Random subgraphs of finite graphs: III. The phase transition for the n-cube

Christian Borgs* Jennifer T. Chayes* Remco van der Hofstad†
Gordon Slade‡ Joel Spencer§

September 14, 2004

Abstract

We study random subgraphs of the n-cube $\{0,1\}^n$, where nearest-neighbor edges are occupied with probability p. Let $p_c(n)$ be the value of p for which the expected size of the component containing a fixed vertex attains the value $\lambda 2^{n/3}$, where λ is a small positive constant. Let $\epsilon = n(p-p_c(n))$. In two previous papers, we showed that the largest component inside a scaling window given by $|\epsilon| = \Theta(2^{-n/3})$ is of size $\Theta(2^{2n/3})$, below this scaling window it is at most $2(\log 2)n\epsilon^{-2}$, and above this scaling window it is at most $O(\epsilon 2^n)$. In this paper, we prove that for $p - p_c(n) \geq e^{-cn^{1/3}}$ the size of the largest component is at least $\Theta(\epsilon 2^n)$, which is of the same order as the upper bound. The proof is based on a method that has come to be known as "sprinkling," and relies heavily on the specific geometry of the n-cube.

1 Introduction and results

1.1 History

The study of the random graph $G(N,p)$, defined as subgraphs of the complete graph on N vertices in which each of the possible $\binom{N}{2}$ edges is present with probability p, was initiated by Erdős and Rényi in 1960 [13]. They showed that for $p = N^{-1}(1+\epsilon)$ there is a phase transition at $\epsilon = 0$ in the sense that the size of the largest component is $\Theta(\log N)$ for $\epsilon < 0$, $\Theta(N)$ for $\epsilon > 0$, and has the nontrivial behavior $\Theta(N^{2/3})$ for $\epsilon = 0$.

The results of Erdős and Rényi were substantially strengthened by Bollobás [8] and Łuczak [20]. In particular, they showed that the model has a scaling window of width $N^{-1/3}$, in the sense

*Microsoft Research, One Microsoft Way, Redmond, WA 98052, USA. borgs@microsoft.com, jchayes@microsoft.com

†Department of Mathematics and Computer Science, Eindhoven University of Technology, P.O. Box 513, 5600 MB Eindhoven, The Netherlands. rhofstad@win.tue.nl

‡Department of Mathematics, University of British Columbia, Vancouver, BC V6T 1Z2, Canada. slade@math.ubc.ca

§Department of Computer Science, Courant Institute of Mathematical Sciences, New York University, 251 Mercer St., New York, NY 10012, U.S.A. spencer@cs.nyu.edu

that if $p = N^{-1}(1+\Lambda_N N^{-1/3})$ then the size of the largest component is $\Theta(N^{2/3})$ when Λ_N remains uniformly bounded in N, is less than $\Theta(N^{2/3})$ when $\Lambda_N \to -\infty$, and is greater than $\Theta(N^{2/3})$ when $\Lambda_N \to +\infty$. It is also known that inside the scaling window the expected size of the component containing a given vertex is $\Theta(N^{1/3})$.

In this paper, we consider random subgraphs of the n-cube $\mathbb{Q}_n = \{0,1\}^n$, where each of the nearest-neighbor edges is occupied (selected) with probability p. We emphasize the role of the volume (number of vertices) of $\mathbb{Q}_n$ by writing

$$V = |\mathbb{Q}_n| = 2^n. \tag{1.1}$$

This model was first analyzed in 1979 by Erdős and Spencer [14], who showed that the probability that the random subgraph is connected tends to 0 for $p < 1/2$, e^{-1} for $p = 1/2$, and 1 for $p > 1/2$. More interestingly for our purposes, they showed that for $p = n^{-1}(1+\epsilon)$ the size of the largest component is $o(V)$ for $\epsilon < 0$, and conjectured that it is $\Theta(V)$ for $\epsilon > 0$.

The conjecture of Erdős and Spencer was proved in 1982 by Ajtai, Komlós and Szemerédi [4], who thereby established a phase transition at $\epsilon = 0$. Their results apply for $p = n^{-1}(1+\epsilon)$ with ϵ *fixed*. For ϵ fixed and negative, the largest component can be shown by comparison with the Poisson branching process to be a.a.s. of size $O(n)$. (We say that E_n occurs *a.a.s.* if $\lim_{n\to\infty} \mathbb{P}(E_n) = 1$.) On the other hand, for ϵ fixed and positive the largest component is at least of size cV for some positive $c = c(\epsilon)$. To prove the latter, Ajtai, Komlós and Szemerédi introduced a method, now known as "sprinkling," which is very similar to methods introduced at roughly the same time in the context of percolation on $\mathbb{Z}^d$ by Aizenman, Chayes, Chayes, Fröhlich and Russo [2]. We will use a variant of sprinkling in this paper. Very recently, Alon, Benjamini and Stacey [5] used the sprinkling technique to extend the Ajtai, Komlós and Szemerédi result [4] to subgraphs of transitive finite graphs of high girth. These results, while applicable to much more than the n-cube, also hold only for ϵ fixed.

A decade after the Ajtai, Komlós and Szemerédi work, Bollobás, Kohayakawa and Łuczak [10] substantially refined their result, in particular studying the behavior of the largest component as $\epsilon \to 0$. Let $|C(x)|$ denote the size of the component of x, let $\mathcal{C}_{\max}$ denote a component of maximal size, and let

$$|\mathcal{C}_{\max}| = \max\{|C(x)| : x \in \mathbb{Q}_n\} \tag{1.2}$$

denote the maximal component size. We again take $p = n^{-1}(1+\epsilon)$, and now assume that $\epsilon \to 0$ as $n \to \infty$. In [10, Corollary 16, Theorem 28] (with somewhat different notation), it is proved that for $\epsilon \le -(\log n)^2(\log\log n)^{-1}n^{-1/2}$,

$$|\mathcal{C}_{\max}| = \frac{2\log V}{\epsilon^2}\Big(1+o(1)\Big) \quad a.a.s., \tag{1.3}$$

and that for $\epsilon \ge 60(\log n)^3 n^{-1}$,

$$|\mathcal{C}_{\max}| = 2\epsilon V\Big(1+o(1)\Big) \quad a.a.s. \tag{1.4}$$

Thus, $\epsilon \ge 60(\log n)^3 n^{-1}$ is supercritical. In addition, it is shown in [10, Theorem 9] that the right side of (1.3) is an upper bound on $|\mathcal{C}_{\max}|$ if ϵ is defined instead by $p = (n-1)^{-1}(1+\epsilon)$, provided that $\epsilon = o(1)$ and $\epsilon < -e^{-o(n)}$, and hence all p such that $p < (n-1)^{-1} - e^{-o(n)}$ are subcritical.

In recent work [11, 12], we developed a general theory of percolation on connected transitive finite graphs that applies to $\mathbb{Q}_n$ and to various high-dimensional tori. The theory is based on the view that the phase transition on many high-dimensional graphs should have similar features to the phase transition on the complete graph. In particular, the largest component should have size $\Theta(V^{2/3})$ in a scaling window of width $\Theta(V^{-1/3})$, and should have size $o(V^{2/3})$ below the window and size $\Theta(V)$ above the window. Note that the bounds of [10], while much sharper than those established in [4], are still far from establishing this behavior.

We will review the results of [11, 12] in detail below, as they apply to $\mathbb{Q}_n$. These results do not give a lower bound on the largest component above the scaling window, and our primary purpose in this paper is to provide such a bound. We define a critical threshold $p_c(n)$ and prove that the largest component has size $\Theta([p - p_c(n)]nV)$ for $p - p_c(n) \geq e^{-cn^{1/3}}$. This falls short of proving a bound for all p above a window of width $V^{-1/3}$, but it greatly extends the range of p covered by the Bollobás, Kohayakawa and Łuczak bound (1.4).

1.2 The critical threshold

The starting point in [11] is to define the critical threshold in terms of the *susceptibility* $\chi(p)$, which is defined to be the expected size of the component of a given vertex:

$$\chi(p) = \mathbb{E}_p|C(0)|. \tag{1.5}$$

For percolation on $\mathbb{Z}^d$, $\chi(p)$ diverges to infinity as p approaches the critical point from below. On $\mathbb{Q}_n$, the function χ is strictly monotone increasing on the interval $[0, 1]$, with $\chi(0) = 1$ and $\chi(1) = V$. In particular, $\chi(p)$ is finite for all p.

For $G(N, p)$, the susceptibility is $\Theta(N^{1/3})$ in the scaling window. For $\mathbb{Q}_n$, the role of N is played by $V = 2^n$, so we could expect by analogy that $p_c(n)$ for the n-cube should be roughly equal to the p that solves $\chi(p) = V^{1/3} = 2^{n/3}$. In [11], we defined the critical threshold $p_c = p_c(n) = p_c(n; \lambda)$ by

$$\chi(p_c) = \lambda V^{1/3}, \tag{1.6}$$

where λ is a small positive constant. The flexibility in the choice of λ in (1.6) is connected with the fact that the phase transition in a finite system is smeared over an interval rather than occurring at a sharply defined threshold, and any value in the transition interval could be chosen as a threshold. Our results show that p_c defined by (1.6) really is a critical threshold for percolation on $\mathbb{Q}_n$.

For the n-cube, the results of [11, Theorems 1.1, 1.5] and [12] imply that there is a $\lambda_0 > 0$ and a $b_0 > 0$ (depending on λ_0) such that if $0 < \lambda \leq \lambda_0$ then

$$1 - \lambda^{-1}2^{-n/3} \leq np_c(n) \leq 1 + b_0 n^{-1}. \tag{1.7}$$

In particular, $p_c(n) = n^{-1} + O(n^{-2})$ as long as $\lambda \geq n2^{-n/3}$. This asymptotic formula is improved in [18], where it is shown that there exist rational numbers a_i ($i \geq 1$), independent of λ, such that for each $s \geq 1$ and for *every* fixed $\lambda > 0$ (not necessarily small),

$$p_c(n; \lambda) = \sum_{i=1}^{s} a_i n^{-i} + O(n^{-(s+1)}). \tag{1.8}$$

Thus $p_c(n;\lambda)$ has an asymptotic expansion to all orders in n^{-1}, with coefficients that do not depend on λ as long as it is fixed. It follows from (1.7) that $a_1 = 1$, and it is shown in [17] that $a_2 = 1$ and $a_3 = \frac{7}{2}$, so that (1.8) gives

$$p_c(n;\lambda) = \frac{1}{n} + \frac{1}{n^2} + \frac{7}{2n^3} + O(n^{-4}). \tag{1.9}$$

1.3 In and around the scaling window

Given $p \in [0,1]$, let $\epsilon = \epsilon(p) \in \mathbb{R}$ be defined by

$$p = p_c(n) + \frac{\epsilon}{n}. \tag{1.10}$$

We say that p is *below* the window (subcritical) if $\epsilon V^{1/3} \to -\infty$, *above* the window (supercritical) if $\epsilon V^{1/3} \to \infty$, and *inside* the window if $|\epsilon| V^{1/3}$ is uniformly bounded in n. In this section, we summarize and rephrase the results stated in [11, Theorems 1.2–1.5], as they apply to $\mathbb{Q}_n$. We give only a selection of the results obtained in [11], in order to simplify the statements.

The results of [11] are contingent on the *triangle condition*. The triangle condition plays an important role in the analysis of percolation on $\mathbb{Z}^d$ for large d [3, 7, 15], as well as on infinite non-amenable graphs [23]. For $x, y \in \mathbb{Q}_n$, let $\{x \leftrightarrow y\}$ denote the event that x and y are in the same component, and let $\tau_p(x,y) = \mathbb{P}_p(x \leftrightarrow y)$. The *triangle diagram* is defined by

$$\nabla_p(x,y) = \sum_{w,z \in \mathbb{Q}_n} \tau_p(x,w)\tau_p(w,z)\tau_p(z,y). \tag{1.11}$$

The *triangle condition* is the statement that

$$\max_{x,y \in \mathbb{Q}_n} \left[\nabla_{p_c(n)}(x,y) - \delta_{x,y} \right] \le a_0 \tag{1.12}$$

where a_0 is less than a sufficiently small constant. The *stronger triangle condition* is the statement that there are positive constants K_1 and K_2 such that $\nabla_p(x,y) \le \delta_{x,y} + a_0$ uniformly in $p \le p_c(n)$, with

$$a_0 = a_0(p) = K_1 n^{-1} + K_2 \chi^3(p) V^{-1}. \tag{1.13}$$

If we choose λ sufficiently small and n sufficiently large, then the stronger triangle condition implies the triangle condition.

The results stated in the following four theorems were proved in [11] assuming the triangle condition (or the stronger triangle condition for [11, Theorem 1.5]), and the stronger triangle condition was established for $\mathbb{Q}_n$ in [12] if we take $\lambda \le \lambda_0$, with λ_0 sufficiently small, in (1.6).

Theorem 1.1 (Below the window). *Let $\lambda \le \lambda_0$ and $p = p_c(n) - \epsilon n^{-1}$, with $\epsilon \ge 0$ and $\epsilon V^{1/3} \to \infty$ as $n \to \infty$. Then*

$$\chi(p) = \frac{1}{\epsilon}[1 + o(1)], \tag{1.14}$$

$$\frac{1}{3600\epsilon^2} \le |\mathcal{C}_{\max}| \le \frac{2\log V}{\epsilon^2}[1 + o(1)] \quad a.a.s. \tag{1.15}$$

Since $p_c(n) > (n-1)^{-1}$ for large n by (1.9), the upper bound of (1.15) extends the range of p covered by the Bollobás, Kohayakawa and Łuczak upper bound of (1.3) from $p < (n-1)^{-1} - e^{-o(n)}$ to all p below the window.

Theorem 1.2 (Inside the window). *Let $\lambda \le \lambda_0$ and $\Lambda < \infty$. Let $p = p_c + \epsilon n^{-1}$ with $|\epsilon| \le \Lambda V^{-1/3}$. There are finite positive constants b_6, b_7, b_8 such that the following statements hold. If $\omega \ge 1$, then*

$$\mathbb{P}_p\Big(\omega^{-1}V^{2/3} \le |\mathcal{C}_{\max}| \le \omega V^{2/3}\Big) \ge 1 - \frac{b_6}{\omega}, \tag{1.16}$$

and

$$b_7 V^{1/3} \le \chi(p) \le b_8 V^{1/3}. \tag{1.17}$$

The constant b_8 depends on Λ and not λ, and the constants b_6 and b_7 depend on both λ and Λ.

Theorem 1.3 (Above the window). *Let $\lambda \le \lambda_0$ and $p = p_c + \epsilon n^{-1}$ with $\epsilon V^{1/3} \to \infty$. Then*

$$\chi(p) \le 162\epsilon^2 V, \tag{1.18}$$

and, for all $\omega > 0$,

$$\mathbb{P}_p\Big(|\mathcal{C}_{\max}| \ge \omega\epsilon V\Big) \le \frac{\text{const}}{\omega}. \tag{1.19}$$

A refinement of (1.19) will be given in Section 2.1.

To see that there is a phase transition at $p_c(n)$, we need an upper bound on the maximal component size in the subcritical phase and a lower bound in the supercritical phase. The former is given in Theorem 1.1 but the latter is not part of Theorem 1.3.

1.4 Main result

Our main result is the following theorem, which is proved in Section 2. Theorem 1.4 provides the missing lower bound for $\epsilon \ge e^{-cn^{1/3}}$. This restriction on ϵ is an artifact of our proof and we believe the theorem remains valid as long as $\epsilon V^{1/3} \to \infty$; see Conjecture 3.2. To fully establish the picture that there is a scaling window of width $\Theta(V^{-1/3})$, it would be necessary to extend Theorem 1.4 to cover this larger range of ϵ.

Theorem 1.4. *There are $c, c_1 > 0$ and $\lambda_0 > 0$ such that the following hold for all $0 < \lambda \le \lambda_0$ and all $p = p_c + \epsilon n^{-1}$ with $e^{-cn^{1/3}} \le \epsilon \le 1$:*

$$|\mathcal{C}_{\max}| \ge c_1 \epsilon 2^n \quad a.a.s., \tag{1.20}$$

$$\chi(p) \ge (c_1\epsilon)^2 2^n \quad \text{as } n \to \infty. \tag{1.21}$$

It is interesting to examine the approach to the critical point with $|p - p_c(n)|$ of order n^{-s} for different values of s. Our results give a hierarchy of bounds as s is varied. For example, it follows from Theorem 1.1 that for $p = p_c(n) - \delta n^{-s}$ with $s > 0$ and $\delta > 0$,

$$\chi(p) = n^{s-1}\delta^{-1}[1 + o(1)]. \tag{1.22}$$

On the other hand, for $p = p_c(n) + \delta n^{-s}$ with $s > 0$ and $\delta > 0$, it follows from Theorems 1.3 and 1.4 that

$$\chi(p) = \Theta(\delta^2 n^{2(1-s)} 2^n). \tag{1.23}$$

Related bounds follow for $|\mathcal{C}_{\max}|$. Thus there is a phase transition on scale n^{-s} for any $s \ge 1$.

2 Proof of Theorem 1.4

In this section, we prove Theorem 1.4 by showing that there is a $c_1 > 0$ such that when $e^{-cn^{1/3}} \leq \epsilon \leq 1$,

$$|\mathcal{C}_{\max}| \geq c_1 \epsilon 2^n \quad \text{a.a.s.} \tag{2.1}$$

and

$$\chi(p) \geq (c_1\epsilon)^2 2^n. \tag{2.2}$$

The proof of (2.1) is based on the method of sprinkling and is given in Sections 2.1–2.2. The bound (2.2) is an elementary consequence of (2.1) and is proved in Section 2.3.

2.1 The percolation probability

For percolation on $\mathbb{Z}^d$, the value of p for which $\chi(p)$ becomes infinite is the same as the value of p where the percolation probability $\mathbb{P}_p(|C(0)| = \infty)$ becomes positive [1, 21]. For $\mathbb{Q}_n$, there can be no infinite component, and the definition of the percolation probability must be modified. For $p = p_c(n) + \epsilon n^{-1}$ with $\epsilon > 0$, we defined the percolation probability in [11] by

$$\theta_\alpha(p) = \mathbb{P}_p(|C(0)| \geq N_\alpha), \tag{2.3}$$

where

$$N_\alpha = N_\alpha(p) = \frac{1}{\epsilon^2}\Big(\epsilon V^{1/3}\Big)^\alpha \tag{2.4}$$

and α is a fixed parameter in $(0, 1)$. The definition (2.3) is motivated as follows. According to Conjectures 3.2–3.3 (see below), above the window the largest component has size $|\mathcal{C}_{\max}| = 2\epsilon V[1 + o(1)]$ a.a.s., while the second largest has size $|\mathcal{C}_2| = 2\epsilon^{-2} \log V[1 + o(1)]$. According to this, above the window $|\mathcal{C}_2| \ll N_\alpha \ll |\mathcal{C}_{\max}|$, so that a component of size at least N_α should in fact be maximal, and $\theta_\alpha(p)$ should correspond to the probability that the origin is in the maximal component. (The above reasoning suggests the range $0 < \alpha < 3$ rather than $0 < \alpha < 1$, but the analysis of [11] requires the latter restriction.)

Let $0 < \alpha < 1$. The combination of [11, Theorem 1.6] with the verification of the triangle condition in [12] implies that there are positive constants b_9, b_{10} such that

$$b_{10}\epsilon \leq \theta_\alpha(p) \leq 27\epsilon, \tag{2.5}$$

where the lower bound holds when $b_9 V^{-1/3} \leq \epsilon \leq 1$ and the upper bound holds when $\epsilon \geq V^{-1/3}$. In addition, there are positive b_{11}, b_{12} such that if $\max\{b_{12}V^{-1/3}, V^{-\eta}\} \leq \epsilon \leq 1$, where $\eta = \frac{1}{3}\frac{3-2\alpha}{5-2\alpha}$, then

$$\mathbb{P}_p\Big(|\mathcal{C}_{\max}| \leq [1 + (\epsilon V^\eta)^{-1}]\theta_\alpha(p)V\Big) \geq 1 - \frac{b_{11}}{(\epsilon V^\eta)^{3-2\alpha}}. \tag{2.6}$$

In the above statements, the constants b_9, b_{10}, b_{11} and b_{12} depend on both α and λ. Note that although (2.6) does not obtain the precise constant of (1.4) found by Bollobás, Kohayakawa and Łuczak, it does extend the range of p from $p \geq n^{-1} + 60(\log^2 n)n^{-2}$ to $p \geq p_c(n) + 2^{-\eta' n}$ for any $\eta' < \eta$. Also, note that the combination of (2.6) and (2.5) gives a refinement of (1.19).

Let

$$Z_{\geq N_\alpha} = \sum_{x \in \mathbb{Q}_n} I[|C(x)| \geq N_\alpha] \tag{2.7}$$

denote the number of vertices in "moderately" large components. Then $\mathbb{E}_p(Z_{\geq N_\alpha}) = \theta_\alpha(p)V$ and hence, by (2.5), above the window $\mathbb{E}_p(Z_{\geq N_\alpha}) = \Theta(\epsilon V)$. In the proof of [11, Theorem 1.6(ii)], it is shown that

$$\mathbb{P}_p\Big(|Z_{\geq N_\alpha} - V\theta_\alpha(p)| \geq (\epsilon V^\eta)^{-1} V\theta_\alpha(p)\Big) \leq \frac{b_{11}}{(\epsilon V^\eta)^{3-2\alpha}} \tag{2.8}$$

for percolation on an arbitrary finite connected transitive graph that obeys the triangle condition, and hence (2.8) holds for $\mathbb{Q}_n$. For $\epsilon \geq e^{-cn^{1/3}}$ and fixed $\alpha \in (0,1)$, there are therefore positive constants η_1, η_2, A such that

$$\mathbb{P}_p\Big(|Z_{\geq N_\alpha} - V\theta_\alpha(p)| \geq V^{1-\eta_1}\theta_\alpha(p)\Big) \leq AV^{-\eta_2}. \tag{2.9}$$

This shows that $Z_{\geq N_\alpha}$ is typically close to its expected value $V\theta_\alpha(p)$. For V sufficiently large and $e^{-cn^{1/3}} \leq \epsilon \leq 1$, it follows from (2.9) and the lower bound of (2.5) that

$$\mathbb{P}_p\Big(Z_{\geq N_\alpha} \leq \tfrac{1}{2} b_{10}\epsilon V\Big) \leq AV^{-\eta_2}. \tag{2.10}$$

2.2 Sprinkling

In Proposition 2.5 below, we prove the lower bound on $|\mathcal{C}_{\max}|$ of (2.1). The proof of (2.1) is based on the following sketch. Let $p = p_c(n) + \epsilon n^{-1}$ with $e^{-cn^{1/3}} \leq \epsilon \leq 1$. Let $p^+ = \epsilon/(2n)$, and define p^- by $p^- + p^+ - p^- p^+ = p$, so that $p^- = p_c(n) + \epsilon/(2n) + o(\epsilon/n)$. Then a percolation configuration with bond density p can be regarded as the union of two independent percolation configurations having bond densities p^- and p^+. The additional bonds due to the latter are regarded as having been "sprinkled" onto the former. For percolation with bond density p^-, it follows from (2.10) that a positive fraction of the vertices lie in moderately large components. We then use the specific geometry of $\mathbb{Q}_n$, in a crucial way, to argue that after a small sprinkling of additional bonds a positive fraction of these vertices will be joined together into a single giant component, no matter how the vertices in the large components are arranged. Our restriction $\epsilon \geq e^{-cn^{1/3}}$ enters in this last step.

In preparation for Proposition 2.5, we state four lemmas. The first lemma uses a very special geometric property of $\mathbb{Q}_n$. For its statement, given any $X \subseteq \mathbb{Q}_n$ and positive integer d, we denote the ball around X of radius d by

$$B[X,d] = \{y \in \mathbb{Q}_n : \exists x \in X \text{ such that } \rho(x,y) \leq d\}, \tag{2.11}$$

where $\rho(x,y)$ denotes the graph distance between x and y.

Lemma 2.1 (Isoperimetric Inequality). *If $X \subseteq \mathbb{Q}_n$ and $|X| \geq \sum_{i \leq u} \binom{n}{i}$ then*

$$|B[X,d]| \geq \sum_{i \leq u+d} \binom{n}{i}. \tag{2.12}$$

Lemma 2.1 is proved in Harper [16]. Bollobás [9] is a very readable and more modern reference. The result of Lemma 2.1 may be seen to be best possible by taking $X = B[\{v\}, d']$ for any fixed $v \in \mathbb{Q}_n$, so that $B[X,d] = B[\{v\}, d'+d]$. For asymptotic calculations we use the inequality of the following lemma.

Lemma 2.2 (Large Deviation). *For $\Delta > 0$,*

$$\sum_{i \le \frac{n-\Delta}{2}} \binom{n}{i} = \sum_{i \ge \frac{n+\Delta}{2}} \binom{n}{i} \le 2^n e^{-\Delta^2/2n}. \tag{2.13}$$

Proof. The first two terms are equal by the symmetry of Pascal's triangle. Dividing by 2^n, the inequality may be regarded as the large deviation inequality

$$\Pr[S_n \ge \Delta] \le e^{-\Delta^2/2n}, \tag{2.14}$$

where $S_n = \sum_{i=1}^n X_i$ with the X_i independent random variables with $\mathbb{P}(X_i = \pm 1) = \frac{1}{2}$. A simple proof of this basic inequality is given in [6, Theorem A.1.1]. □

Lemma 2.3 (Big Overlap). *Let $\Delta, \epsilon > 0$ satisfy $e^{-\Delta^2/2n} < \frac{\epsilon}{2}$. Let $S, T \subseteq \mathbb{Q}_n$ with $|S|, |T| \ge \epsilon 2^n$. Then*

$$|B[S, \Delta] \cap T| \ge \frac{1}{2}|T|. \tag{2.15}$$

Proof. From Lemma 2.2, $|S| \ge \sum_{i \le (n-\Delta)/2} \binom{n}{i}$. Hence, by Lemma 2.1, $|B[S, \Delta]| \ge \sum_{i \le (n+\Delta)/2} \binom{n}{i}$ (intuitively, we have crossed the equator). Therefore, by Lemma 2.2,

$$|\mathbb{Q}_n \setminus B[S, \Delta]| \le \sum_{i < \frac{n-\Delta}{2}} \binom{n}{i} < \frac{\epsilon}{2} 2^n \le \frac{1}{2}|T|, \tag{2.16}$$

and so $B[S, \Delta]$ must overlap at least half of T. □

Lemma 2.4 (Many Paths). *Let $\Delta, \epsilon > 0$ satisfy $e^{-\Delta^2/2n} < \frac{\epsilon}{2}$. Let $S, T \subseteq \mathbb{Q}_n$ with $|S|, |T| \ge \epsilon 2^n$. Then there is a collection of $\frac{1}{2}\epsilon 2^n n^{-2\Delta}$ vertex disjoint paths from S to T, each of length at most Δ.*

Proof. Set $T_1 = B[S, \Delta] \cap T$. By Lemma 2.3, $|T_1| \ge \frac{1}{2}\epsilon 2^n$. Let $T_2 \subseteq T_1$ be a maximal subset such that no $x, y \in T_2$ are within distance 2Δ of each other. Every $y \in T_1$ must lie in a ball of radius 2Δ around some $x \in T_2$ and each such ball has size at most $n^{2\Delta}$ (using a crude upper bound), so $|T_2| \ge n^{-2\Delta}|T_1|$. For each $x \in T_2$ there is a path of length at most Δ to some $z \in S$ and the paths from $x, y \in T_2$ must be disjoint as otherwise x, y would be at distance at most 2Δ. □

Now we use Lemma 2.4 and (2.10) to prove (2.1).

Proposition 2.5 (Sprinkling). *There are absolute positive constants c_1, β such that*

$$\mathbb{P}(|\mathcal{C}_{\max}| \le c_1 \epsilon 2^n) \le 2^{-\beta n}. \tag{2.17}$$

whenever $e^{-cn^{1/3}} \le \epsilon \le 1$. In particular, $|\mathcal{C}_{\max}| \ge c_1 \epsilon 2^n$ a.a.s.

Proof. As usual, we write $p = p_c(n) + \epsilon n^{-1}$. Let p^- be such that

$$p^- + \frac{\epsilon}{2n} - \frac{\epsilon}{2n} p^- = p. \tag{2.18}$$

Note that $p^- = p_c(n) + \epsilon/(2n) + o(\epsilon/n)$. We consider the random subgraph of $\mathbb{Q}_n$ with probability p as the union of the random subgraph G^- with probability p^- and the random subgraph H (the sprinkling) with probability $\epsilon/(2n)$. Crucially, G^- and H are chosen independently.

Let C_i ($i \in I$) denote the components of G^- of size at least $2^{\alpha n/3}$. Set

$$D = \bigcup_{i \in I} C_i, \tag{2.19}$$

and note that

$$N_\alpha = \frac{1}{\epsilon^{2-\alpha}} V^{\alpha/3} \geq V^{\alpha/3} = 2^{\alpha n/3}. \tag{2.20}$$

Since $|D| \geq Z_{\geq N_\alpha}$ by (2.20), it follows from (2.10) that there is an absolute positive constant c_2 such that

$$\mathbb{P}_{p^-}(|D| \geq c_2 \epsilon 2^n) \geq \mathbb{P}_{p^-}(Z_{\geq N_\alpha} \geq c_2 \epsilon 2^n) \to 1 \tag{2.21}$$

exponentially rapidly in n. Thus we may assume that G^- has $|D| \geq c_2 \epsilon 2^n$. It suffices to show that the probability that at least $\frac{1}{3}|D|$ vertices of D lie in a single component of $G^- \cup H$ tends to 1 exponentially rapidly in n (intuitively, that the sprinkling H joins together the disparate components of G^-). We will prove this by estimating the complementary probability, which we will show is in fact much smaller than exponential in n.

Suppose that it is not the case that at least $\frac{1}{3}|D|$ vertices of D lie in a single component of $G^- \cup H$ (sprinkling fails). Then there exists $J \subseteq I$ such that the union of C_j over $j \in J$ has between $\frac{1}{3}|D|$ and $\frac{2}{3}|D|$ vertices of D. Set $K = I \backslash J$ for convenience and let C_J, C_K denote the union of the C_i over $i \in J$, $i \in K$ respectively. Then $D = C_J \cup C_K$, and each of C_J and C_K has size between $\frac{1}{3}|D|$ and $\frac{2}{3}|D|$. Fix such J, K. Critically, there must be no path from C_J to C_K in H.

Let Δ be such that $e^{-\Delta^2/2n} < \frac{1}{6} c_2 \epsilon$, and set $c_3 = \frac{1}{6} c_2$. By Lemma 2.4, there are at least $c_3 \epsilon 2^n n^{-2\Delta}$ disjoint paths from C_J to C_K in $\mathbb{Q}_n$, each of length at most Δ. Each path is in H with probability at least $(\epsilon/2n)^\Delta$. Disjointness implies independence and the probability that H has none of these paths is at most

$$\left[1 - (\epsilon/2n)^\Delta\right]^{c_3 \epsilon 2^n n^{-2\Delta}} \leq \exp\left[-\epsilon^\Delta c_3 \epsilon 2^n n^{-3\Delta} 2^{-\Delta}\right]. \tag{2.22}$$

The above quantity bounds the probability that sprinkling fails for a particular J, K. Since each component C_i is of size at least $2^{\alpha n/3}$, the number of components $|I|$ is at most $2^{n(1-\alpha/3)}$. The number of choices for J (and hence $K = I \backslash J$) is bounded by 2 to this number. Thus the total probability that sprinkling fails is bounded from above by

$$2^{2^{n(1-\alpha/3)}} \exp\left[-\epsilon^\Delta c_3 \epsilon 2^n n^{-3\Delta} 2^{-\Delta}\right] = \exp\left[(\log 2) 2^{n(1-\alpha/3)} - \epsilon^\Delta c_3 \epsilon 2^n n^{-3\Delta} 2^{-\Delta}\right]. \tag{2.23}$$

Finally, we make some computations to estimate (2.23). Fix a positive $\tau < \alpha/3$. Our restriction on ϵ is due to the fact that we require $e^{-\Delta^2/2n} < \epsilon$ (dropping an unimportant factor $\frac{1}{2}$ on the right hand side) and $\epsilon^\Delta \geq 2^{-\tau n}$. Replacing these inequalities by equalities gives a value $e^{-cn^{1/3}}$ for ϵ, where $c = \tau^{2/3} 2^{-1/3} (\log 2)^{2/3}$. Accordingly, we choose ϵ such that $\epsilon \geq e^{-cn^{1/3}}$, with this value of c. Then we may take $\Delta \sim c' n^{2/3}$, for any constant c' such that $c' > \sqrt{2c}$. Then $\epsilon^\Delta c_3 \epsilon 2^n n^{-3\Delta} 2^{-\Delta} \geq 2^{(1-\tau+o(1))n}$, since the factors $c_3 \epsilon n^{-3\Delta} 2^{-\Delta}$ are absorbed into the $o(1)$. Therefore, as required,

(2.23) is exponentially small (in fact, doubly so). Thus we have shown that the probability that sprinkling fails for a particular J, K is much smaller than the reciprocal of the number $2^{|I|}$ of such J, K, and hence a.a.s. the sprinkling succeeds, no J, K exist, and there is a component of size at least $\frac{1}{3}|D|$. $\square$

2.3 The expected component size

Finally, we prove the lower bound on $\chi(p)$ stated in (2.2). Let $\mathcal{C}_1, \mathcal{C}_2, \dots, \mathcal{C}_{2^n}$ denote the components in $\mathbb{Q}_n$ arranged in decreasing order:

$$|\mathcal{C}_1| = |\mathcal{C}_{\max}| \geq |\mathcal{C}_2| \geq \dots, \tag{2.24}$$

with $\mathcal{C}_i = \varnothing$ if there are fewer than i components. By translation invariance,

$$\begin{aligned}\chi(p) &= 2^{-n} \sum_{x \in \mathbb{Q}_n} \mathbb{E}_p |C(x)| = 2^{-n} \sum_{x \in \mathbb{Q}_n} \mathbb{E}_p \Big[\sum_{i=1}^{2^n} |\mathcal{C}_i| \; I[x \in \mathcal{C}_i] \Big] \\ &= 2^{-n} \mathbb{E}_p \Big[\sum_{i=1}^{2^n} |\mathcal{C}_i|^2 \Big] \geq 2^{-n} \mathbb{E}_p \Big[|\mathcal{C}_{\max}|^2 \Big]. \end{aligned} \tag{2.25}$$

By (2.17),

$$\chi(p) \geq 2^{-n} [c_1 \epsilon 2^n]^2 \mathbb{P}_p(|\mathcal{C}_{\max}| \geq c_1 \epsilon 2^n) = [1 + o(1)](c_1 \epsilon)^2 2^n. \tag{2.26}$$

By redefining c_1 to absorb the factor $[1 + o(1)]$, this completes the proof of Theorem 1.4.

3 Conjectures

We conclude with some conjectures.

We conjecture that the upper bound of (1.15) is actually sharp for all p that are not exponentially close to $p_c(n)$. This was proved in [10] (see (1.3)) for $\epsilon \geq (\log n)^2 (\log \log n)^{-1} n^{-1}$. We also conjecture that this behavior can be extended appropriately to cover a larger range of ϵ, as follows.

Conjecture 3.1. *Let $p = p_c(n) - \epsilon n^{-1}$ with $\lim_{n \to \infty} \epsilon = 0$ and $\lim_{n \to \infty} \epsilon e^{\delta n} = \infty$ for every $\delta > 0$. Then*

$$|\mathcal{C}_{\max}| = \frac{2 \log V}{\epsilon^2} [1 + o(1)] \qquad a.a.s. \tag{3.1}$$

If we assume instead that $\lim_{n \to \infty} \epsilon = 0$ and $\lim_{n \to \infty} \epsilon 2^{n/3} = \infty$, then

$$|\mathcal{C}_{\max}| = \Theta\Big(\frac{2 \log(\epsilon^3 V)}{\epsilon^2} \Big) \qquad a.a.s. \tag{3.2}$$

Note that (3.2) reduces to (3.1) when ϵ is not exponentially small. The asymptotic behavior (3.2), with V replaced by N, is known to apply to the random graph $G(N, p)$ for $N^{-1/3} \ll N(p_c - p) \ll 1$; see [19, Theorem 5.6].

For $\epsilon \geq 60 (\log n)^3 n^{-1}$ with $\epsilon = o(1)$, Bollobás, Kohayakawa and Łuczak [10] proved (see (1.4)) that $|\mathcal{C}_{\max}| = 2\epsilon 2^n (1 + o(1))$ a.a.s. We conjecture that this formula holds for all ϵ above the window. This behavior has been proven for the random graph above the scaling window; see [19, Theorem 5.12].

Conjecture 3.2. *Let $p = p_c(n) + \epsilon n^{-1}$ with $\epsilon > 0$, $\lim_{n\to\infty} \epsilon = 0$ and $\lim_{n\to\infty} \epsilon V^{1/3} = \infty$. Then*

$$|\mathcal{C}_{\max}| = 2\epsilon V[1 + o(1)] \qquad a.a.s., \tag{3.3}$$

$$\chi(p) = 4\epsilon^2 V[1 + o(1)]. \tag{3.4}$$

The constants in the above conjecture can be motivated by analogy to the Poisson branching process with mean λ. There the critical point is $\lambda = 1$ whereas we have a critical point $p_c(n)$. An increase in p beyond $p_c(n)$ by ϵn^{-1} increases the average number of neighbors of a vertex by ϵ, which we believe corresponds to the Poisson branching process with mean $1 + \epsilon$. This process is infinite with probability $\sim 2\epsilon$. When we generate the component of x in $\mathbb{Q}_n$ it cannot, of course, be infinite, but with probability $\sim 2\epsilon$ it will not die quickly. Consider components of $x, y \in \mathbb{Q}_n$ that do not die quickly. We believe that these components will not avoid each other. Rather, all of them will coalesce to form a component $\mathcal{C}_{\max}$ of size $2\epsilon V$. Finally, with probability $\sim 2\epsilon$ a given vertex 0 lies in $\mathcal{C}_{\max}$ and this contributes $\sim (2\epsilon)(2\epsilon V)$ to $\chi(p)$, which we believe is the dominant contribution. Note also that for any positive *constant* ϵ, it is shown in [10, Theorem 29] that $|\mathcal{C}_{\max}| \sim aV$ where a is the probability that the Poisson branching process with mean $1 + \epsilon$ is infinite.

For the random graph $G(N, p)$, outside the scaling window there is an intriguing *duality* between the subcritical and supercritical phases (see [6, Section 10.5], and, for a more general setting, see [22]). Let $p = N^{-1}(1+\epsilon)$ lie above the scaling window for $G(N, p)$, and to avoid unimportant issues, assume that $\epsilon = o(1)$. Almost surely, there is a dominant component of size $\sim 2\epsilon N$. Remove this component from the graph, giving G^-. Then G^- behaves like the random graph in the subcritical phase with probability $p' = p_c(1 - \epsilon)$. In particular, the size of the largest component of G^- (the second largest component of G) is given asymptotically by the size of the largest component for p'.

We believe that this duality holds for random subgraphs of $\mathbb{Q}_n$ as well. Let $\mathcal{C}_2$ denote the second largest component, and set $p = n^{-1}(1 + \epsilon)$. Bollobás, Kohayakawa and Łuczak [10] showed that if $\epsilon \to 0$ and $\epsilon \geq 60(\log n)^3 n^{-1}$ then $|\mathcal{C}_2| \sim (2 \log V)\epsilon^{-2}$ a.a.s. Note that this matches the behavior of $|\mathcal{C}_{\max}|$ for $p = p_c(n) - \epsilon n^{-1}$ given in Conjecture 3.1. The following conjecture, which we are far from able to show using our present methods even for $\epsilon \geq e^{-cn^{1/3}}$, claims an extension of the result of [10] to all p above the window.

Conjecture 3.3. *Let $p = p_c(n) + \epsilon n^{-1}$ with $\epsilon > 0$, $\lim_{n\to\infty} \epsilon = 0$ and $\lim_{n\to\infty} \epsilon V^{1/3} = \infty$. Then as $n \to \infty$, the size of the second largest component $\mathcal{C}_2$ is*

$$|\mathcal{C}_2| = \frac{2 \log(\epsilon^3 V)}{\epsilon^2}[1 + o(1)] \qquad a.a.s. \tag{3.5}$$

Finally, we consider the largest component inside the scaling window. For $G(N, p)$, it is known that the largest component inside the window has size $XN^{2/3}$ where X is a positive random variable with a particular distribution. Similarly, we expect that for $\mathbb{Q}_n$ the largest component inside the window has size $YV^{2/3}$ for some positive random variable Y. Our current methods are not sufficient to prove this.

Acknowledgements

This work began during a conversation at afternoon tea, while RvdH, GS and JS were visiting Microsoft Research. The authors thank Benny Sudakov for bringing the question of the critical point of the n-cube to our attention, and for telling us about the recent paper [5]. The work of RvdH was carried out in part at the University of British Columbia and in part at Delft University of Technology. The work of GS was supported in part by NSERC of Canada.

References

[1] M. Aizenman and D.J. Barsky. Sharpness of the phase transition in percolation models. *Commun. Math. Phys.*, **108**:489–526, (1987).

[2] M. Aizenman, J.T. Chayes, L. Chayes, J. Fröhlich, and L. Russo. On a sharp transition from area law to perimeter law in a system of random surfaces. *Commun. Math. Phys.*, **92**:19–69, (1983).

[3] M. Aizenman and C.M. Newman. Tree graph inequalities and critical behavior in percolation models. *J. Stat. Phys.*, **36**:107–143, (1984).

[4] M. Ajtai, J. Komlós, and E. Szemerédi. Largest random component of a k-cube. *Combinatorica*, **2**:1–7, (1982).

[5] N. Alon, I. Benjamini, and A. Stacey. Percolation on finite graphs and isoperimetric inequalities. *Ann. Probab.*, **32**:1727–1745, (2004).

[6] N. Alon and J.H. Spencer. *The Probabilistic Method*. Wiley, New York, 2nd edition, (2000).

[7] D.J. Barsky and M. Aizenman. Percolation critical exponents under the triangle condition. *Ann. Probab.*, **19**:1520–1536, (1991).

[8] B. Bollobás. The evolution of random graphs. *Trans. Amer. Math. Soc.*, **286**:257–274, (1984).

[9] B. Bollobás. *Combinatorics: set systems, hypergraphs, families of vectors and combinatorial probability*. Cambridge University Press, Cambridge, (1986).

[10] B. Bollobás, Y. Kohayakawa, and T. Łuczak. The evolution of random subgraphs of the cube. *Random Struct. Alg.*, **3**:55–90, (1992).

[11] C. Borgs, J.T. Chayes, R. van der Hofstad, G. Slade, and J. Spencer. Random subgraphs of finite graphs: I. The scaling window under the triangle condition. To appear in *Random Struct. Alg.*

[12] C. Borgs, J.T. Chayes, R. van der Hofstad, G. Slade, and J. Spencer. Random subgraphs of finite graphs: II. The lace expansion and the triangle condition. To appear in *Ann. Probab.*

[13] P. Erdős and A. Rényi. On the evolution of random graphs. *Magyar Tud. Akad. Mat. Kutató Int. Közl.*, **5**:17–61, (1960).

[14] P. Erdős and J. Spencer. Evolution of the n-cube. *Comput. Math. Appl.*, **5**:33–39, (1979).

[15] T. Hara and G. Slade. Mean-field critical behaviour for percolation in high dimensions. *Commun. Math. Phys.*, **128**:333–391, (1990).

[16] L.H. Harper. Optimal numberings and isoperimetric problems on graphs. *J. Combinatorial Theory*, **1**:385–393, (1966).

[17] R. van der Hofstad and G. Slade. Expansion in n^{-1} for percolation critical values on the n-cube and $\mathbb{Z}^n$: the first three terms. To appear in *Combin. Probab. Comput.*

[18] R. van der Hofstad and G. Slade. Asymptotic expansions in n^{-1} for percolation critical values on the n-cube and $\mathbb{Z}^n$. Preprint, (2003).

[19] S. Janson, T. Łuczak, and A. Ruciński. *Random Graphs*. John Wiley and Sons, New York, (2000).

[20] T. Łuczak. Component behavior near the critical point of the random graph process. *Random Structures Algorithms*, **1**:287–310, (1990).

[21] M.V. Menshikov. Coincidence of critical points in percolation problems. *Soviet Mathematics, Doklady*, **33**:856–859, (1986).

[22] M. Molloy and B. Reed. The size of the giant component of a random graph with a given degree sequence. *Combin. Probab. Comput.*, **7**:295–305, (1998).

[23] R.H. Schonmann. Multiplicity of phase transitions and mean-field criticality on highly non-amenable graphs. *Commun. Math. Phys.*, **219**:271–322, (2001).