

Phase Transitions in Partially Structured Random Graphs

Oskar Sandberg*

January 29, 2008

Abstract

We study a one parameter family of random graph models that spans a continuum between traditional random graphs of the Erdős-Rényi type, where there is no underlying structure, and percolation models, where the possible edges are dictated exactly by a geometry. We find that previously developed theories in the fields of random graphs and percolation have, starting from different directions, covered almost all the models described by our family. In particular, the existence or not of a phase transition where a giant cluster arises has been proved for all values of the parameter but one. We prove that the single remaining case behaves like a random graph and has a single linearly sized cluster when the expected vertex degree is greater than one.

1 Random Graphs and Finite Percolation

The $\mathcal{G}(n, p)$ family of random graphs, originally due to Erdős and Rényi, is constructed by letting the vertex set $V = [n] = \{0, 1, \dots, n - 1\}$, and letting every possible edge $\{u, v\}$ belong to the edge set E independently with probability p . Bernoulli bond percolation models, on the other hand, are typically constructed by starting with a finite degree lattice and retaining only the edges therein in the same manner, and not any others (retained edges are called open).

The study of random graphs is thus the study of models without an underlying geometry, whereas percolation models depend heavily on the geometry and structure of the lattice on which they are defined. A reasonable question is to ask what happens if one relaxes the influence of the structure in percolation models – for example by allowing open edges to form between vertices more than one step from each other in the lattice. Such models, known as *long-range percolation*, have been studied previously (see [10] [1] and below for more references).

*The Department of Mathematical Sciences, Chalmers University of Technology and Göteborg University. ossa@math.chalmers.se.

Likewise, starting from the other direction, one may ask what happens to $\mathcal{G}(n, p)$ like graphs when structure is introduced – making some possible edges more likely to appear than others. Results about this can be gleaned from recent generalized random graph models [4], and show how much structure can be introduced while keeping the behavior of the model more or less intact.

In fact, in terms of the distance dependence of the edges, known models for long-range percolation and generalized random graphs come very close to covering the whole spectrum. We will discuss a family of random graph models with a single real parameter α that regulates the influence of an underlying structure. We will see that the cases when $\alpha < 1$ fall in the category of previously analyzed random graphs, while the cases where $\alpha > 1$ fall in the category of long-range percolation models. We present some connectivity results for the final, critical, case where $\alpha = 1$.

1.1 Notation

As is common, we will use G to denote both specific graph realizations and the random graphs, though we strive to make the difference clear by context. Where \mathcal{G} is a random graph family, $G \sim \mathcal{G}$ means that G is distributed according to this family. $C_1 = C_1(G)$ will denote the biggest connected component of G .

As usual, a series of events A_1, A_2, \dots occurring *asymptotically almost surely* (a.a.s.) means that $\lim_{n \rightarrow \infty} \mathbf{P}(A_n) = 1$.

1.2 Organization

The paper is organized as follows. In Section 2 we introduce our “ α -model” of random graphs, and in Section 3 we go through how the behavior of the model varies, discussing the known results for most values of α . Finally Section 4 contains our analytic contribution.

2 The α -model

A major difference between the analysis of random graphs and percolation models is whether it is done in a finite or infinite setting. Questions about percolation are typically asked about the behavior of clusters on a infinite grid – the most basic question being whether an infinite cluster remains open. $\mathcal{G}(n, p)$ is, on the other hand, almost always studied for finite values of n – this for the simple reason that if $n = \infty$ and $p > 0$ the graph is a.s. not locally finite. The typical approach is instead to scale the value of p with n – in particular, sparse random graphs are ones where $p = c/n$ for some fixed c (meaning that expected degree is essentially constant for all n). The question is then, rather than asking whether an infinite cluster exists, to look at the relative size of the largest cluster (compared to n) as a function of c .

We take the latter approach here, using a degree normalizer and studying finite graphs, but note that it largely intersects with normalizer free models in




Figure 1: A realization of the α -model $\mathcal{G}_\alpha(n, c)$ with $\alpha = 1$, $n = 100$, and $c = 2$.

cases where the degree is already limited by the structure. We will also restrict ourselves to a one dimensional geometry. One dimension is not an interesting environment for standard Bernoulli percolation, but long-range percolation can be fruitful here. Our geometry is based around the following metric:

$$d(u, v) = \min(|u - v|, n - |u - v|).$$

This is equivalent to placing the vertices in a ring and using the geodesic distance (see Figure 1).

Definition 2.1. (The α -model) For $\alpha \in [0, \infty]$ the family of random graphs $\mathcal{G}_\alpha(n, c)$ are graphs $G = (V, E)$, where $V = [n]$ and for $u, v \in V$

$$p_{u,v} = \mathbf{P}(\{u, v\} \in E) = \frac{c}{h_{\alpha,n} d(u, v)^\alpha}$$

where $h_{\alpha,n} = \sum_{u \in V: u \neq 0} 1/d(u, 0)^\alpha$, independently for all disjoint $\{u, v\}$.

For $\alpha = 0$ this equivalent to $G(n, p)$ with $p = c/(n - 1)$. When $\alpha = \infty$

$$p_{uv} = \begin{cases} c/2 & \text{if } d(u, v) = 1 \\ 0 & \text{otherwise.} \end{cases}$$

which is standard percolation, for which an infinite cluster cannot exist if $c < 2$.

Random graphs with edge probabilities given by a power-law of the distance are not new, and have appeared in more or less exactly this form elsewhere. See the pioneering work of Aizenman, Newman, and Schulman [10] [1], the ideas of Kleinberg [8], as well as later work by other authors [2], [5].

3 Regimes of the α -model : The Emergence of Structure

3.1 $\alpha = 0$: $\mathcal{G}(n, p)$ Random Graph

When $\alpha = 0$, the distance between the points does not affect connectivity, and the α -model is exactly the same as $\mathcal{G}(n, p)$ with $p = c/(n - 1)$. This case has, of course, been extensively studied, see [6] as well as the book length discussions in [7] [3]. With regard to connectivity, it is known to undergo a phase transition at $c = 1$: the largest connected cluster is of size $\theta(\log n)$ in the subcritical phase, $\theta(n^{2/3})$ in the critical phase, and $\theta(n)$ in the supercritical phase.

Theorem 3.1. (Erdős, Rényi) *Let $G \sim \mathcal{G}(n, p)$ for $p = c/(n - 1)$, and ρ be the survival probability of a Galton-Watson branching process with Poisson(c) offspring distribution. Then*

$$|C_1|/n \xrightarrow{p} \rho$$

as $n \rightarrow \infty$.

Much more is known regarding the distribution sequence of component sizes, see the above books for details.

3.2 $0 < \alpha < 1$: Essentially a Random Graph

When $\alpha > 0$ the geographic structure of the model starts affecting the edges. However, as long as $\alpha < 1$ much of the general behavior is retained. While the random graph results above cannot be directly applied, this regime falls within a more recent general random graph model of Bollobás, Janson, and Riordan [4].

In the BJR model $\mathcal{G}(n, \kappa)$, one is given a “ground space” \mathcal{S} which we take to be $[0, 1]$, and a Borel measurable kernel $\kappa : \mathcal{S} \times \mathcal{S} \mapsto \mathbb{R}^+$. For a sequence of points (x_1, x_2, \dots, x_n) of \mathcal{S} corresponding the n vertices, edges are added independently between each pair with probability:

$$p_{uv} = \min(\kappa(x_u, x_v)/n, 1).$$

The model contains a lot of freedoms which we will not require. For instance the sequence of points may be random, in which case the limiting distribution obviously matters greatly. We will not need this, and indeed may set $x_u = u/n$. We then let:

$$\kappa(x, y) = \begin{cases} c/|x - y|^\alpha & \text{for } x \neq y \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

where $|x - y|$ is again interpreted as circular distance, this time on $[0, 1]$.

In order for the results of BJR to hold, κ must adhere to certain conditions which the authors call being “graphical” (Definition 2.7 in [4]). The troublesome condition for our κ is that it must belong to $L^1(\mathcal{S} \times \mathcal{S}, \mu \times \mu)$ (where

μ in our case is the Lebesgue measure on $[0, 1]$). Clearly this is true only for $\alpha < 1$.

For this case however, the BJR model with the above kernel gives $p_{ij} = c/(n^{1-\alpha}d(i, j)^\alpha)$ which is asymptotically equivalent to the α model but for a slightly different value of c . BJR prove that $\mathcal{G}(n, \kappa)$ behaves more or less like the classical $\mathcal{G}(n, p)$ – in particular, most of the proofs are very similar barring the technical difficulties incurred by the greater generality. With regard to connectivity, the same result as above holds, that the phase transition at which a $\theta(n)$ cluster emerges is at $c = 1$.

Theorem 3.2. (Bollobás, Janson, Riordan) *If G is a random graph of size n from the BJR model with kernel κ , then*

$$|C_1|/n \xrightarrow{p} \rho_\kappa$$

where ρ_κ is the survival probability of a multitype Galton-Watson process with the offspring distribution of x given by a Poisson point process on \mathcal{S} with intensity $\kappa(x, \cdot)$.

Because our geometries are transitive, this reduces to a single-type Galton-Watson process with Poisson offspring. It is thus well known that $\rho > 0$ exactly when the expected number of offspring is greater than 1.

3.3 $\alpha = 1$: The Small World

When $\alpha = 1$, the resulting model no longer falls within the that of BJR since the κ implied by (1) is no longer integrable at 0. On the other hand, it does not fall within the long-range percolation models discussed below, this time since $1/x$ is not integrable at infinity. The connectivity of the resulting “in-between” model has to our knowledge not been studied elsewhere.

This model is of added interest due to the results of Kleinberg regarding small-world models and further work that has followed, see [8] [9]. These show that the value of α is intimately related with possibility of decentralized routing (path-finding) in graphs, and that it is exactly when the relation between distance and edge prevalence is as at $\alpha = 1$ that such routing will find short paths.

In Section 4 below, we prove that this case also undergoes at phase transition equivalent to the previous – when the expected degree of each vertex is greater than one, C_1 has order n , otherwise it is sublinear.

3.4 $1 < \alpha < 2$: Long-Range Percolation

When we move to $\alpha > 1$ the model undergoes another radical change. In this case the normalizing constant $h_{\alpha, n} \rightarrow h_\alpha < \infty$ as $n \rightarrow \infty$. The model is thus very similar to the long-range percolation models studied in the 1980s within mathematical physics (see [10] and [1] for rigorous results). It was found that these models do undergo a phase transition similar to the regimes above: for

c less than a certain value there is no percolation, while for large c it can be shown to occur.

While the theorem is stated for percolation on \mathbb{Z} in [10], the arguments in the proofs use only finite subsets, and may be stated as

Theorem 3.3. (Newman, Schulman) *For $1 < \alpha < 2$ there exists $c < \infty$ and $\phi > 0$ such that for $\mathcal{G}_\alpha(n, c)$*

$$|C_1| \geq \phi n \quad a.a.s.$$

In all the regimes of the α -model discussed in Sections 3.1 – 3.3, the critical value of c was found to be one. This is not possible when $\alpha > 1$.

Proposition 3.4. *For any $\alpha > 1$, there exists an $\epsilon > 0$ such that for $c = 1 + \epsilon$ in $\mathcal{G}_\alpha(n, c)$ $|C_1|/n \rightarrow 0$ a.a.s.*

Proof. Let $\epsilon = 1/(6h_\alpha^3)$. For a given graph, let $K(v)$ be the clustering number of a vertex v . That is

$$K(v) = \text{number of 3-cycles containing } v$$

For a given vertex v , let Y be the number of second order neighbors.

$$Y \leq \sum_{u \in N(v)} (|N(v)| - 1) - K(v)$$

which implies

$$\mathbf{E}[Y] \leq \mathbf{E}[|N(v)|] \mathbf{E}[|N(u)|] - \mathbf{E}[K(v)].$$

Since $\mathbf{E}[K(v)] \geq \mathbf{P}(v \leftrightarrow v+1, v \leftrightarrow v+2, v+1 \leftrightarrow v) = 1/(2h_\alpha k^3)$, it holds that $\mathbf{E}[Y] < 1$. It follows that any two-step exploration process on the graph is dominated by a subcritical Galton-Watson process. \square

3.5 $\alpha = 2$: The Second Critical Value

The situation when $\alpha = 2$ is also handled in [10] and [1]. Here it turns out that for the α -model, as we have defined it, there is no giant cluster unless c is large enough that $p_{u,u+1} = 1$ (in which case everything is of course trivially connected). However, the authors show that there are in fact distributions with this tail decay where you will get a giant cluster, however, it is necessary that:

$$\liminf_{d \rightarrow \infty} d^\alpha p_{u,u+d} \geq 1$$

In our case, this implies that $c = h_\alpha$, which of course is the trivial case where $p_{i,i+1} = 1$. For more general long-range percolation formulations one can define non-trivial situations which do percolate with this tail.

3.6 $2 < \alpha < \infty$: Essentially Percolation

When $\alpha > 2$, the model shows behavior similar to standard percolation. A way to see this is to consider a renormalization of the vertex space into large blocks. For $c < h_\alpha$ let $B_1, B_2, \dots, B_{n/m}$ be contiguous blocks of vertices, each of size m . Let $X_{i,j}$ be the number of edges between blocks i and j .

$$\mathbf{E}[X_{i,i+1}] \leq \sum_{x=-\infty}^0 \sum_{y=1}^{\infty} \frac{1}{h_\alpha |x-y|^\alpha} \leq \infty.$$

Since $X_{i,i+1}$ is the sum of independent events each occurring with probability smaller than one, that it has bounded expectation means that

$$\mathbf{P}(X_{i,i+1} > 0) \leq p < 1$$

where p is independent of n and m . On the other hand, if $|i-j| > 1$, then

$$\mathbf{P}(X_{i,j} > 0) \leq \mathbf{E}[X_{i,j}] < h'/m^{\alpha-2} \rightarrow 0$$

as $m \rightarrow \infty$ (h' is a constant).

This means that if m is sufficiently large, and we view two blocks as connected if there is any edge between them, then the system of connected blocks will look more or less like standard Bernoulli percolation.

Formally, let $k = \log n$ and $m = (\log n)^{3/(\alpha-2)}$. Then the integral bounds give that for any vertex x

$$\begin{aligned} \mathbf{P}(x \text{ is connected to at least } mk \text{ vertices}) &\leq \\ \mathbf{P}(k \text{ adjacent blocks are connected in either direction from } x) &+ \\ + \mathbf{P}(\text{one of those } 2k \text{ blocks has a non-adjacent connection}) &\leq 1/\log n. \end{aligned}$$

for n sufficiently large. It follows that $\mathbf{E}[\#\text{ vertices in clusters larger than } mk] < n/\log n$ whence $\mathbf{P}(|C_1|/n > \rho) \rightarrow 0$ for any $\rho > 0$.

3.7 $\alpha = \infty$: Percolation

As noted above, at $\alpha = \infty$ the α -model is exactly percolation with probability $c/2$ that each edge is open.

4 Analysis of case $\alpha = 1$

In this section, we prove a result similar to Theorems 3.1 and 3.2 for the ‘‘Small World’’ case where $\alpha = 1$.

To $\mathcal{G}_1(n, c)$ we associate a Galton-Watson branching process $\{Z_t^n(c)\}_{t \in \mathbb{Z}^+}$ where the child distribution is the same as the marginal distribution of each vertices degree in $\mathcal{G}_1(n, c)$. Let $\rho_n = \rho_n(c)$ be the survival probability of this process, which we know tends to 0 if $c \leq 1$ and a positive value otherwise.

By the “law of rare events”, the distribution δ_i converges to a Poisson(c) distribution. Before proceeding we prove that this implies that $\rho_n(c) \rightarrow \rho(c)$ as $n \rightarrow \infty$, where $\rho(c)$ is the survival probability of GW process with Poisson(c) offspring. This continuity result isn’t new¹ and seems to be assumed by some works on random graphs, but since we need it explicitly at several points below, we include it here.

Lemma 4.1. *Let for each $k \in \mathbb{N}$, $\{Z_i^k\}_{i=0}^\infty$ be a Galton-Watson branching process with offspring given by a non-degenerate distribution with probability generating function f_k and extinction probability q_k (so that $q_k = f(q_k)$).*

If $f_k \rightarrow f$ as $n \rightarrow \infty$ pointwise, then $q_k \rightarrow q$, the smallest value in $[0, 1]$ for which $f(q) = q$.

Proof. Recall that the probability generating functions involved are convex, and take the value 1 at 1. Choose a subsequence k_i such that $q_{k_i} \rightarrow \bar{q}$.

$$\bar{q} \leftarrow q_{k_i} = f_{k_i}(q_{k_i}) \rightarrow f(\bar{q})$$

all as $n \rightarrow \infty$. It follows that \bar{q} is a fixed point of f . If $q = 1$ then that is f ’s only fixpoint in $[0, 1]$, and it follows directly that $\bar{q} = q$ and the result is established. If $q < 1$, f now has two fixpoints in $[0, 1]$ by convexity: 1 and q . We must rule out the case $\bar{q} = 1$.

Assume that $q_{k_i} \rightarrow 1$. Let $q < s < q_{k_i}$, which implies that $f_{k_i}(s) > s$. Letting $i \rightarrow \infty$ for all $q < s < 1$, $f(s) \geq s$. But then there cannot exist a fixpoint $q < 1$ such that $f(q) = q$, which contradicts our assumption. \square

Theorem 4.2. *If $G \sim \mathcal{G}_1(n, c)$ then*

$$\frac{|C_1(G)|}{n} \xrightarrow{p} \rho \text{ as } n \rightarrow \infty$$

The proof largely follows the proof for Erdős-Rényi type graphs (see [7], and also [4]). We follow the proof of Lemma 9.6 in the latter without significant deviation up to the proof of Claim 2 in the latter half). The difference is that, because of the clustering, the branching process coupling breaks down sooner. Therefore we need a different argument for why all “large” clusters are in fact the same.

Proof. Choose c' and ϵ such that $1 < c' < c$ and $0 < \epsilon < 1 - 1/c'$.

We construct the graph using the following well known coupling: For every pair of vertices u and v , let $U_{u,v}$ be a random variable uniformly distributed on $[0, 1]$. We add an edge between u and v if $U_{u,v} < p_{u,v}$.

Let $G \sim \mathcal{G}_1(n, c')$ constructed in this manner. Later we will increase the $p_{u,v}$ by

$$\frac{\delta}{h_{1,n}d(u, v)} \tag{2}$$

¹Thanks to Peter Jagers for helping me with this proof.

where $\delta = c - c'$. The resulting graph is distributed the same as $\mathcal{G}_1(n, c)$.

Note that

$$\log(n) \leq h_{1,n} \leq 2 \log(n) \quad (3)$$

Consider a standard exploration process on G starting at a vertex x . We terminate the exploration either when the explored set becomes larger than some function $\omega(n) \leq n^\epsilon$ where $\omega(n) \rightarrow \infty$ as $n \rightarrow \infty$ or when the exploration process dies. Following [4] (but with a modified definition) we call such functions *admissible* if for any $\gamma > 0$

$$\omega(n)/n^\gamma \rightarrow 0$$

as $n \rightarrow \infty$. Let the set $B = B_\omega$ be set of x for which the process stopped for the former reason (B thus contains all the vertices in components larger than $\omega(n)$).

Since $\omega(n) \ll n^\epsilon$, the exploration may be coupled between two branching processes. From above, we can couple it with $Z^n(c')$, and from below by a similar process but where the offspring are given a random variable Y , the degree of x only counting neighbors more than $\omega(n)$ steps away (since the worst case is that we have already explored the $\omega(n)$ nearest vertices). Using (3) this gives

$$\mathbf{E}[Y] \geq c' \left(1 - \frac{\log \omega(n)}{\log n} \right) \geq c'(1 - \epsilon).$$

Let $\rho' = \rho'(c')$ be the survival probability of this process. It then follows that,

$$\rho' \leq \mathbf{P}(x \in B) \leq \rho(c') + o(1).$$

By selecting ϵ sufficiently small, we can make ρ' arbitrarily close to $\rho(c')$ (Lemma 4.1), while the coupling still holds for n sufficiently large. Thus $\mathbf{P}(x \in B) \rightarrow \rho(c')$ as $n \rightarrow \infty$. Addition over all the vertices gives

$$\frac{1}{n} \mathbf{E}|B| \rightarrow \rho(c'). \quad (4)$$

What remains is to show two things:

1. For all admissible $\omega(n)$, $|B|/n \xrightarrow{p} \rho(c')$ as $n \rightarrow \infty$.
2. For some admissible $\omega(n)$, B consists of only one component.

To prove the first claim, we note that the derivation of (4) did not depend on the choice of $\omega(n)$, and thus holds for all admissible functions. Given such a function, let $\omega'(n)$ be one strictly larger, and let B and B' be their respective sets of vertices in large components (note that $B' \subset B$).

$$\frac{\mathbf{E}|B \setminus B'|}{n} = \frac{\mathbf{E}|B| - \mathbf{E}|B'|}{n} \rightarrow 0. \quad (5)$$

It follows that if $|B|/n \xrightarrow{p} \rho(c')$ holds for B , it must also hold for B' , since

$$\begin{aligned} \mathbf{P}\left(\left|\frac{|B'|}{n} - \rho\right| > \epsilon_1\right) &= \mathbf{P}\left(\left|\frac{|B| - |B \setminus B'|}{n} - \rho\right| > \epsilon_1\right) \\ &\leq \mathbf{P}\left(\left|\frac{|B|}{n} - \rho\right| > \epsilon_1/2\right) \\ &\quad + \mathbf{P}\left(\frac{|B \setminus B'|}{n} > \epsilon_1/2\right) \rightarrow 0 \end{aligned}$$

The second term of the convergence follows by (5) and the first moment method.

Now let $\omega(n) \leq \log \log(n)$. We will show the claim for this $\omega(n)$, and use the previous result to establish it for any faster growing $\omega(n)$. We now explore from two vertices x and y . Start the exploration from x first. At the end of this, we have found a connected subset $C(x)$ of vertices around x . Because both the expected degree of each vertex and the variance is constant, a Chebyshev bound shows that the probability that we should encounter a vertex with more than $\omega(n)$ neighbors is $o(1)$. Thus we can assume that $|C(x)| \leq 2\omega(n)$.

$$\mathbf{P}(y \in C(x)) \leq \frac{2\omega(n)}{\log n} = o(1)$$

since at each step of the exploration, the probability that any vertex which is not y is connected to it is less than $1/2 \log n$ by (3). Next we explore from y and until we have constructed a $C(y)$. In each step of the exploration, the probability we draw a vertex in $C(x)$ next is bounded from above by $\log \log \log n / \log n$, so

$$\mathbf{P}(C(x) \cap C(y) \neq \emptyset) \leq 2 \log \log n \frac{\log \log \log n}{\log n} + o(1) = o(1).$$

It follows that:

$$\rho' \rho' - o(1) \leq P(x, y \in B) \leq \rho(c') \rho(c') + o(1)$$

whence $P(x, y \in B) = \rho^2(c')$ and

$$\frac{1}{n^2} \mathbf{E}[|B|^2] \rightarrow \rho^2$$

as $n \rightarrow \infty$. This means that $\mathbf{Var}(|B|/n) \rightarrow 0$ and thus $|B|/n \xrightarrow{p} \rho(c')$. The first claim is thus established.

For the second claim, we will add the additional edges that we withheld in the beginning by increasing the threshold for edge existence by (2), and show that this connects all large clusters. We start by letting $\omega(n) = \log^4(n)$ and B be as before. We condition on the graph G constructed with the c' threshold, which we may assume has $|B| \geq (\rho - \epsilon_2)n$ (for ϵ_2 arbitrarily small) by the above. From G we create the graph G^c by completing all the connected clusters of G -

note that while this adds edges, it does not change the connectivity properties of the graph. In particular, if adding the additional edges makes B a connected component in G^c , it does so also in G .

Now select from B as many subsets K_1, K_2, \dots, K_m as possible, such that each K_i is a clique in G^c , and each $|K_i| = \log^3(n)$. Since B consists only of connected clusters of size at least $\log^4(n)$ in G^c , we can select these so that $m = (\rho - \epsilon_2 - o(1))n / \log^3(n)$.

Consider now the graph H , created by taking the K_i as vertices, as connecting K_i and K_j if a new edge is created between any two constituent vertices when the δn edges are added. Since for any vertices x and y , $d(x, y) \leq n/2$

$$\mathbf{P}(x \text{ and } y \text{ are connected by the new edges}) \geq \frac{\delta}{2n \log n}$$

It follows that the number of connections created between K_i and K_j dominates a random variable X which is $\text{Bin}(\log^6(n), \delta/2n \log n)$ distributed. From a simple second moment estimate, one gets

$$\mathbf{P}(X > 0) \geq \frac{\delta \log^6(n)}{4n \log n} = \left(\frac{\delta \log^2(n)}{4} \right) \frac{\log^3(n)}{n}.$$

It follows that H is dominated by a graph of the form

$$\mathcal{G} \left((1 - \epsilon_2 - o(1)) \frac{n}{\log^3(n)}, \left(\frac{\delta \log^2(n)}{4} \right) \frac{\log^3(n)}{n} \right)$$

of the standard Erdős-Rényi $\mathcal{G}(n, p)$ family. But the threshold for $\mathcal{G}(n, p)$ being completed connected a.a.s. is $p \gg \log n/n$, which holds here. Thus H is a.a.s. connected, from which it follows that B is a.a.s. connected in the completed graph. This establishes the result. \square

5 Conclusion

The α -model spans the spectrum from structure-free random graph models to ordinary percolation. When $\alpha \leq 1$, the connectivity results more or less mirror those of random graphs, whereas for greater values they behave more like percolation.

We note that the cases where $\alpha \leq 1$ are exactly those where $p_{u,v} \rightarrow 0$ as $n \rightarrow \infty$ for all $u \neq v$. An interesting question is to further explore this territory and see if this property, under some regularity (perhaps monotonicity) requirements, is sufficient for “random graph” type behavior, or if there are cases where this holds, but where the critical value is not one.

Acknowledgments

Thanks to my advisers Olle Häggström, and Devdatt Dubhashi, as well as Oliver Riordan for some early discussion about the subject, and Peter Jagers for some technical assistance. Also thanks to Svante Janson who pointed out an error in an earlier version of this manuscript.

References

- [1] M. Aizenman and C. M. Newman. Discontinuity of the percolation density in one dimensional $1/|x - y|^2$ percolation models. *Communications in Mathematical Physics*, 107:611–647, 1986.
- [2] I. Benjamini and N. Berger. The diameter of long-range percolation clusters on finite cycles. *Random Structures and Algorithms*, 19:102–111, 2001.
- [3] B. Bollobas. *Random Graphs*. Academic Press, 1985.
- [4] B. Bollobas, S. Janson, and Oliver Riordan. The phase transition in inhomogeneous random graphs. *Random Structure and Algorithms*, 31:3–122, 2007.
- [5] D. Coppersmith, D. Gamarik, and M. Sviridenko. The diameter of a one-dimensional long-range percolation graph. In *Proceedings of the thirteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 329 – 337, 2002.
- [6] P. Erdős and A. Renyi. On random graphs. *Publicationes Mathematicae*, 6:290–297, 1959.
- [7] S. Janson, T. Luczak, and A. Rucinski. *Random Graphs*. Wiley, 2000.
- [8] J. Kleinberg. The small-world phenomenon: an algorithmic perspective. In *Proceedings of the 32nd ACM Symposium on Theory of Computing*, 2000.
- [9] J. Kleinberg. Complex networks and decentralized search algorithms. In *Proceedings of the International Congress of Mathematicians (ICM)*, 2006.
- [10] C.M. Newman and L.S. Schulman. One dimensional $1/|j - i|^s$ percolation models: The existence of a transition for $s \leq 2$. *Communications in Mathematical Physics*, 104:547–571, 1986.