right of (4.1) becomes

$$\alpha^{-k} n^k e^{-\alpha \frac{k}{n}} Q_{n,k,r} = e^{o(n)} \inf_{s>0} \exp\{n\Theta_r(s, k/n)\}, \tag{4.14}$$

where

$$\Theta_r(s, \varrho) = -\varrho \log \alpha - \varrho \log \varrho + \varrho + \varrho \log F_r(s) - \log s. \tag{4.15}$$

Here we should interpret (4.14) as an upper bound for $r = n$ and a lower bound for fixed r . It is clear that, regardless of r , the sum is dominated by $k = \lfloor \varrho n \rfloor$ for which $\varrho \mapsto \inf_{s>0} \Theta_r(s, \varrho)$ is maximal. Such values are characterized as follows:

Lemma 4.2. *Let $\alpha > 0$ and $r \geq 2$. Then there is a unique $(s_r, \varrho_r) \in [0, \infty) \times [1/r, 1]$ for which*

$$\Theta_r(s_r, \varrho_r) = \sup_{1/r \leq \varrho \leq 1} \inf_{s>0} \Theta_r(s, \varrho). \tag{4.16}$$

Moreover, we always have $s_r \in (0, \infty)$ and $\varrho_r \in (1/r, 1)$ and, furthermore,

$$\lim_{r \rightarrow \infty} \Theta_r(s_r, \varrho_r) = \begin{cases} 1 + \alpha/2 - \log \alpha, & \text{if } \alpha \leq 1, \\ 1 + \frac{1}{2\alpha}, & \text{if } \alpha > 1. \end{cases} \tag{4.17}$$

Proof. We begin by ruling out the “boundary values” of s and ϱ . First, if $\varrho = 1/r$, then the infimum over s is actually achieved by $s = \infty$. In that case $F_r(s) = \infty$ and the (one-sided) derivative with respect to ϱ is infinite, i.e., $\varrho = 1/r$ is a strict local minimum of $\varrho \mapsto \inf_{s>0} \Theta_r(s, \varrho)$. Similarly, for $\varrho = 1$ the infimum over $s > 0$ is achieved at $s = 0$ but then the ϱ -derivative of $\varrho \mapsto \inf_{s>0} \Theta_r(s, \varrho)$ is at negative infinity, i.e., also $\varrho = 1$ is a strict local minimum. It follows that any (s_r, ϱ_r) satisfying (4.16) necessarily lies in $(0, \infty) \times (1/r, 1)$.

Setting the partial derivatives with respect to s and ϱ to zero shows that any minimizing pair is the solution of the equations

$$F_r(s) = \alpha \varrho \quad \text{and} \quad sF'_r(s) = \alpha. \tag{4.18}$$

In light of monotonicity of $s \mapsto sF'_r(s)$, the solution is actually unique. To figure out the asymptotic as $r \rightarrow \infty$, we note that for $s \leq 1/e$,

$$sF'_r(s) = \sum_{\ell=1}^r a_\ell \frac{s^\ell}{(\ell-1)!} \xrightarrow{r \rightarrow \infty} W(s), \tag{4.19}$$

where W is the unique number in $[0, 1/e]$ such that $We^{-W} = s$. (Incidentally, W is closely related to the survival probability of the Galton-Watson branching process with Poisson offspring distribution.) If $s > 1/e$, then $sF'_r(s) \rightarrow \infty$ as $r \rightarrow \infty$. Using the relation between $sF'_r(s)$ and α , we thus get

$$s_r \xrightarrow{r \rightarrow \infty} \begin{cases} \alpha e^{-\alpha}, & \text{if } \alpha \leq 1, \\ 1/e, & \text{if } \alpha > 1. \end{cases} \tag{4.20}$$

Integrating the derivative of F_r now shows that $F_r(s_r) \rightarrow \alpha(1 - \alpha/2)$ for $\alpha \leq 1$. Using that $F'_r(s)$ is bounded for $s \leq s_r$, we also find that $F_r(s_r) \rightarrow 1/2$ for $\alpha \geq 1$. This yields

$$\varrho_r \xrightarrow{r \rightarrow \infty} \begin{cases} 1 - \alpha/2, & \text{if } \alpha \leq 1, \\ \frac{1}{2\alpha}, & \text{if } \alpha > 1. \end{cases} \tag{4.21}$$

Noting that $\Theta(s_r, \varrho_r) = \varrho_r - \log s_r$ we now get (4.17). ■

Proof of Theorem 2.4. By the fact that the supremum over ϱ in (4.16) is achieved at an interior point, we can control the difference between the maximizing k/n and its continuous counterpart ϱ . Thence

$$P_{n,\alpha}(L \cap B_r) = q_{n,r} n! \left(\frac{\alpha}{n}\right)^n e^{-\alpha n/2} \exp\{n\Theta_r(s_r, \varrho_r)\}, \tag{4.22}$$

where

$$\frac{\tilde{c}_r}{\sqrt{n}} \leq q_{n,r} \leq n \tag{4.23}$$

for some positive constants \tilde{c}_r which may depend on r and α . Since B_n contains every realization of $\mathcal{G}(n, \alpha/n)$, taking $r = n$ and applying Lemma 4.2 directly shows that $P_{\alpha,n}(L) \leq e^{n\Psi(\alpha)+o(n)}$. To get a corresponding lower bound, we fix $r \geq 2$ and apply $P_{\alpha,n}(L) \geq P_{\alpha,n}(L \cap B_r)$. Taking $1/n$ -th power and letting $n \rightarrow \infty$ then yields

$$\lim_{n \rightarrow \infty} P_{n,\alpha}(L \cap B_r)^{1/n} = \alpha e^{-1-\alpha/2+\Theta_r(s_r, \varrho_r)}. \tag{4.24}$$

As we have just checked, the right-hand side tends to $e^{\Psi(\alpha)}$ as $r \rightarrow \infty$. ■

Corollary 4.3. *We have*

$$\lim_{r \rightarrow \infty} \lim_{n \rightarrow \infty} P_{\alpha,n}(B_r \cap L)^{1/n} = \lim_{n \rightarrow \infty} P_{\alpha,n}(L)^{1/n}. \tag{4.25}$$

Proof. This summarizes the last step of the previous proof. ■

5. NO BIG = NO CYCLES

Here we will prove that absence of large component has a comparable cost to absence of cycles, at least on an exponential scale. To achieve this goal, apart from Corollary 4.3, we will need the following upper bound:

Lemma 5.1. *Let B_r be the event that $\mathcal{G}(n, \alpha/n)$ has no components larger than r , and let L be the event that all connected components of $\mathcal{G}(n, \alpha/n)$ are trees. Then for all $r \geq 1$,*

$$P_{n,\alpha}(B_r) \leq P_{n,\alpha}(L) \left(1 - \frac{\alpha}{n}\right)^{-\frac{1}{2}rn}. \tag{5.1}$$

Proof. Let C be the restriction of $\mathcal{G}(n, \alpha/n)$ to a set $S \subset \{1, \dots, n\}$. Let T be a tree on S . Then

$$\frac{P_{n,\alpha}(C = T)}{P_{n,\alpha}(C \supset T)} = \left(1 - \frac{\alpha}{n}\right)^{\binom{|S|}{2} - |S| + 1} \geq \left(1 - \frac{\alpha}{n}\right)^{\frac{1}{2}|S|^2}. \tag{5.2}$$

Hence

$$P_{n,\alpha}(C \text{ is connected}) \leq \sum_T P_{n,\alpha}(C \supset T) \leq \left(1 - \frac{\alpha}{n}\right)^{-\frac{1}{2}|S|^2} P_{n,\alpha}(C \text{ is a tree}). \tag{5.3}$$

Now, if L_r is the event that no component of $\mathcal{G}(n, \alpha/n)$ of size larger than r has cycles, then $B_r \subset L_r$ and so $P_{n,\alpha}(B_r) \leq P_{n,\alpha}(L_r)$. Let $\{S_j\}$ be a partition of $\{1, \dots, n\}$ and let $P_{n,\alpha}(\{S_j\})$ denote the probability that $\{S_j\}$ are the connected components of $\mathcal{G}(n, \alpha/n)$. Then

$$P_{n,\alpha}(L_r) = \sum_{\{S_j\}} P_{n,\alpha}(\{S_j\}) P_{n,\alpha}(L_r | \{S_j\}), \tag{5.4}$$

where $P_{n,\alpha}(L_r | \{S_j\})$ is the conditional probability of L_r given that $\{S_j\}$ are the connected components of $\mathcal{G}(n, \alpha/n)$.

Letting C_j represent the restriction of $\mathcal{G}(n, \alpha/n)$ to S_j , the bound (5.3) tells us that

$$\begin{aligned} P_{n,\alpha}(L_r | \{S_j\}) &= \prod_{j: |S_j| \geq r} P_{n,\alpha}(C_j \text{ is a tree} \mid C_j \text{ is connected}) \\ &\leq \prod_j P_{n,\alpha}(C_j \text{ is a tree} \mid C_j \text{ is connected}) \prod_{j: |S_j| < r} \left(1 - \frac{\alpha}{n}\right)^{-\frac{1}{2}|S_j|^2} \end{aligned} \tag{5.5}$$

Using that $|S_j| < r$ for every S_j contributing to the second product and applying that the sum of $|S_j|$ over the components with $|S_j| < r$ gives at most n , we then get

$$P_{n,\alpha}(L_r | \{S_j\}) \leq P_{n,\alpha}(L | \{S_j\}) \left(1 - \frac{\alpha}{n}\right)^{-\frac{1}{2}rn}. \tag{5.6}$$

Plugging this back in (5.4), the desired bound follows. ■

Proof of Theorem 2.5. By Lemma 5.1 we have

$$\limsup_{n \rightarrow \infty} P_{n,\alpha}(B_{\epsilon n})^{1/n} \leq e^{\epsilon/2} \lim_{n \rightarrow \infty} P_{n,\alpha}(L)^{1/n}. \tag{5.7}$$

On the other hand, the inclusion $B_r \supset B_r \cap L$ and Corollary 4.3 yield

$$\liminf_{n \rightarrow \infty} P_{n,\alpha}(B_r)^{1/n} \geq \liminf_{n \rightarrow \infty} P_{n,\alpha}(B_r \cap L)^{1/n} \xrightarrow{r \rightarrow \infty} \lim_{n \rightarrow \infty} P_{n,\alpha}(L)^{1/n}. \tag{5.8}$$

Since $P_{n,\alpha}(B_r) \leq P_{n,\alpha}(B_{\epsilon n})$ eventually for any fixed $r \geq 1$ and $\epsilon > 0$, all limiting quantities are equal provided we take $r \rightarrow \infty$ and/or $\epsilon \downarrow 0$ after $n \rightarrow \infty$. ■

6. PROOF OF MAIN RESULT

Before we start proving our main result, Theorem 2.1, we need to ensure that if a large component is present in the graph, then it is unique. The statement we need is as follows:

Lemma 6.1. *Let $K_{\epsilon,2}$ be the event that $\mathcal{G}(n, \alpha/n)$ is either connected or has exactly two connected components, each of which is of size at least ϵn , and recall that K is the event that $\mathcal{G}(n, \alpha/n)$ is connected. Then for all $\alpha_0 > 0$ and $\epsilon_0 > 0$ there exists $c_1 = c_1(\alpha_0, \epsilon_0) < 1$, such that for all $\epsilon \geq \epsilon_0$ and all $\alpha \leq \alpha_0$,*

$$\limsup_{n \rightarrow \infty} P_{\alpha,n}(K^c | K_{\epsilon,2})^{1/n} < c_1. \tag{6.1}$$

Proof. It clearly suffices to show that the ratio of $P_{\alpha,n}(K_{\epsilon,2} \setminus K)$ and $P_{\alpha,n}(K)$ decays to zero exponentially with n , with a rate that is uniformly bounded in $\epsilon \geq \epsilon_0$ and $\alpha \leq \alpha_0$. In light of Theorem 2.3 and the fact that $K_{\epsilon,2}$ only admits components that grow linearly with n , we have

$$\frac{P_{\alpha,n}(K_{\epsilon,2} \setminus K)}{P_{\alpha,n}(K)} = e^{o(n)} \sum_{\epsilon n \leq k \leq n - \epsilon n} \binom{n}{k} \frac{\pi_1(\alpha^{k/n})^k \pi_1(\alpha(1 - k/n))^{n-k}}{\pi_1(\alpha)^n} \left(1 - \frac{\alpha}{n}\right)^{k(n-k)}, \tag{6.2}$$

where $o(n)/n$ tends to zero uniformly in $\epsilon \geq \epsilon_0$ and $\alpha \leq \alpha_0$. Writing ϱ for k/n , the expression under the sum can be bounded by $e^{n[\Xi(\varrho) - \Xi(0)]}$, where

$$\Xi(\varrho) = -S(\varrho) + \varrho \log \pi_1(\alpha\varrho) + (1 - \varrho) \log \pi_1(\alpha(1 - \varrho)) - \alpha\varrho(1 - \varrho). \tag{6.3}$$

Since ϱ is restricted to the interval $[\epsilon, 1 - \epsilon]$, the right-hand side of (6.2) will be exponentially small if we can show $\Xi(\varrho) < \Xi(0)$ for all $\varrho \in (0, 1)$ and all α .

As is easy to check, the function $\varrho \mapsto \Xi(\varrho)$ is symmetric about the point $\varrho = 1/2$. Hence, if we can prove that it is strictly convex throughout $[0, 1]$, then it is maximized at the endpoints. Introducing the function

$$G(\eta) = \eta \log \frac{\pi_1(\eta)}{\eta} \tag{6.4}$$

we have

$$\alpha \Xi(\eta/\alpha) = G(\eta) + G(\alpha - \eta) + \eta(\alpha - \eta). \tag{6.5}$$

In order to prove strict convexity of Ξ , it thus suffices to show that $G''(\eta) + 1 > 0$ for all $\eta > 0$. Introducing yet another abbreviation $q(\eta) = \eta/(1 - e^{-\eta})$, a tedious but straightforward differentiation yields

$$G''(\eta) + 1 = \frac{1}{q}(q' - q)(q e^{-\eta} - 1). \tag{6.6}$$

A direct evaluation now shows that both $q' - q$ and $q e^{-\eta} - 1$ are negative once $\eta > 0$. ■

We will use the above lemma via the following simple conclusion:

Lemma 6.2. *Let N_r denote the number of connected components of $\mathcal{G}(n, \alpha/n)$ of size at least r and let \mathcal{V}_r be the set of vertices contained in these components. Then for all $\alpha > 0$ and $\varrho > \epsilon > 0$ there exists $c = c(\epsilon, \varrho, \alpha) > 0$ such that*

$$P_{\alpha,n}(|\mathcal{V}_{\epsilon n}| = \lfloor \varrho n \rfloor \ \& \ N_{\epsilon n} = 1) \geq (1 - e^{-c n}) P_{\alpha,n}(|\mathcal{V}_{\epsilon n}| = \lfloor \varrho n \rfloor). \tag{6.7}$$

Proof. Clearly, (6.7) will follow if we can prove that

$$P_{\alpha,n}(|\mathcal{V}_{\epsilon n}| = \lfloor \varrho n \rfloor \ \& \ N_{\epsilon n} > 1) \leq e^{-c n} P_{\alpha,n}(|\mathcal{V}_{\epsilon n}| = \lfloor \varrho n \rfloor). \tag{6.8}$$

Let $\mathcal{V}(x)$ denote the connected component of $\mathcal{G}(n, \alpha/n)$ containing the vertex x and let $x \leftrightarrow y$ denote the event that $x, y \in \mathcal{V}_{\epsilon n}$ but $\mathcal{V}(x) \cap \mathcal{V}(y) = \emptyset$. Then (6.8) will be proved once we show

$$P_{\alpha,n}(|\mathcal{V}_{\epsilon n}| = \lfloor \varrho n \rfloor \ \& \ x \leftrightarrow y) \leq e^{-2c n} P_{\alpha,n}(|\mathcal{V}_{\epsilon n}| = \lfloor \varrho n \rfloor). \tag{6.9}$$

(Indeed, the sum over x, y adds only a multiplicative factor of n^2 on the right-hand side.) By conditioning on the set $\mathcal{V}_{\epsilon n}$ and the set $\mathcal{V}(x) \cup \mathcal{V}(y)$, this inequality will in turn follow from

$$P_{\alpha,n}(x \leftrightarrow y \ \& \ \mathcal{V}(x) \cup \mathcal{V}(y) = \mathcal{V}) \leq e^{-2c n} P_{\alpha,n}(\mathcal{V}(x) \cup \mathcal{V}(y) = \mathcal{V}). \tag{6.10}$$

Indeed, let us multiply both sides by the probability that \mathcal{V} is disconnected from the rest of the graph and that all components disjoint from \mathcal{V} of size at least ϵn take the total volume $\lfloor \varrho n \rfloor - |\mathcal{V}|$. The sum over all admissible \mathcal{V} reduces from (6.10) to (6.9).

We will deduce (6.10) from Lemma 6.1. Recall that K is the event that the graph is connected and $K_{\epsilon,2}$ is the event that it has at most two components, each of which is of size at least ϵn . We will now use these events for the restriction of $\mathcal{G}(n, \alpha/n)$ to \mathcal{V} : Let $m = |\mathcal{V}|$, $\tilde{\alpha} = \alpha \frac{m}{n}$ and $\tilde{\epsilon} = \epsilon \frac{n}{m}$. Then we have

$$\{x \leftrightarrow y\} \cap \{\mathcal{V}(x) \cup \mathcal{V}(y) = \mathcal{V}\} \subset K^c \cap K_{\tilde{\epsilon},2}, \tag{6.11}$$

while for the event on the right-hand side of (6.10) we simply get

$$\{\mathcal{V}(x) \cup \mathcal{V}(y) = \mathcal{V}\} = K_{\tilde{\epsilon},2}. \tag{6.12}$$

By Lemma 6.1 and the fact that $\tilde{\alpha} \leq \alpha$ and $\tilde{\epsilon} \geq \epsilon$,

$$P_{\tilde{\alpha},m}(K^c \mid K_{\tilde{\epsilon},2}) \leq e^{-c_1 m}, \tag{6.13}$$

once n is sufficiently large. But $m \geq 2\epsilon n$ and so (6.10) holds with $c = \epsilon c_1$. ■

Now we have finally amassed all ingredients needed for the proof of our main result.

Proof of Theorem 2.1. The case $\varrho = 0$ is quickly reduced to Theorems 2.4–2.5 while $\varrho = 1$ boils down to Theorem 2.3. Thus, we are down to the cases $\varrho \in (0, 1)$. Let $\epsilon \in (0, \varrho)$. By Lemma 6.2, we can focus on the situations with $N_{\epsilon n} = 1$. To make our notation simple, let us assume that ϱn is an integer. Then we have

$$P_{\alpha,n}(|\mathcal{V}_{\epsilon n}| = \varrho n \ \& \ N_{\epsilon n} = 1) = \binom{n}{\varrho n} P_{\varrho n, \alpha \varrho}(K) P_{n-\varrho n, \alpha(1-\varrho)}(B_{\epsilon n}) \left(1 - \frac{\alpha}{n}\right)^{\varrho n(1-\varrho)n}. \tag{6.14}$$

The terms on the right-hand side represent the following: the number of ways to choose the unique component of size qn , the probability that this component is connected, the probability that the complement contains no component of size larger than ϵn and, finally, the probability that the two parts of the graph do not have any edge between them. Invoking Stirling's formula to deal with the binomial term, and plugging explicit expressions for $P_{qn,\alpha q}(K)$ and $P_{n-qn,\alpha(1-q)}(B_{\epsilon n})$ from Theorems 2.3–2.5, the result reduces to a simple calculation. ■

ACKNOWLEDGEMENTS

We thank anonymous referees for advice on style and literature.

REFERENCES

- [1] K. Alexander, J. T. Chayes, and L. Chayes, The Wulff construction and asymptotics of the finite cluster distribution for two-dimensional Bernoulli percolation, *Commun Math Phys* 131 (1990), 1–51.
- [2] N. Alon and J. H. Spencer, The probabilistic method, (Wiley-Interscience Series in discrete mathematics and optimization), 2nd ed., Wiley-Interscience, New York, 2000.
- [3] E. A. Bender, E. R. Canfield, and B. D. McKay, The asymptotic number of labeled connected graphs with a given number of vertices and edges, *Random Struct Alg* 1 (1990), 127–169.
- [4] E. A. Bender, E. R. Canfield, and B. D. McKay, Asymptotic properties of labeled connected graphs, *Random Struct Alg* 3 (1992), 183–202.
- [5] M. Biskup and L. Chayes, Rigorous analysis of discontinuous phase transitions via mean-field bounds, *Commun Math Phys* 238 (2003), 53–93.
- [6] M. Biskup, L. Chayes, and N. Crawford, Mean-field driven first-order phase transitions in systems with long-range interactions, *J Statist Phys* 122 (2006), 1139–1193.
- [7] M. Biskup, L. Chayes, and R. Kotecký, Critical region for droplet formation in the two-dimensional Ising model, *Commun Math Phys* 242 (2003), 137–183.
- [8] T. Bodineau, The Wulff construction in three and more dimensions, *Commun Math Phys* 207 (1999), 197–229.
- [9] T. Bodineau, D. Ioffe, and Y. Velenik, Rigorous probabilistic analysis of equilibrium crystal shapes, *J Math Phys* 41 (2000), 1033–1098.
- [10] B. Bollobás, The evolution of random graphs, *Trans Amer Math Soc* 286 (1984), 257–274.
- [11] B. Bollobás, *Random graphs*, (Cambridge studies in advanced mathematics, Vol. 73), 2nd ed., Cambridge University Press, Cambridge, 2001.
- [12] B. Bollobás, G. Grimmett, and S. Janson, The random-cluster model on the complete graph, *Probab Theor Relat Field* 104 (1996), 283–317.
- [13] C. Borgs, J. T. Chayes, H. Kesten, and J. Spencer, The birth of the infinite cluster: Finite-size scaling in percolation, *Commun Math Phys* 224 (2001), 153–204.
- [14] R. Cerf, Large deviations for three dimensional supercritical percolation, *Astérisque* 267 (2000), vi + 177.
- [15] R. Cerf and A. Pisztora, On the Wulff crystal in the Ising model, *Ann Probab* 28 (2000), 947–1017.
- [16] M. Costeniuc, R. S. Ellis, and H. Touchette, Complete analysis of phase transitions and ensemble equivalence for the Curie-Weiss-Potts model, *J Math Phys* 46 (2005), 063301.

- [17] R. L. Dobrushin, R. Kotecký, and S. B. Shlosman, Wulff construction, A global shape from local interaction, Amer Math Soc, Providence, RI, 1992.
- [18] R. S. Ellis, Entropy, large deviations, and statistical mechanics (Grundlehren der Mathematischen Wissenschaften, Vol. 271), Springer-Verlag, New York, 1985.
- [19] A. Engel, R. Monasson, and A. K. Hartmann, On large deviation properties of Erdős-Rényi random graphs, J Statist Phys 117 (2004), 387–426.
- [20] P. Erdős and A. Rényi, On the evolution of random graphs, Publ Math Inst Hungar Acad Sci 5 (1960), 17–61.
- [21] G. R. Grimmett, Percolation, (Grundlehren der Mathematischen Wissenschaften, Vol. 321), 2nd ed., Springer-Verlag, Berlin, 1999.
- [22] D. Ioffe and R. H. Schonmann, Dobrushin-Kotecký-Shlosman theorem up to the critical temperature, Commun Math Phys 199 (1998), 117–167.
- [23] S. Janson, T. Łuczak, and A. Ruciński, Random graphs, Wiley, New York, 2000.
- [24] R. M. Karp, The transitive closure of a random digraph, Random Struct Alg 1 (1990), 73–93.
- [25] D. E. Knuth and A. Schönhage, The expected linearity of a simple equivalence algorithm, Theor Comput Sci 6 (1978), 281–315.
- [26] M. V. Lomonosov and V. P. Poleskiĭ, An upper bound for the reliability of information networks, Problemy Peredači Informacii 7 (1971), 78–81.
- [27] T. Łuczak, Component behavior near the critical point of the random graph process, Random Struct Alg 1 (1990), 287–310.
- [28] M. Łuczak and T. Łuczak, The phase transition in the cluster-scaled model of a random graph, Random Struct Alg 28 (2006), 215–246.
- [29] T. Łuczak, B. Pittel, and J. C. Wierman, The structure of a random graph at the point of the phase transition, Trans Amer Math Soc 341 (1994), 721–748.
- [30] N. O’Connell, Some large deviation results for sparse random graphs, Probab Theor Relat Field 110 (1998), 277–285.
- [31] B. Pittel, On the largest component of the random graph at a nearcritical stage, J Combin Theor B 82 (2001), 237–269.
- [32] B. Pittel, J. Spencer, and N. Wormald, Sudden emergence of a giant k -core in a random graph, J Combin Theor B 67 (1996), 111–151.
- [33] L. Russo, On the critical percolation probabilities, Z Wahrsch Verw Gebiete 56 (1981), 229–237.
- [34] V. E. Stepanov, The probability of the connectedness of a random graph $\mathcal{G}_m(t)$, (Russian) Teor Verojatnost i Primenen 15 (1970), 58–68.
- [35] A. Takács, A generalization of an inequality of Stepanov, J Combin Theor B 48 (1990), 289–293.