Inhomogeneous random graphs

Svante Janson
Uppsala University

Singapore
May, 2006

Background

One of the most studied random graphs is $G(n,p)$, which has n vertices that can be taken as the integers $1,\dots,n$, and where each pair of vertices is connected by an edge with probability p, independently of all other edges.

We consider the case $p = c/n$ for some constant $c > 0$, and let $n \to \infty$. The degree of a given vertex has a binomial distribution $\mathrm{Bi}(n-1,c/n) \approx \mathrm{Po}(c)$. This is a strongly concentrated distribution with an exponentially decreasing tail. Many graphs from "real life" has much larger tails, for example power-law tails, and it is therefore important to study also random graph models with such behaviour.

We will describe one class of random graphs that generalize $G(n, c/n)$ but also allow many less homogeneous examples, for example natural examples of 'scale-free' random graphs, where the degree distribution has a power-law tail. We believe that when it comes to modelling real-world graphs with, for example, observed power-laws for vertex degrees, our model provides an interesting and flexible alternative to existing models.

Nevertheless, we will see that many properties of $G(n, c/n)$ extend to these random graphs. In particular, we consider the question whether there exist a giant component or not, and we will, typically, find a phase transition similar to what happens for $G(n, c/n)$. There are, however, some interesting twists for some examples.

We are interested in graphs with a large number of vertices, and in particular in asymptotics as the number tends to infinity. The graphs we consider are such that the average degree stays bounded, so they are rather sparse.

Reference

These lectures are based on

B. Bollobás, S. Janson & O. Riordan (2006+), The phase transition in inhomogeneous random graphs. *Random Structures and Algorithms*, to appear.
`arXiv:math.PR/0504589`

See this paper for proofs and further details.

Some examples

We will study random graphs where the edges appear independently, as in $G(n,p)$, but where the probability of an edge may differ between different pairs of vertices. We begin with some examples of the type of random graph we consider.

Example. Let, as above, the vertices be the n integers $1,\dots,n$. For every pair $\{i,j\}$ with $i \neq j$, independently of all other pairs, connect i and j by an edge with probability

$$p_{ij} = \frac{\kappa(i/n, j/n)}{n}$$

for a given symmetric function κ on $(0,1]^2$.

Taking $\kappa = c$ constant, we obtain $G(n, c/n)$, but other functions κ give many other interesting random graphs.

A specific interesting example is given by the choice $\kappa(x,y) = c/\max\{x,y\}$; we then connect i and j by an edge with probability

$$p_{ij} = \frac{c}{\max\{i,j\}}.$$

Here $c > 0$ is a parameter that will be kept constant as n varies. We assume for simplicity that $c \le 2$; otherwise we have to define

$$p_{ij} = \min\left(\frac{c}{\max\{i,j\}}, 1\right).$$

This example is the *uniformly grown random graph*, or c/j*-graph*, $G_n^{1/j}(c)$. The graph $G_n^{1/j}(c)$ is thus the graph on $\{1,2,\ldots,n\}$ in which edges are present independently, and the probability that for $i \neq j$ the edge ij is present is $p_{ij} = c/\max\{i,j\}$, or simply c/j if $i < j$.

Remark

In principle, we may take any symmetric non-negative function κ on $(0,1]^2$ in this example. However, the function κ is evaluated only at rational points, so in order to make sense of having κ defined on the entire square, it is reasonable to impose a continuity condition.

To assume that κ is continuous (and thus bounded) on the closed unit square $[0,1]^2$ would be convenient, but too strong for our purposes since it excludes the example $\kappa(x,y) = c/\max\{x,y\}$ just given.

To assume that κ is continuous on the open unit square $(0,1)^2$ is enough, but it turns out that it suffices to assume that κ is *continuous almost everywhere* in the unit square.

Example. Let $M \geq 2$ be a fixed integer and divide $(0, 1]$ into the M intervals $I_k = ((k-1)/M, k/M]$, $k = 1, \ldots, M$. Let κ be constant on each square $I_k \times I_l$.

This means that the vertices are of M different types, and that the probability of an edge ij depends on the types of i and j. If n is a multiple of M, there are n/M vertices of each type.

Variations

We have obtained our edge probabilities p_{ij} by evaluating κ at the point $(i/n, j/n)$. An interesting alternative is to let $x_1, \dots, x_n$ be n *random* points in $(0, 1]$, independent and uniformly distributed, and then take $p_{ij} = \kappa(x_i, x_j)/n$. (We divide by n in order to keep the average degree bounded.)

If the vertex labels $1, \dots, n$ do not have any special significance, we may order the sequence $x_1, \dots, x_n$. The ordered sequence is then close to $\frac{1}{n}, \frac{2}{n}, \dots, 1$, so it is not surprising that we will obtain a random graph with the same asymptotic properties as in Example above.

As we will see in the general definition, we can also allow $x_1, \dots, x_n$ to be random and dependent; asymptotically, only the density of the points matter.

There is nothing special with the interval $(0, 1]$ here; it can be replaced by another space. For example, the finite-type case in Example is simpler described by using a finite type space $\{1, \ldots, M\}$.

This leads to the following general (but long and somewhat technical) definition.

Definition

The general inhomogeneous random graph $G^{\mathcal{V}}(n, \kappa)$ is defined as follows, with a slight extension of the version in Bollobás, Janson and Riordan (2006+). We proceed in two steps, constructing first the vertices and then the edges. Note that n is a parameter measuring the size of the graph, and we are primarily interested in asymptotics as $n \to \infty$.

In many instances, n is the number of vertices, but that is not always the case; in general, the number of vertices may be random, and we require only that it is roughly proportional to n. Moreover, there is in general no need for the parameter n to be integer valued.

A *vertex space* $\mathcal{V}$ is a triple $(\mathcal{S}, \mu, (\mathbf{x}_n)_{n\geq 1})$, such that the following holds.

(i) $\mathcal{S}$ is a separable metric space.

(ii) μ is a (positive) Borel measure on $\mathcal{S}$ with $0 < \mu(\mathcal{S}) < \infty$.

(iii) For each n, $\mathbf{x}_n$ is a random sequence $(x_1, x_2, \ldots, x_{N_n})$ of N_n points of $\mathcal{S}$ (where N_n may be deterministic or random).

(Formally, we should write $\mathbf{x}_n = (x_1^{(n)}, \ldots, x_{N_n}^{(n)})$, say, as we assume no relationship between the elements of $\mathbf{x}_n$ for different n, but we omit this extra index.)

Let $M(\mathcal{S})$ be the space of all (positive) finite Borel measures on $\mathcal{S}$, and equip $M(\mathcal{S})$ with the standard weak topology: $\nu_n \to \nu$ iff $\int f\, d\nu_n \to \int f\, d\nu$ for all bounded continuous functions $f : \mathcal{S} \to \mathbb{R}$. Let

$$\mu_n := \frac{1}{n} \sum_{i=1}^{N_n} \delta_{x_i}$$

where δ_x is the Dirac measure at $x \in \mathcal{S}$; thus μ_n is a random element of $M(\mathcal{S})$. We will further assume that

(iv) $\mu_n \xrightarrow{\mathsf{p}} \mu$, as elements of $M(\mathcal{S})$.

Recall that a set $A \subseteq \mathcal{S}$ is a *μ-continuity set* if A is (Borel) measurable and $\mu(\partial A) = 0$, where ∂A is the boundary of A. The convergence condition (iv) is equivalent to the condition that for every μ-continuity set A,

$$\mu_n(A) := \#\{i \le N_n : x_i \in A\}/n \xrightarrow{\mathsf{p}} \mu(A).$$

Let $\mathcal{V}$ be a vertex space, and let κ be a symmetric non-negative (Borel) measurable function on $\mathcal{S}\times\mathcal{S}$. (We call such a function a *kernel*.) We define the random graph $G^{\mathcal{V}}(n,\kappa)$ by first letting the vertex set be $\{1,\dots,N_n\}$.

(It is sometimes more convenient to identify the vertices with the points $x_1,\dots,x_{N_n}$ in $\mathcal{S}$ rather than integers, but note that this must be done with care if there are repetitions among the points x_i.)

We then add edges as follows. Given the sequence $\mathbf{x}_n$, we consider each pair of vertices $\{i,j\}$ with $i\neq j$ separately, and let there be an edge between i and j with probability

$$p_{ij}=\min\biggl\{\frac{\kappa(i/n,j/n)}{n},1\biggr\}.$$

This random choice is done independently for all pairs $\{i,j\}$, conditioned on $\mathbf{x}_n$.

In order to avoid pathologies, we finally assume

(v) κ is continuous a.e. on $\mathcal{S} \times \mathcal{S}$;

(vi) $\kappa \in L^1(\mathcal{S} \times \mathcal{S}, \mu \times \mu)$, i.e.,

$$\iint_{\mathcal{S}^2} \kappa(x,y)\, d\mu(x)\, d\mu(y) < \infty;$$

(vii)

$$\frac{1}{n}\, \mathbb{E}\, e\Big(G^{\mathcal{V}}(n,\kappa)\Big) \to \tfrac{1}{2} \iint_{\mathcal{S}^2} \kappa(x,y)\, d\mu(x)\, d\mu(y).$$

We say that the kernel κ is *graphical* on the vertex space $\mathcal{V} = (\mathcal{S}, \mu, (\mathbf{x}_n)_{n\geq 1})$ when these hold.

It can be shown that (vii) follows from the other assumptions if κ is bounded.

Note that the number of vertices $v(G^{\mathcal{V}}(n,\kappa))$ is roughly proportional to n; more precisely,

$$\frac{v(G^{\mathcal{V}}(n,\kappa))}{n} = \frac{N_n}{n} = \mu_n(\mathcal{S}) \xrightarrow{\mathrm{p}} \mu(\mathcal{S}).$$

Remark. It is assumed in Bollobás, Janson and Riordan (2006+) that μ is a probability measure, i.e. that $\mu(\mathcal{S}) = 1$. (This is equivalent to $N_n/n \xrightarrow{\mathrm{p}} 1$.) The extension here to $0 < \mu(\mathcal{S}) < \infty$ is convenient, but on a formal level it is only a matter of notation since we may normalize μ: for any $c > 0$, the random graph $G^{\mathcal{V}}(n,\kappa)$ is unchanged if we replace n by cn, κ by $c\kappa$ and μ by $c^{-1}\mu$, and we may choose $c = \mu(\mathcal{S})$. All results in Bollobás, Janson and Riordan (2006+) thus hold in our setting too, possibly with trivial modifications.

We often suppress the dependence on $\mathcal{V}$, writing $G(n,\kappa)$ for $G^{\mathcal{V}}(n,\kappa)$.

More generally, we may consider a sequence (κ_n) of kernels on $\mathcal{S}$ and the corresponding random graphs $G^{\mathcal{V}}(n, \kappa_n)$. We say that the sequence (κ_n) of kernels on $(\mathcal{S}, \mu)$ is *graphical on $\mathcal{V}$ with limit κ* if

(v) κ is continuous a.e. on $\mathcal{S} \times \mathcal{S}$;

(v') for a.e. $(y, z) \in \mathcal{S}^2$, $y_n \to y$ and $z_n \to z$ imply that $\kappa_n(y_n, z_n) \to \kappa(y, z)$;

(vi) $\kappa \in L^1(\mathcal{S} \times \mathcal{S}, \mu \times \mu)$;

(vii')

$$\frac{1}{n} \mathbb{E}\, e(G^{\mathcal{V}}(n, \kappa_n)) \to \tfrac{1}{2} \iint_{\mathcal{S}^2} \kappa(x, y)\, d\mu(x)\, d\mu(y).$$

We will see that most results for $G^{\mathcal{V}}(n, \kappa)$ hold for $G^{\mathcal{V}}(n, \kappa_n)$ too.

Further variations

A common variation of the construction above is that the edges are generated with probability

$$p_{ij} := 1 - \exp\Big(-\kappa(x_i, x_j)/n\Big),$$

rather than by $\Big\{\kappa(x_i, x_j)/n, 1\Big\}$. This is the result if we regard κ as intensities of Poisson processes of edges, and construct a multi1graph by adding a Poisson number of edges between i and j, with mean $\kappa(x_i, x_j)/n$, and then merge multiple edges. (There will be very few multiple edges, and for many purposes it does not matter whether we leave them or not.)

We can treat this version by regarding it as an instance of $G^{\mathcal{V}}(n, \kappa_n)$ with

$$\kappa_n(x, y) := n\Big(1 - \exp(-\kappa(x, y)/n)\Big).$$

It is easily seen that the sequence (κ_n) is graphical with limit κ if κ is graphical.

Another alternative, studied by Britton, Deijfen and Martin-Löf (2006+) in a special case, is to let

$$\frac{p_{ij}}{1-p_{ij}} = \frac{\kappa(x_i, x_j)}{n},$$

i.e., to take

$$p_{ij} := \frac{\kappa(x_i, x_j)}{n + \kappa(x_i, x_j)}.$$

Again, the sequence (κ_n) is graphical with limit κ if κ is graphical.

This version is sometimes simpler than our standard one. In particular, if the sequence $\mathbf{x}_n$ is deterministic, and the kernel κ has rank 1, i.e., $\kappa(x, y) = \psi(x)\psi(y)$ for some function ψ, then the probability of obtaining a specific graph G on the given vertex set equals

$$C \prod_i \psi(x_i)^{d_i}$$

where d_i is the degree of vertex i. The probability thus depends on the degree sequence only.

More examples

Example. The Erdős-Rényi random graph. If $\kappa = c$ is constant, then any choice of $\mathcal{S}$ and any choice of $x_1, \ldots, x_n$ gives the classical Erdős–Rényi random graph $G(n, c/n)$.

The simplest choice is to let $\mathcal{S}$ consist of a single point.

Example. The finite-type case.
Let $\mathcal{S} = \{s_1, \ldots, s_r\}$ be finite. Then κ is an $r \times r$ matrix. In this case, $G(n, \kappa)$ has vertices of r different types (or colours), say n_i vertices of type i, with two vertices of types i and j joined by an edge with probability $n^{-1}\kappa(i, j)$ (for $n \geq \max \kappa$). The condition (iv) means that $n_i/n \to \mu_i$ for each i (in probability if the n_i are random), where $\mu_i := \mu\{i\} \geq 0$.

This case has been studied by Söderberg who noted our main result on the phase transition in this case.

Example. I.i.d. vertices. For any $\mathcal{S}$ and μ, we can take $x_1, \ldots, x_n$ to be i.i.d. random points in $\mathcal{S}$ with distribution μ. (This has been proposed by Söderberg.)

Example. Poisson process graph. For any $\mathcal{S}$ and μ, and any $\lambda > 0$, let $x_1, \dots, x_N$ be the points of a Poisson process on $\mathcal{S}$ with intensity measure $\lambda\mu$. In other words, N has a Poisson distribution $\mathrm{Po}(\lambda)$, and, given N, the points are i.i.d. as in the preceding example. We consider the random graph $\widetilde{G}_\lambda(\kappa)$, where we use λ as a parameter instead of n.

Conditioned on $N = m$, this random graph is just $G(m, \tilde{\kappa})$, with $x_1, \dots, x_m$ as in the preceding example and $\tilde{\kappa} := (m/\lambda)\kappa$.

Results

We sometimes need one additional condition.

A kernel κ on a ground space $(\mathcal{S}, \mu)$ is *reducible* if $\exists A \subset \mathcal{S}$ with $0 < \mu(A) < \mu(\mathcal{S})$ such that $\kappa = 0$ a.e. on $A \times (\mathcal{S} \setminus A)$; otherwise κ is *irreducible*.

Thus κ is irreducible if $A \subseteq \mathcal{S}$ with $\kappa = 0$ a.e. on $A \times (\mathcal{S} \setminus A)$ implies that $\mu(A) = 0$ or $\mu(\mathcal{S})$.

Roughly speaking, κ is reducible if the vertex set of $G^{\mathcal{V}}(n, \kappa)$ can be split into two parts so that the probability of an edge from one part to the other is zero, and irreducible otherwise.

A branching process

A main tool to study components is a branching process approximation. We use the multi-type Galton–Watson branching process with type space $\mathcal{S}$, where a particle of type $x \in \mathcal{S}$ is replaced in the next generation by a set of particles distributed as a Poisson process on $\mathcal{S}$ with intensity $\kappa(x,y)\,d\mu(y)$. (Thus, the number of children with types in a subset $A \subseteq \mathcal{S}$ has a Poisson distribution with mean $\int_A \kappa(x,y)\,d\mu(y)$, and these numbers are independent for disjoint sets A and for different particles.) We denote this branching process, started with a single particle of type x, by $\mathfrak{X}_\kappa(x)$.

Let $\rho(\kappa;x)$ be the probability that the branching process survives for eternity.

We further define

$$\rho(\kappa) := \int_{\mathcal{S}} \rho(\kappa;x)\,d\mu(x).$$

An integral operator

Let T_κ be the integral operator on $(\mathcal{S}, \mu)$ with kernel κ, defined by

$$(T_\kappa f)(x) = \int_{\mathcal{S}} \kappa(x, y) f(y)\, d\mu(y),$$

for any (measurable) function f such that this integral is defined (finite or $+\infty$) for a.e. x. Note that $T_\kappa f$ is defined for every $f \geq 0$, with $0 \leq T_\kappa f \leq \infty$. If $\kappa \in L^1(\mathcal{S} \times \mathcal{S})$, as we assume, then $T_\kappa f$ is also defined for every bounded f; in this case $T_\kappa f \in L^1(\mathcal{S})$ and thus $T_\kappa f$ is finite a.e.

We define

$$\|T_\kappa\| := \sup\Big\{ \|T_\kappa f\|_2 : f \geq 0,\, \|f\|_2 \leq 1 \Big\} \leq \infty.$$

This is also the spectral radius.

Define the non-linear operator Φ_κ by

$$\Phi_\kappa f := 1 - e^{-T_\kappa f}$$

for $f \geq 0$. For such f we have $0 \leq T_\kappa f \leq \infty$, and thus $0 \leq \Phi_\kappa f \leq 1$. We can characterize $\rho(\kappa; x)$, and thus $\rho(\kappa)$, in terms of the non-linear operator Φ_κ.

There is a (necessarily unique) maximum solution $\tilde{\rho}_\kappa$ to

$$\Phi_\kappa(\tilde{\rho}_\kappa) = \tilde{\rho}_\kappa,$$

i.e., a solution that pointwise dominates all other solutions. Furthermore, $\rho(\kappa; x) = \tilde{\rho}_\kappa(x)$ for a.e. x, and

$$\Phi_\kappa(\rho_\kappa) = \rho_\kappa \quad a.e.,$$

where the function ρ_κ is defined by $\rho_\kappa(x) := \rho(\kappa; x)$.

If $\|T_\kappa\| \leq 1$, then $\tilde{\rho}_\kappa$ is identically zero, and this is thus the only solution. If $\|T_\kappa\| > 1$, then $\tilde{\rho}_\kappa$ is positive on a set of positive measure. Thus $\rho(\kappa) > 0$ if and only if $\|T_\kappa\| > 1$.

Main results

We denote the orders of the components of a graph G by $C_1(G) \geq C_2(G) \geq \ldots$, with $C_j(G) = 0$ if G has fewer than j components. We let $N_k(G)$ denote the total number of vertices in components of order k, and write $N_{\geq k}(G)$ for $\sum_{j \geq k} N_j(G)$, the number of vertices in components of order at least k. Our results are asymptotic, and all unspecified limits are taken as $n \to \infty$.

Theorem 1. *Let (κ_n) be a graphical sequence of kernels on a vertex space $\mathcal{V}$ with limit κ.*

(i) *If $\|T_\kappa\| \le 1$, then $C_1\big(G^{\mathcal{V}}(n,\kappa_n)\big) = o_p(n)$, while if $\|T_\kappa\| > 1$, then $C_1\big(G^{\mathcal{V}}(n,\kappa_n)\big) = \Theta(n)$* **whp**.

(ii) *For any $\varepsilon > 0$,* **whp** *we have*

$$\frac{1}{n} C_1\big(G^{\mathcal{V}}(n,\kappa_n)\big) \le \rho(\kappa) + \varepsilon.$$

(iii) *If κ is irreducible, then*

$$\frac{1}{n} C_1(G^{\mathcal{V}}(n,\kappa_n)) \xrightarrow{\mathrm{p}} \rho(\kappa).$$

In all cases $\rho(\kappa) < 1$; furthermore, $\rho(\kappa) > 0$ if and only if $\|T_\kappa\| > 1$.

As customary, we say that a sequence of random graphs G_n (with n vertices in G_n) has a *giant component* if $C_1(G_n) = \Theta(n)$ **whp**.

Corollary 1. *Let κ be a graphical kernel on a vertex space $\mathcal{V}$, and consider the random graphs $G^{\mathcal{V}}(n, c\kappa)$ where $c > 0$ is a constant. Then the threshold for the existence of a giant component is $c = \|T_\kappa\|^{-1}$. More precisely, if $c \le \|T_\kappa\|^{-1}$, then $C_1\Big(G^{\mathcal{V}}(n, c\kappa)\Big) = o_p(n)$, while if $c > \|T_\kappa\|^{-1}$ and κ is irreducible, then $C_1\Big(G^{\mathcal{V}}(n, c\kappa)\Big) = \rho(c\kappa)n + o_p(n) = \Theta_p(n)$.*

Corollary 2. *Let κ be a graphical kernel on a vertex space $\mathcal{V}$. Then the property that $G^{\mathcal{V}}(n, c\kappa)$ has* **whp** *a giant component holds for every $c > 0$ if and only if $\|T_\kappa\| = \infty$. Otherwise it has a finite threshold $c_0 > 0$.*

In the light of the results above, we say that a kernel κ is *subcritical* if $\|T_\kappa\| < 1$, *critical* if $\|T_\kappa\| = 1$, and *supercritical* if $\|T_\kappa\| > 1$. We use the same expressions for a random graph $G(n, \kappa)$ and a branching process $\mathfrak{X}_\kappa$.

The number of edges in the graph at the point where the giant component emerges is maximal in the classical Erdős–Rényi case. the proof.

Theorem 2. *Let κ_n be a graphical sequence of kernels on a vertex space $\mathcal{V}$ with limit κ, and assume that κ is critical, i.e. $\|T_\kappa\| = 1$. Then $\frac{1}{n}e(G^{\mathcal{V}}(n,\kappa_n)) \xrightarrow{\mathrm{p}} \frac{1}{2}\iint \kappa \leq 1/2$, with equality in the uniform case $\kappa = 1$.*

We can also determine the asymptotic number of edges in the giant component. As this is not always uniquely defined, for any graph G, let $\mathcal{C}_1(G)$ be the *largest component* of G, i.e., the component with most vertices, chosen according to any fixed rule if there is a tie. In order to state the next result concisely, let

$$\zeta(\kappa) := \frac{1}{2}\iint_{\mathcal{S}^2} \kappa(x,y)\Big(\rho(\kappa;x) + \rho(\kappa;y) - \rho(\kappa;x)\rho(\kappa;y)\Big)\, d\mu(x)\, d\mu(y).$$

Theorem 3. *Let (κ_n) be a graphical sequence of kernels on a vertex space $\mathcal{V}$ with irreducible limit κ. Then*

$$\frac{1}{n} e\Big(\mathcal{C}_1(G^{\mathcal{V}}(n,\kappa_n))\Big) \xrightarrow{\mathrm{p}} \zeta(\kappa).$$

The giant component is **whp** unique when it exists; the second largest component is much smaller. Indeed, only $o_p(n)$ vertices are in 'large' components other than the largest.

Theorem 4. *Let (κ_n) be a graphical sequence of kernels on a vertex space $\mathcal{V}$ with irreducible limit κ, and let $G_n = G^{\mathcal{V}}(n, \kappa_n)$. If $\omega(n) \to \infty$ and $\omega(n) = o(n)$, then*

$$\sum_{j \geq 2:\, C_j(G_n) \geq \omega(n)} C_j(G_n) = o_p(n).$$

In particular,

$$C_2(G_n) = o_p(n).$$

Remark. If κ and κ' are two kernels with $\kappa' \geq \kappa$, then one can couple the corresponding graphs or branching processes so that $G(n,\kappa) \subseteq G(n,\kappa')$ or $\mathfrak{X}_\kappa \subseteq \mathfrak{X}_{\kappa'}$. Thus $\rho(\kappa) \leq \rho(\kappa')$.

If κ is irreducible and $\rho(\kappa) > 0$, then $\rho(\kappa') > \rho(\kappa)$ unless $\kappa' = \kappa$ a.e.

Similarly, the threshold $c_0(\kappa') := \|T_{\kappa'}\|^{-1}$ is at most $c_0(\kappa) := \|T_\kappa\|^{-1}$.

Here, however, somewhat surprisingly, we may have $c_0(\kappa') = c_0(\kappa)$ even if $\kappa' > \kappa$. On the other hand, it is easily seen that if T_κ is compact and $\kappa' > \kappa$ on a set of positive measure, then $\|T_{\kappa'}\| > \|T_\kappa\|$ and thus $c_0(\kappa') < c_0(\kappa)$.

Stability

The giant component of $G_n = G^{\mathcal{V}}(n, \kappa_n)$ is stable in the sense that its size does not change much if we add or delete a few vertices or edges. Note that the vertices or edges added or deleted do not have to be random or independent of the existing graph; they can be chosen by an adversary after inspecting the whole of G_n.

Theorem 5. *Let (κ_n) be a graphical sequence of kernels on a vertex space $\mathcal{V}$ with irreducible limit κ, and let $G_n = G^{\mathcal{V}}(n, \kappa_n)$. For every $\varepsilon > 0$ there is a $\delta > 0$ (depending on κ) such that,* **whp**,

$$(\rho(\kappa) - \varepsilon)n \le C_1(G'_n) \le (\rho(\kappa) + \varepsilon)n$$

for every graph G'_n that may be obtained from G_n by deleting at most δn vertices and their incident edges, and then adding or deleting at most δn edges.

In particular, if G'_n is a graph on $[n] = 1, \ldots, n$ with $e(G'_n \triangle G_n) = o_p(n)$ then

$$C_1(G'_n) = C_1(G_n) + o_p(n) = \rho(\kappa)n + o_p(n).$$

As pointed out by Britton and Martin-Löf [**?**], the theorem has the following interpretation: suppose that G_n represents the network of contacts that may allow the spread of an infectious disease from person to person, and that we wish to eliminate the possibility of an epidemic by vaccinating some of the population. Even if the entire network of contacts is known, if the source of the infection is not known, a significant (constant, as $n \to \infty$) proportion of the population must be vaccinated: otherwise, there is still a giant component in the graph on the unvaccinated people, and if the infection starts at one of its vertices, it spreads to $\Theta(n)$ people.

More on the phase transition

Fix a graphical kernel κ on a vertex space $\mathcal{V}$, and study $G^{\mathcal{V}}(n, c\kappa)$ for a real parameter $c > 0$ as in the corollary above.

By Theorem 1, the size of the largest component of $G^{\mathcal{V}}(n, c\kappa)$ is described by the function $\rho(c\kappa)$, which is 0 for $c \le c_0 := \|T_\kappa\|^{-1}$ and strictly positive for larger c. With $\mathcal{V}$ and κ fixed, let us denote this function by $\rho(c)$, $c > 0$. It turns out that $\rho(c)$ is continuous on $(0, \infty)$.

Since $\rho(c) = 0$ for $c \le c_0$ but not for larger c, the function ρ is not analytic at c_0; in physical terminology, *there is a phase transition at* c_0.

For the classical Erdős–Rényi random graph $G(n, c/n)$ (obtained with $\kappa = 1$), it is well-known that ρ is continuous but the first derivative has a jump at $c_0 = 1$; more precisely, ρ'

jumps from 0 to $\rho'_+(c_0) = 2$. For finite d, we shall say that the phase transition in $G^{\mathcal{V}}(n, \kappa)$ has *exponent* d if $\rho(c_0 + \varepsilon) = \Theta(\varepsilon^d)$ as $\varepsilon \searrow 0$. As we have just noted, in $G(n, c/n)$ the phase transition has exponent 1. If $\rho(c_0 + \varepsilon) = o(\varepsilon^d)$ for all d, we say that the phase transition has *infinite exponent*. We are deliberately avoiding the physical term 'order', as it is not used in a consistent way in this context.

It has been shown (Dorogovtsev, Mendes and Samukhin (2001), Durrett (2003) and Bollobás, Janson and Riordan (2005)) that in the case $\mathcal{S} = (0, 1]$ and $\kappa(x, y) = 1/(x \vee y)$, the phase transition 'is of infinite order', i.e., has infinite exponent. We shall later see in that it is also possible to have a phase transition with any finite exponent larger than 1 (including non-integer values).

The next theorem shows that the phase transition has exponent 1 for a wide class of kernels κ, including all bounded κ.

Theorem 6. *Let κ be a kernel on a ground space $(\mathcal{S}, \mu)$. Suppose that κ is irreducible, and that*

$$\sup_x \int_{\mathcal{S}} \kappa(x,y)^2 \, d\mu(y) < \infty.$$

(i) *The function $c \mapsto \rho(c) := \rho(c\kappa)$ is analytic except at $c_0 := \|T_\kappa\|^{-1}$.*

(ii) *Furthermore, T_κ has an eigenfunction ψ of eigenvalue $\|T_\kappa\| < \infty$, and every such eigenfunction is bounded and satisfies*

$$\rho(c_0+\varepsilon) = 2c_0^{-1} \frac{\int_{\mathcal{S}} \psi \int_{\mathcal{S}} \psi^2}{\int_{\mathcal{S}} \psi^3} \varepsilon + O(\varepsilon^2), \qquad \varepsilon > 0,$$

so $\rho'_+(c_0) = 2c_0^{-1} \int_{\mathcal{S}} \psi \int_{\mathcal{S}} \psi^2 / \int_{\mathcal{S}} \psi^3 > 0$ and ρ has a phase transition at c_0 with exponent 1.

Corollary 3. *Let κ be an irreducible kernel such that (6) holds, and let $c_0 := \|T_\kappa\|^{-1} > 0$. Then $c_0 \rho'_+(c_0) \le 2$, with equality in the classical Erdős–Rényi case.*

Bounds on the small components

For the classical random graph $G(n, c/n)$ it is well-known that in the subcritical ($c < 1$) case, $C_1 = O(\log n)$ **whp**, and that in the supercritical ($c > 1$) case, $C_2 = O(\log n)$ **whp**. If we add some conditions, we obtain similar results. As before, we write G_n for $G^{\mathcal{V}}(n, \kappa_n)$.

Theorem 7. *Let (κ_n) be a graphical sequence of kernels on a vertex space $\mathcal{V}$ with limit κ.*

(i) *If κ is subcritical, i.e., $\|T_\kappa\| < 1$, and $\sup_{x,y,n} \kappa_n(x, y) < \infty$, then $C_1(G_n) = O(\log n)$* **whp**.

(ii) *If κ is supercritical, i.e., $\|T_\kappa\| > 1$, κ is irreducible, and either $\inf_{x,y,n} \kappa_n(x, y) > 0$ or $\sup_{x,y,n} \kappa_n(x, y) < \infty$, then $C_2(G_n) = O(\log n)$* **whp**.

Note that in part (ii) we draw the same conclusion from the very different assumptions $\inf_{x,y,n} \kappa_n(x,y) > 0$ and $\sup_{x,y,n} \kappa_n(x,y) < \infty$. There is no similar result for the subcritical case assuming only that $\inf_{x,y,n} \kappa_n(x,y) > 0$. Example:

The random graph $G_n^{1/j}(c)$ with $0 < c < 1/4$ is subcritical and satisfies $C_1(G_n^{1/j}(c)) = {}_{n^{\Theta(1)}}$ **whp**.

Degree sequence

The degree of a vertex of a given type x is asymptotically Poisson with a mean

$$\lambda(x) := \int_{\mathcal{S}} \kappa(x,y)\, d\mu(y)$$

that depends on x. This leads to a mixed Poisson distribution for the degree D of a random vertex of $G^{\mathcal{V}}(n,\kappa_n)$. We write Z_k for the number of vertices of $G^{\mathcal{V}}(n,\kappa_n)$ with degree k.

Theorem 8. *Let (κ_n) be a graphical sequence of kernels on a vertex space $\mathcal{V}$ with limit κ. Define $\lambda(x)$ as above, and let Ξ have the mixed Poisson distribution $\int_{\mathcal{S}} \mathrm{Po}(\lambda(x))\, d\mu(x)$. Then, for any fixed $k \geq 0$,*

$$Z_k/n \xrightarrow{\mathrm{p}} \mathbb{P}(\Xi = k) = \int_{\mathcal{S}} \frac{\lambda(x)^k}{k!} e^{-\lambda(x)}\, d\mu(x).$$

In other words, if D is the degree of a random vertex of $G^{\mathcal{V}}(n,\kappa_n)$, then

$$\mathcal{L}(D \mid G^{\mathcal{V}}(n,\kappa_n)) \xrightarrow{\mathrm{p}} \mathcal{L}(\Xi) = \int_{\mathcal{S}} \mathrm{Po}(\lambda(x))\, d\mu(x).$$

Distances between vertices

Let us write $d(v,w)$ for the graph distance between two vertices of $G_n = G^{\mathcal{V}}(n,\kappa_n)$, which we take to be infinite if they lie in different components.

Theorem 9. *Let κ_n be a graphical sequence of kernels on a vertex space $\mathcal{V}$ with limit κ, with $\|T_\kappa\| > 1$. Let $G_n = G^{\mathcal{V}}(n,\kappa_n)$, and let v and w be two independently chosen random vertices in G_n.*

(i) *If κ is irreducible and $1 < \|T_\kappa\| < \infty$, then*

$$\Big(d(v,w)/\log n \mid d(v,w) < \infty\Big) \xrightarrow{\mathsf{p}} 1/\log \|T_\kappa\|$$

(ii) *If κ is irreducible and $\|T_\kappa\| = \infty$, then there is a function $f(n) = o(\log n)$ such that*

$$\mathbb{P}\Big(d(v,w) \le f(n) \mid d(v,w) < \infty\Big) \to 1.$$

Define the 'diameter' of G_n as

$$\operatorname{diam}(G_n) := \max\{d(v,w) \,:\, v,w \in V(G),\, d(v,w) < \infty\},$$

the maximum of the diameters of the components of G_n.

Theorem 10. *Let κ be a kernel on a finite ground space $(\mathcal{S},\mu)$, $\mathcal{S} = \{1,2,\ldots,r\}$, with $\mu\{i\} > 0$ for each i. If $0 < \|T_\kappa\| < 1$, then*

$$\frac{\operatorname{diam}(G_n)}{\log n} \xrightarrow{\mathrm{p}} \frac{1}{\log \|T_\kappa\|^{-1}}$$

as $n \to \infty$, where $G_n = G^{\mathcal{V}}(n,\kappa)$. If $\|T_\kappa\| > 1$ and κ is irreducible, then

$$\frac{\operatorname{diam}(G_n)}{\log n} \xrightarrow{\mathrm{p}} \frac{2}{\log \|T_{\widehat{\kappa}}\|^{-1}} + \frac{1}{\log \|T_\kappa\|},$$

where $\widehat{\kappa}$ is the dual kernel to κ, defined by

$$\widehat{\kappa}(x,y) = (1 - \rho(\kappa))\kappa(x,y)$$

on $(\mathcal{S},\widehat{\mu})$, where the measure $\widehat{\mu}$ is defined by

$$d\widehat{\mu}(x) = (1 - \rho(\kappa;x))/(1 - \rho(\kappa))\, d\mu(x).$$

More examples

Example. The homogeneous case. More generally, let the ground space $(\mathcal{S}, \mu)$ be arbitrary, and let κ be irreducible and such that $\int_{\mathcal{S}} \kappa(x,y)\, d\mu(y)$ is independent of $x \in \mathcal{S}$, i.e., that

$$\int_{\mathcal{S}} \kappa(x,y)\, d\mu(y) = c \qquad \text{for every } x,$$

for some constant c. (This says roughly that, asymptotically, all vertices have the same average degree.)

Then $T_\kappa 1 = c$, so the constant function 1 is a positive eigenfunction with eigenvalue c, and thus $\|T_\kappa\| = c$, and by Theorem 1 there is a giant component (and $\rho(\kappa) > 0$) if and only if $c > 1$.

In the branching process, the number of children of each particle has a Po(c) distribution. Hence, ignoring the types of the particles, the distributions of the process $\mathfrak{X}_\kappa$ and the single-type process $\mathfrak{X}_c$ are the same. In particular, $\rho(\kappa) = \rho(c)$, so $\rho(\kappa) = \rho(c)$ is given by the same equation as in the Erdős–Rényi case.

Thus, the global behaviour of $G(n, \kappa)$ is exactly the same as that of $G(n, c/n)$, at least in terms of the size of the giant component. The local behaviour can be quite different, though. For example, $G(n, \kappa)$ may have many more triangles or other small cycles than $G(n, c/n)$. On the other hand, the vertex degrees have an asymptotic Po(c) distribution just as in $G(n, c/n)$.

A natural example of such a homogeneous κ is given by taking $\mathcal{S}$ as $(0, 1]$ (now better regarded as the circle $\mathbb{T}$), μ as Lebesgue measure, and $\kappa(x, y) = h(x - y)$ for an even function $h \geq 0$ of period 1. For example, h can be constant on a small interval $(-\delta, \delta)$ and vanish outside it; this gives a modification of $G(n, c/n)$ where only "short" edges are allowed.

More generally, $\mathcal{S}$ can be any compact homogeneous space, for example a sphere, with Haar measure μ and an invariant metric d, and $\kappa(x, y)$ a function of the distance $d(x, y)$.

Example. Take $\mathcal{S} = (0, 1]$ with μ the Lebesgue measure, and let $x_i = i/n$. Set $\kappa(x, y) = 1[x + y \leq 1]$ and consider the kernel $c\kappa$, so that

$$p_{ij} = \begin{cases} c/n, & i + j \leq n; \\ 0, & i + j > n. \end{cases}$$

Thus $G(n, c\kappa)$ can be obtained from the random graph $G(n, c/n)$ by deleting all edges ij with $i + j > n$.

The operator T_κ is compact, and it easy to see that it has eigenvalues $(-1)^k \omega_k^{-1}$ and eigenfunctions $\cos(\omega_k x)$, with $\omega_k = (k + 1/2)\pi$, $k = 0, 1, \ldots$. Hence $\|T_\kappa\| = 2/\pi$ and the critical value is $c_0 = \pi/2$. Theorem 6 shows that at the critical value we have $c_0 \rho'_+(c_0) = 3/2$.

Example. Edge percolation. Let κ be an irreducible graphical kernel on a vertex space $\mathcal{V}$ with $\|T_\kappa\| > 1$, and let $0 < p \le 1$. Independently of everything else, keep each edge in $G(n,\kappa)$ with probability p and delete it with probability $1-p$. Denote the resulting graph by $G^{\langle p\rangle}(n,\kappa)$.

This random graph $G^{\langle p\rangle}(n,\kappa)$ is nothing but $G(n,\tilde\kappa_n)$, where

$$\tilde\kappa_n(x,y) := p\Big(\kappa(x,y)\wedge n\Big).$$

Clearly, $x_n \to x$ and $y_n \to y$ imply $\tilde\kappa_n(x_n,y_n) \to p\kappa(x,y)$, provided (x,y) is a point of continuity of κ. Furthermore, $\mathbb{E}\, e\Big(G^{\langle p\rangle}(n,\kappa)\Big) = p\,\mathbb{E}\, e\Big(G(n,\kappa)\Big) \to p\frac{1}{2}\iint\kappa$. Hence, Theorem 1 applies with κ replaced by $p\kappa$, so

$$n^{-1}C_1\Big(G^{\langle p\rangle}(n,\kappa)\Big) \xrightarrow{\mathrm{p}} \rho(p\kappa).$$

In particular, $G^{\langle p\rangle}(n,\kappa)$ has **whp** a component of order $\Theta(n)$ if and only if $\|T_{p\kappa}\| > 1$, i.e., if $p > \|T_\kappa\|^{-1}$.

Example. Vertex percolation. Independently of everything else, keep each vertex in $G(n,\kappa)$ with probability p and delete it with probability $1-p$. Denote the resulting graph by $G^{[p]}(n,\kappa)$.

$G^{[p]}(n,\kappa)$ is the graph $G(m,\kappa_n)$ obtained from a random sample $\widetilde{x}_1,\dots,\widetilde{x}_m$ of the points $x_1,\dots,x_n$, rather than from all of them.

Theorem 1 applies (with μ replaced by $p\mu$), and it follows that

$$n^{-1}C_1\Big(G^{[p]}(n,\kappa)\Big) \xrightarrow{\mathsf{p}} p\rho(p\kappa).$$

In particular, $G^{[p]}(n,\kappa)$ has **whp** a component of order $\Theta(n)$ if and only if $\|T_{p\kappa}\|>1$, i.e. if $p>\|T_\kappa\|^{-1}$. We thus obtain the same threshold for vertex percolation in $G(n,\kappa)$ as for edge percolation.

Dubins' model

A common setting is the following: the vertex space $\mathcal{V}$ is $(\mathcal{S}, \mu, (\mathbf{x}_n)_{n\geq 1})$, where $\mathcal{S} = (0,1]$, μ is the Lebesgue measure, and $\mathbf{x}_n = (x_1, \ldots, x_n)$ with $x_i = i/n$. In this case, we have $p_{ij} = \kappa(i/n, j/n)/n \wedge 1$ for the probability of an edge between vertices i and j. We shall consider several choices of κ in some detail.

Observe first that if κ is a positive function on $(0,\infty)^2$ that is homogeneous of degree -1, then $p_{ij} = \kappa(i,j) \wedge 1$. Since this does not depend on n, in this case we can also consider the infinite graph $G(\infty, \kappa)$, defined in the same way as $G_n = G^{\mathcal{V}}(n,\kappa)$ but on the vertex set $\{1, 2, \ldots\}$. Note that the graphs $G^{\mathcal{V}}(n,\kappa)$ are induced subgraphs of $G(\infty,\kappa)$ and that we can construct them by successively adding new vertices, and for each new vertex an appropriate random set of edges to earlier vertices.

We first consider $\kappa(x, y) = c/(x \vee y)$ with $c > 0$, so that if $j \geq c$ then

$$p_{ij} = c/j \qquad \text{for } i < j.$$

In this case we can regard $G^{\mathcal{V}}(n, \kappa)$ as a sequence of graphs grown by adding new vertices one at a time where, when vertex k is added, it gets $\mathrm{Bi}(k-1, c/k)$ edges, whose other endpoints are chosen uniformly among the other vertices. (We might instead take $\mathrm{Po}(c) \wedge (k-1)$ new edges, without any difference in the asymptotic results below.)

This infinite graph $G(\infty, \kappa)$ was considered by Dubins in 1984, who asked when $G(\infty, \kappa)$ is a.s. connected. Dubins' question was answered partially by Kalikow and Weiss (1988). A little later Shepp (1989) proved that $G(\infty, \kappa)$ is a.s. connected if and only if $c > 1/4$. This result was generalized to more general homogeneous kernels by Durrett and Kesten (1990).

The finite random graph $G^{\mathcal{V}}(n, \kappa)$ with this κ, has been studied by Durrett (2003), who points out that it has the same critical value $c = 1/4$ for the emergence of a giant component as the infinite version has for connectedness, and by Bollobás, Janson and Riordan (2005) who rigorously show that this example has a phase transition with infinite exponent. More precisely, denoting $\rho(\kappa)$ by $\rho(c)$, it was shown by Riordan (2006+) that

$$\rho(1/4 + \varepsilon) = \exp\Big(-\tfrac{\pi}{2}\varepsilon^{-1/2} + O(\log \varepsilon)\Big).$$

A similar formula for the closely related CHKNS model introduced by Callaway, Hopcroft, Kleinberg, Newman and Strogatz (2001) had been given earlier by Dorogovtsev, Mendes and Samukhin (2001) using non-rigorous methods.

To find the critical value by our methods, we have to find the norm of T_κ on $L^2(0,1)$. Using the isometry $U : f \mapsto e^{-x/2} f(e^{-x})$ of $L^2(0,1)$ onto $L^2(0,\infty)$, we may instead consider $\widetilde{T}_\kappa := U T_\kappa U^{-1}$, which by a simple calculation is the integral operator on $L^2(0,\infty)$ with kernel

$$\begin{aligned}\tilde{\kappa}(x,y) &= e^{-x/2}\kappa(e^{-x},e^{-y})e^{-y/2} \\ &= ce^{-x/2-y/2+x\wedge y} = ce^{-|x-y|/2}.\end{aligned}$$

Hence $\widetilde{T}_\kappa$ is the restriction to $(0,\infty)$ of the convolution with $h(x) := ce^{-|x|/2}$. Because of translation invariance, it is easily seen that $\widetilde{T}_\kappa$ has the same norm as convolution with h on $L^2(-\infty,\infty)$, and taking the Fourier transform we find

$$\begin{aligned}\|T_\kappa\| = \|\widetilde{T}_\kappa\| &= \|f \mapsto h * f\|_{L^2(-\infty,\infty)} \\ &= \sup_{\xi\in\mathbb{R}} |\hat{h}(\xi)| = \int_{-\infty}^{\infty} h(x)\,dx = 4c.\end{aligned}$$

Thus, Theorem 1 shows that there is a giant component if and only if $c > 1/4$.

To find the size of the giant component is more challenging. It is easy to see that T_κ is a non-compact operator, and that it has no eigenfunctions at all in L^2. We suspect that this is connected to the fact that the phase transition has infinite exponent.

The mean-field scale-free model

Another interesting case with a homogeneous kernel as in Subsection is $\kappa(x,y) = c/\sqrt{xy}$ with $c > 0$; then (for $ij \geq c^2$)

$$p_{ij} = c/\sqrt{ij}.$$

This model has been studied in detail by Riordan (2006+). Considering the sequence $G^{\mathcal{V}}(n,\kappa)$ as a growing graph, in this case, together with each new vertex we add a number of edges that has approximately a Poisson Po($2c$) distribution; the other endpoint of each edge is chosen with probability proportional to $i^{-1/2}$, which is approximately proportional to the degree of vertex i. Hence, this random graph model resembles the growth with preferential attachment model of Barabási and Albert (1999), which was made precise as the *LCD model* by Bollobás and Riordan (2004).

In fact, up to a factor of $1 + o(i^{-1})$ in the edge probabilities, this model is the 'mean-field' version of the Barabási–Albert model, having the same individual edge probabilities, but with edges present independently.

In this case, T_κ is an unbounded operator, because $x^{-1/2} \notin L^2(0,1)$, and thus there is no threshold. In other words, $\rho(c) := \rho(\kappa) > 0$ for every $c > 0$.

As shown by Riordan (2006+), $\rho(c)$ grows very slowly at first in this case too; more precisely,

$$\rho(c) \sim 2e^{1-\gamma} \exp\bigl(-1/(2c)\bigr) \qquad \text{as } c \to 0,$$

where γ is Euler's constant. The result in Riordan (2006+), for the Barabási–Albert model is different, showing that in this model the dependence between edges is important.

Remark. Random graphs related to the ones defined here but with some dependence between edges can be obtained by adding at each new vertex a number of edges with some other distribution, for example Bi(m, p) for some fixed m and p. Such random graphs have been considered by several authors, and their results show that not only the expected numbers of edges added at each step are important, but also the variances; the edge dependencies shift the threshold.

The CHKNS model

The CHKNS model of Callaway, Hopcroft, Kleinberg, Newman and Strogatz (2001) grows from a single vertex; vertices are added one by one, and after each vertex is added, an edge is added with probability δ; the endpoints are chosen uniformly among all existing vertices. (Multiple edges are allowed; this does not matter for the asymptotics.)

Following Durrett (2003), we consider a modification (which is perhaps at least as natural): after adding each vertex, add a Poisson Po(δ) number of edges to the graph, again choosing the endpoints of these edges uniformly at random. Thus, when vertex k is added, each existing pair of vertices acquires $\mathrm{Po}\left(\delta/\binom{k}{2}\right)$ new edges, and these numbers are independent. When we have reached n vertices, the number of edges between vertices

i and j, with $1 \le i \le j \le n$, is thus Poisson with mean

$$e_{ij} := \sum_{k=j}^{n} \frac{\delta}{\binom{k}{2}} = 2\delta \sum_{k=j}^{n} \frac{1}{k(k-1)} = 2\delta\left(\frac{1}{j-1} - \frac{1}{n}\right),$$

and the probability that there is one or more edges between i and j is $p_{ij} := 1 - \exp(-e_{ij})$.

Hence, ignoring multiple edges, we have a graph G_n of our type, with $\mathcal{S} = (0, 1]$, μ Lebesgue measure, $x_i = i/n$ and

$$\kappa_n(x, y) := n\left(1 - \exp\left(-2\delta\left(\frac{1}{n(x \vee y) - 1} - \frac{1}{n}\right)\right)\right)$$
$$\to \kappa(x, y) := 2\delta\left(\frac{1}{x \vee y} - 1\right).$$

Theorem 1 shows that $C_1(G_n)/n \xrightarrow{\mathrm{p}} \rho(\kappa)$.

The original CHKNS model, $\widetilde{G}_n$, say, can be treated by a comparison argument. It follows that $C_1(\widetilde{G}_n)/n \xrightarrow{\mathrm{p}} \rho(\kappa)$ holds for the CHKNS model too.

In particular, the threshold for the CHKNS model, as well as for Durrett's modification, is given by $\|T_\kappa\| = 1$, or $2\delta = \|T\|^{-1}$, where T is the integral operator with kernel $1/(x \vee y) - 1$ on $L^2(0,1)$. This kernel is strictly smaller that the kernel $1/(x \vee y)$ considered above. However, changing variables as there, we see that T is equivalent to the operator on $L^2(0,\infty)$ with kernel $e^{-|x-y|/2} - e^{-(x+y)/2}$. Using translational invariance of the operator with kernel $e^{-|x-y|/2}$ considered above, it is easily seen that T has the same norm as this operator, namely 4.

Thus the thresholds for the CHKNS model and Durrett's modification are both given by $2\delta = 1/4$, i.e. $\delta = 1/8$, as was found by non-rigorous arguments by Callaway, Hopcroft, Kleinberg, Newman and Strogatz (2001) and Dorogovtsev, Mendes and Samukhin (2003), and first proved rigorously by Durrett (2003).

The rank 1 case

The *rank 1* case, is when the kernel κ has the form $\kappa(x, y) = \psi(x)\psi(y)$ for some function $\psi > 0$ on $\mathcal{S}$. We shall assume that $\int \psi \, d\mu < \infty$, but not necessarily that $\int \psi^2 \, d\mu < \infty$.

This is a special case of our general model that, while very restrictive, is also very natural, and includes or is closely related to many random graph models considered by other authors.

The function $\psi(x)$ can be interpreted as the "activity" of a vertex at x, with the probability of an edge between two vertices proportional to the product of their activities.

In the rank 1 case, $T_\kappa f = \Big(\int f\psi\Big)\psi$, and $\|T_\kappa\| = \|\psi\|_2^2 = \int \psi^2 \, d\mu \leq \infty$. Thus T_κ is bounded if and only if $\psi \in L^2$, in which case T_κ has rank 1, so it is compact, and ψ is the unique (up to multiplication by constants) eigenfunction with non-zero eigenvalue.

By Theorem 8, the distribution of vertex degrees is governed by the distribution of the function $\lambda(x) = (\int \psi \, d\mu)\psi(x)$ on $(\mathcal{S}, \mu)$. In particular, the degree sequence will (asymptotically) have a power-law tail if the distribution of $\lambda(x)$ has; for example, if $\mathcal{S} = (0, 1]$ with μ Lebesgue measure, and $\psi(x) = cx^{-1/p}$.

Another, perhaps more canonical, version is to take $\psi(x) = x$ on $\mathcal{S} = [0, \infty)$, with a suitable finite Borel measure μ. Note that every random graph considered in this example may be defined in this way, since we may map $\mathcal{S}$ to $[0, \infty)$ by $x \mapsto \psi(x)$. Alternatively, we may map by $x \mapsto \lambda(x)$ and have $\psi(x) = cx$ with $c > 0$ and $\lambda(x) = x$.

Random graphs of this type have been studied by several authors, e.g, Chung and Lu, Norros and Reittu, and Britton, Deijfen and Martin-Löf.

To study the phase transition in the rank 1 case, let us now consider the kernel $c\kappa(x,y) = c\psi(x)\psi(y)$, with $c > 0$ a parameter. We study the size of the giant component (if any) as a function of c, and let $\alpha(c) := c\int \psi\rho_{c\kappa}\,d\mu$, where, as before, $\rho_{c\kappa}(x) = \rho(c\kappa; x)$ is the survival probability of the branching process $\mathfrak{X}_{c\kappa}(x)$. Then $\rho_{c\kappa} \searrow 0$ a.e. as $c \searrow c_0$, and thus so, by dominated convergence,

$$\alpha(c)/c \searrow 0 \text{ as } c \searrow c_0.$$

We have $T_{c\kappa}\rho_{c\kappa} = cT_\kappa\rho_{c\kappa} = \alpha(c)\psi$, and thus

$$\rho_{c\kappa} = \Phi_{c\kappa}(\rho_{c\kappa}) = 1 - e^{-T_{c\kappa}\rho_{c\kappa}} = 1 - e^{-\alpha(c)\psi}.$$

Let

$$\beta(t) := \int_{\mathcal{S}} \Bigl(1 - e^{-t\psi(x)}\Bigr)\psi(x)\,d\mu(x), \qquad t \ge 0.$$

Then,

$$\alpha(c) = c\int_{\mathcal{S}} \rho_{c\kappa}\psi\,d\mu = c\beta\Bigl(\alpha(c)\Bigr).$$

so $c = \alpha(c)/\beta\Bigl(\alpha(c)\Bigr)$, i.e., α is the inverse function to $t \mapsto \gamma(t) := t/\beta(t)$.

Let us consider some concrete examples. Take $\mathcal{S} = (0, 1]$ with μ Lebesgue measure, and let $\psi(x) = x^{-1/p}$ where $1 < p \leq \infty$.

Case 1: $1 < p < 2$. In this case, $\|\psi\|_2 = \infty$, so $c_0 = 0$. Calculations yield

$$\rho(c) \sim C_3\alpha(c) = C_3\gamma^{-1}(c) \sim C_4 c^{1/(2-p)} \quad \text{as } c \to 0.$$

Note that this exponent $1/(2-p)$ may be any real number in $(1, \infty)$.

Case 2: $p = 2$. We still have $\|\psi\|_2 = \infty$ and thus $c_0 = 0$. We now find that

$$\rho(c) = e^{-(1+o(1))/2c} \qquad \text{as } c \to 0.$$

More refined estimates can be obtained in the same way.

Case 3: $2 < p < 3$. For $p > 2$ we have $\int \psi^2 \, d\mu < \infty$, and thus $c_0 > 0$, so we have a phase transition. (In fact, $c_0 = 1 - 2/p$.) Calculations yield

$$\rho(c_0+\varepsilon) \sim C_4 \alpha(c_0+\varepsilon) \sim C_5 \varepsilon^{1/(p-2)} \qquad \text{as } \varepsilon \searrow 0.$$

We thus have a phase transition at c_0 with exponent $1/(p-2)$. Note that this exponent may be any real number in $(1, \infty)$. (Taking instead e.g. $\psi(x) = x^{-1/2} \ln^{-1}(e^3/x)$, it is similarly seen that there is a phase transition with infinite exponent.)

Case 4: $p = 3$. Similarly, with $c_0 = 1/3$,

$$\rho(c_0+\varepsilon) \sim C \alpha(c_0+\varepsilon) \sim C_1 \varepsilon / \ln(1/\varepsilon) \qquad \text{as } \varepsilon \searrow 0,$$

so $\rho'(c_0) = 0$.

Case 5: $3 < p \leq \infty$. In this case, $\int \psi^3 \, d\mu < \infty$. We find

$$\rho(c_0 + \varepsilon) \sim C_2 \alpha(c_0 + \varepsilon) \sim C_3 \varepsilon,$$

so we have a phase transition with exponent 1. (This is similar to Theorem 6, although the conditions are not quite satisfied.)

A "quantum random graph"

Ioffe and Levit (2006+) have recently introduced and studied a new random graph model related to quantum theory. To construct the graph, start with n cirles of length $\beta > 0$. break the circles into pieces using independent Poisson processes of intensity λ. Then connect every pair of circles using a Poisson process with intensity $1/n$ of links.

It is easily seen that this is an instance of our model, where $\mathcal{S}$ is the family of all intervals in a circle of length β, μ is a certain measure with total mass $e^{-\lambda\beta} + \lambda\beta$, and

$$\kappa(I, J) = |I \cap J|.$$

A calculation shows that

$$\|T_\kappa\| = \frac{2}{\lambda}\Big(1 - e^{-\lambda\beta}\Big) - \beta e^{-\lambda\beta}.$$

See further Janson (2006+).

Turova's model

Turova (2003, 2004, 2005) has studied a dynamical random graph $G(t)$, $t \geq 0$, defined as follows, using three parameters $\gamma > 0$, $\lambda > 0$ and $\delta \geq 0$. The graph starts with a single vertex at time $t = 0$. Each existing vertex produces new, initially isolated, vertices according to a Poisson process with intensity γ. As soon as there are at least two vertices, each vertex sends out edges according to another Poisson process with intensity λ; the other endpoint is chosen uniformly among all other existing vertices. (Multiple edges are allowed, but this makes little difference.) Vertices live for ever, but edges die with intensity δ, i.e., the lifetime of an edge has an exponential distribution with mean $1/\delta$.

By homogeneity we may assume $\gamma = 1$; the general case follows by replacing λ and δ by λ/γ and δ/γ and changing the time scale.

The vertices proliferate according to a Yule process (binary fission process): writing $N(t)$ for the number of vertices at time t, the probability that a new vertex is added in the infinitesimal time interval $[t, t + dt]$ is $N(t)\,dt$. It is well-known that

$$e^{-t}N(t) \overset{\text{a.s.}}{\to} W \qquad \text{as } t \to \infty$$

for a random variable W with $W > 0$ a.s. (In fact, $W \sim \text{Exp}(1)$.)

We condition on the vertex process. We then take $\mathcal{S} = [0, \infty)$ and let $x_1, \ldots, x_{N(t)}$ be the ages of the particles existing at time t. It is easily checked that this gives a vertex space $(\mathcal{S}, \mu, \mathbf{x}_n)$, where μ is the measure on $[0, \infty)$ given by $d\mu/dx = e^{-x}$ (the exponential distribution).

Moreover, the number of edges at time t between two vertices of ages x_i and x_j has a Poisson distribution with mean $\kappa_t^*(x,y)/N(t)$, where

$$\kappa_t^*(x,y) := 2\lambda \int_0^{x\wedge y} e^{-\delta s} \frac{N(t)}{N(t-s)-1}\, ds.$$

It is easily checked that if $\delta \neq 1$ and $x_t \to x$, $y_t \to y$, then

$$\kappa_t^*(x_t,y_t) \to \kappa_\delta(x,y) = \frac{2\lambda}{1-\delta}\Big(e^{(1-\delta)(x\wedge y)} - 1\Big)$$

For $\delta = 1$, corresponding to $\delta = \gamma$ in the non-rescaled model, let $\kappa_1(x,y) := 2\lambda(x \wedge y)$. Then $\kappa_t^*(x_t,y_t) \to \kappa_\delta(x,y)$ in this case also.

Theorem 1 thus applies to $G(t)$ conditioned on the process $(N(t))_{t\geq 0}$.