Shortest-Weight Paths in Random Graphs

Hamed Amini

EPFL

Nice Random Graphs Workshop , May 2014

Randomized Broadcast

- The classical randomized broadcast model was first investigated by Frieze and Grimmett (1985).

Randomized Broadcast

- The classical randomized broadcast model was first investigated by Frieze and Grimmett (1985).
- Given a graph $G = (V, E)$, initially a piece of information is placed on one of the nodes in V. Then in each time step, every informed node sends the information to another node, chosen independently and uniformly at random among its neighbors.

Randomized Broadcast

- The classical randomized broadcast model was first investigated by Frieze and Grimmett (1985).
- Given a graph $G = (V, E)$, initially a piece of information is placed on one of the nodes in V. Then in each time step, every informed node sends the information to another node, chosen independently and uniformly at random among its neighbors.
- The question now is how many time-steps are needed such that all nodes become informed.

Randomized Broadcast

- The classical randomized broadcast model was first investigated by Frieze and Grimmett (1985).
- Given a graph $G = (V, E)$, initially a piece of information is placed on one of the nodes in V. Then in each time step, every informed node sends the information to another node, chosen independently and uniformly at random among its neighbors.
- The question now is how many time-steps are needed such that all nodes become informed.
- Fountoulakis and Panagiotou (2010) have recently shown that in the case of random r-regular graphs, the process completes in $\left(\frac{1}{\log(2(1-1/r))} - \frac{1}{r\log(1-1/r)}\right)\log n + o(\log n)$ rounds w.h.p.

Asynchronous Broadcasting

Each node has a Poisson clock with rate one. $\mathrm{ABT}(G)$ denotes the time it takes to inform the whole population.

Corollary

Let $G \sim \mathcal{G}(n,r)$ be a random r-regular graph with n vertices. We have w.h.p.

$$\mathrm{ABT}(G) = 2\left(\frac{r-1}{r-2}\right)\log n + o(\log n).$$

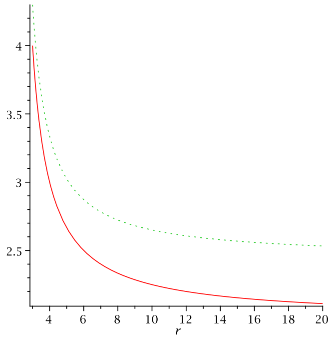

Figure: Comparison of the time to broadcast in the synchronized version (dashed) and in the case with exponential random weights (plain)

Configuration Model

- For $n \in \mathbb{N}$, let $(d_i)_1^n$ be a sequence of non-negative integers such that $\sum_{i=1}^n d_i$ is even.
- We define a random multigraph with given degree sequence $(d_i)_1^n$, denoted by $G^*(n,(d_i)_1^n)$:

Configuration Model

- For $n \in \mathbb{N}$, let $(d_i)_1^n$ be a sequence of non-negative integers such that $\sum_{i=1}^n d_i$ is even.
- We define a random multigraph with given degree sequence $(d_i)_1^n$, denoted by $G^*(n, (d_i)_1^n)$:
  - To each node i we associate d_i labeled half-edges.
  - All half-edges need to be paired to construct the graph, this is done by a uniform random matching.
  - When a half-edge of i is paired with a half-edge of j, we interpret this as an edge between i and j.

Configuration Model

- For $n \in \mathbb{N}$, let $(d_i)_1^n$ be a sequence of non-negative integers such that $\sum_{i=1}^n d_i$ is even.
- We define a random multigraph with given degree sequence $(d_i)_1^n$, denoted by $G^*(n, (d_i)_1^n)$:
  - To each node i we associate d_i labeled half-edges.
  - All half-edges need to be paired to construct the graph, this is done by a uniform random matching.
  - When a half-edge of i is paired with a half-edge of j, we interpret this as an edge between i and j.
- Conditional on the multigraph $G^*(n, (d_i)_1^n)$ being a simple graph, we obtain a uniformly distributed random graph with the given degree sequence, which we denote by $G(n, (d_i)_1^n)$.

Assumptions on the Degree Sequence

(i) $|\{i, d_i^{(n)} = r\}|/n \to p_r$ for every $r \geq 0$ as $n \to \infty$;

(ii) $\lambda := \sum_r r p_r \in (0, \infty)$;

(iii) $\sum_{i=1}^n d_i^2 = O(n)$.

Assumptions on the Degree Sequence

(i) $|\{i,\, d_i^{(n)} = r\}|/n \to p_r$ for every $r \geq 0$ as $n \to \infty$;

(ii) $\lambda := \sum_r r p_r \in (0, \infty)$;

(iii) $\sum_{i=1}^n d_i^2 = O(n)$.

(iii) ensures that $\liminf \mathbb{P}(G^*(n, (d_i)_1^n)$ is simple$) > 0$. Janson (2009)

Local Structure

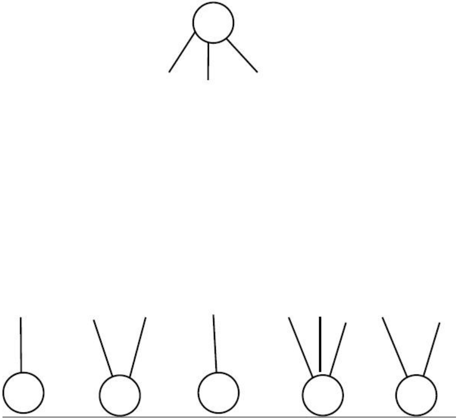

Local Structure

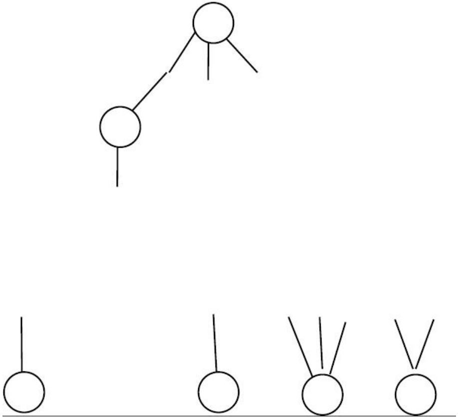

Local Structure

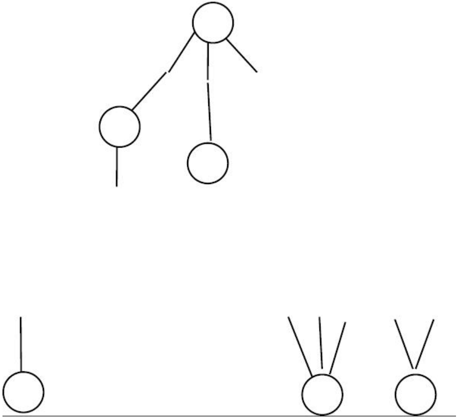

Local Structure

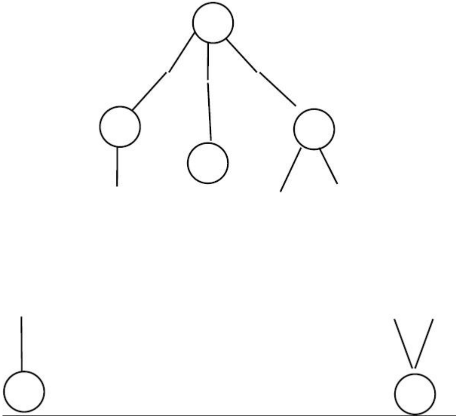

Branching Process Approximation

The first individual has offspring distribution $\{p_k\}$.
The other individuals have offspring distribution $\{q_k\}$:

$$q_k = \frac{(k+1)p_{k+1}}{\lambda}, \text{ and, } \nu = \sum_{k=0}^{\infty} k q_k \in (0,\infty).$$

The mean of the size of generation k is $\lambda\nu^{k-1}$.

Branching Process Approximation

The first individual has offspring distribution $\{p_k\}$.
The other individuals have offspring distribution $\{q_k\}$:

$$q_k = \frac{(k+1)p_{k+1}}{\lambda}, \text{ and, } \nu = \sum_{k=0}^{\infty} kq_k \in (0,\infty).$$

The mean of the size of generation k is $\lambda\nu^{k-1}$. The condition $\nu > 1$ is equivalent to the existence of a giant component, the size of which is proportional to n (Molloy, Reed 1998, and Janson 2009).

Typical Graph Distance

Theorem

For a and b chosen uniformly at random in the giant component, we have

$$\frac{\mathrm{dist}(a,b)}{\log n} \xrightarrow{p} \frac{1}{\log \nu}.$$

Van der Hofstad, Hooghiemstra, Van Mieghem 2005 for configuration model with i.i.d. degrees,
Bollobás, Janson, Riordan 2007 for inhomogeneous ramdom graphs.

Typical Weighted Distance

Theorem

For a and b chosen uniformly at random in $G(n,(d_i)_1^n)$ with $d_{\min} \geq 2$ and with i.i.d. exponential 1 weights on its edges, we have

$$\mathrm{dist}_w(a,b) - \frac{\log n}{\nu - 1} \xrightarrow{d} V.$$

Bhamidi, Van der Hofstad, Hooghiemstra 2009 for configuration model with i.i.d. degrees

Bhamidi, Van der Hofstad, Hooghiemstra 2010 for Erdős-Rényi random graphs

Typical Weighted Distance

Theorem

For a and b chosen uniformly at random in $G(n, (d_i)_1^n)$ with $d_{\min} \geq 2$ and with i.i.d. exponential 1 weights on its edges, we have

$$\mathrm{dist}_w(a, b) - \frac{\log n}{\nu - 1} \xrightarrow{d} V.$$

Bhamidi, Van der Hofstad, Hooghiemstra 2009 for configuration model with i.i.d. degrees

Bhamidi, Van der Hofstad, Hooghiemstra 2010 for Erdős-Rényi random graphs

Recall:

$$\frac{\mathrm{dist}(a, b)}{\log n} \xrightarrow{p} \frac{1}{\log \nu}.$$

Diameter

Generating function of $\{q_k\}_{k=0}^{\infty}$:

$$G_q(z) = \sum_{k=0}^{\infty} q_k z^k.$$

Let $\mathcal{X}_q$ be a Galton-Watson Tree (GWT) with offspring distribution q.
The extinction probability of the branching process, β, is the smallest solution of the fixed point equation

$$\beta = G_q(\beta).$$

Diameter

Define

$$\beta_* := G_q'(\beta) = \sum_{k=1}^{\infty} k q_k \beta^{k-1}.$$

Diameter

Define

$$\beta_* := G_q'(\beta) = \sum_{k=1}^{\infty} k q_k \beta^{k-1}.$$

Let $\mathcal{X}_q^+ \subseteq \mathcal{X}_q$ be the set of particles of $\mathcal{X}_q$ that survive and let D^+ denote the offspring distribution in $\mathcal{X}_q^+$.

We have

$$\mathbb{P}(D^+ = 1) = G_q'(\beta) = \beta_*.$$

Diameter

Define

$$\beta_* := G_q'(\beta) = \sum_{k=1}^{\infty} kq_k\beta^{k-1}.$$

Let $\mathcal{X}_q^+ \subseteq \mathcal{X}_q$ be the set of particles of $\mathcal{X}_q$ that survive and let D^+ denote the offspring distribution in $\mathcal{X}_q^+$.

We have

$$\mathbb{P}(D^+ = 1) = G_q'(\beta) = \beta_*.$$

The probability that the particles in generation k in $\mathcal{X}_q^+$, consists of a single particle, given that the whole process survives, is exactly β_*^k. This event corresponds to the branching process staying *thin* for k generations.

Diameter

$d_{\min} := \min\{\, k \mid p_k > 0 \,\}$ is such that for $k < d_{\min}$; $|\{i,\, d_i = k\}| = 0$, for all n sufficiently large.

Theorem

We have

$$\frac{\mathrm{diam}(G(n,(d_i)_1^n))}{\log n} \xrightarrow{p} \frac{1}{\log \nu} + \frac{\mathbf{1}(d_{\min}=2)}{-\log q_1} + 2\frac{\mathbf{1}(d_{\min}=1)}{-\log \beta_*}.$$

Bollobás, de la Vega 1982 for random regular graphs;
Fernholz, Ramachandran 2007 for configuration model;
Riordan, Wormald 2010 for Erdős-Rényi random graphs,
Bollobás, Janson, Riordan 2007 for inhomogeneous ramdom graphs.

WEIGHTED DIAMETER

Theorem (A., Lelarge)

Consider a random graph $G(n,(d_i)_1^n)$ with i.i.d. exponential 1 *weights on its edges, then*

$$\frac{\mathrm{diam}_w(G(n,(d_i)_1^n))}{\log n} \xrightarrow{p} \frac{1}{\nu-1} + \frac{2}{d_{\min}}\mathbf{1}_{(d_{\min}\geq 3)} + \frac{\mathbf{1}_{(d_{\min}=2)}}{1-q_1} + \frac{2}{1-\beta_*}\mathbf{1}_{(d_{\min}=1)}.$$

Ding, Kim, Lubetzky, Peres 2010 (random regular graphs)

WEIGHTED DIAMETER

Theorem (A., Lelarge)

Consider a random graph $G(n,(d_i)_1^n)$ with i.i.d. exponential 1 weights on its edges, then

$$\frac{\mathrm{diam}_w(G(n,(d_i)_1^n))}{\log n} \xrightarrow{p} \frac{1}{\nu-1} + \frac{2}{d_{\min}}\mathbf{1}_{(d_{\min}\geq 3)} + \frac{\mathbf{1}_{(d_{\min}=2)}}{1-q_1} + \frac{2}{1-\beta_*}\mathbf{1}_{(d_{\min}=1)}.$$

Ding, Kim, Lubetzky, Peres 2010 (random regular graphs)

Recall:

$$\frac{\mathrm{diam}(G(n,(d_i)_1^n))}{\log n} \xrightarrow{p} \frac{1}{\log \nu} + \frac{\mathbf{1}(d_{\min}=2)}{-\log q_1} + 2\frac{\mathbf{1}(d_{\min}=1)}{-\log \beta_*}.$$

Sketch of Proof

- The main idea of the proof consists in growing the balls around each vertex of the graph simultaneously so that the diameter becomes equal to twice the time when the last two balls intersect.
- Instead of taking a graph at random and then analyzing the balls, we use a standard coupling argument in random graph theory which allows to build the balls and the graph at the same time.
- There will be three different cases to consider depending on whether $d_{\min} \geq 3$, $d_{\min} = 2$, or $d_{\min} = 1$. Let

$$s_n := \left(\frac{1}{d_{\min}} \mathbf{1}_{(d_{\min} \geq 3)} + \frac{1}{2(1-q_1)} \mathbf{1}_{(d_{\min}=2)} + \frac{1}{1-\beta_*} \mathbf{1}_{(d_{\min}=1)} \right) \log n.$$

Sketch of Proof

- The proof of upper bound will consist in defining the two parameters α_n and β_n with the following significance:
  - (i) two balls of size at least β_n intersect almost surely,
  - (ii) the time it takes for the balls to go from size α_n to size β_n have all the same asymptotic for all the vertices of the graph, and the asymptotic is half of the typical weighted distance in the graph,
  - (iii) the time it takes for the growing balls centered at a given vertex to reach size at least α_n is upper bounded by $(1+\varepsilon)s_n$ for all $\varepsilon > 0$ w.h.p.

Sketch of Proof

- The proof of upper bound will consist in defining the two parameters α_n and β_n with the following significance:
  - (i) two balls of size at least β_n intersect almost surely,
  - (ii) the time it takes for the balls to go from size α_n to size β_n have all the same asymptotic for all the vertices of the graph, and the asymptotic is half of the typical weighted distance in the graph,
  - (iii) the time it takes for the growing balls centered at a given vertex to reach size at least α_n is upper bounded by $(1+\varepsilon)s_n$ for all $\varepsilon > 0$ w.h.p.
- To obtain the lower bound, we show that w.h.p.
  - (iv) there are at least two nodes with degree $d_{\min}$ such that the time it takes for the balls centered at these vertices to achieve size at least α_n is worst than the other vertices, and is lower bounded by $(1-\varepsilon)s_n$, for all $\varepsilon > 0$.

For a vertex $a \in V$ and a real number $t > 0$, the t-radius neighborhood of a is defined as

$$B_w(a,t) := \{\, b, \mathrm{dist}_w(a,b) \leq t \,\}.$$

The first time t where the ball $B_w(a,t)$ reaches size $k+1 \geq 1$ will be denoted by $T_a(k)$, i.e.,

$$T_a(k) := \min \{\, t : |B_w(a,t)| \geq k+1 \,\}, \qquad T_a(0) = 0.$$

We use I_a to denote the size of the component containing a in the graph minus one,

$$I_a := \max \{\, |B_w(a,t)|, t \geq 0 \,\} - 1.$$

so that for all $k > I_a$, we set $T_a(k) = \infty$.

Upper Bound

$$\alpha_n = \log^3 n, \text{ and } \beta_n = 3\sqrt{\frac{\lambda}{\nu - 1} n \log n}.$$

Upper Bound

$$\alpha_n = \log^3 n, \text{ and } \beta_n = 3\sqrt{\frac{\lambda}{\nu - 1} n \log n}.$$

Proposition

We have w.h.p.

$$\mathrm{dist}_w(u, v) \leq T_u(\beta_n) + T_v(\beta_n), \textit{ for all } u \textit{ and } v.$$

Upper Bound

$$\alpha_n = \log^3 n, \text{ and } \beta_n = 3\sqrt{\frac{\lambda}{\nu - 1} n \log n}.$$

Proposition

We have w.h.p.

$$\mathrm{dist}_w(u, v) \leq T_u(\beta_n) + T_v(\beta_n), \textit{ for all } u \textit{ and } v.$$

Proposition

For a uniformly chosen vertex u and any $\varepsilon > 0$, we have

$$\mathbb{P}\left(T_u(\beta_n) - T_u(\alpha_n) \geq \frac{(1+\varepsilon)\log n}{2(\nu - 1)} \mid I_u \geq \alpha_n \right) = o(n^{-1}).$$

Upper Bound: $d_{\min} \geq 3$

Lemma

We have $\mathbb{P}(I_a \geq \alpha_n) \geq 1 - o(n^{-3/2})$.

Upper Bound: $d_{\min} \geq 3$

Lemma

We have $\mathbb{P}(I_a \geq \alpha_n) \geq 1 - o(n^{-3/2})$.

Lemma

For a uniformly chosen vertex a, and any $\varepsilon, \ell > 0$, *we have*

$$\mathbb{P}\left(T_a(\alpha_n) \geq \varepsilon \log n + \ell \right) = o(n^{-1} + e^{-d_{\min}\ell}).$$

Upper Bound: $d_{\min} = 2$

Lemma

For a uniformly chosen vertex a, any $x > 0$*, and any* $\ell = O(\log n)$*, we have*

$$\mathbb{P}\left(T_a(\alpha_n \wedge l_a) \geq x \log n + \ell \right) \leq o(n^{-1}) + o(e^{-2(1-q_1)\ell}).$$

Upper Bound: $d_{\min} = 1$

- Let $\mathcal{C}_a$ the event that a is connected to the 2-core.
- The condition $\nu > 1$ ensures that the 2-core has size $\Omega(n)$, w.h.p.
- We consider the graph $\tilde{G}_n(a)$ obtained by removing all vertices of degree one except a until no such vertices exist.
- We consider two cases depending on whether both the vertices a and b are connected to the 2-core (i.e., the events $\mathcal{C}_a$ and $\mathcal{C}_b$ both hold), or both the vertices a and b belong to the same tree component of the graph.

Upper Bound: $d_{\min} = 1$

Lemma

$$\mathbb{P}\left(\tilde{T}_a(\alpha_n \wedge \tilde{I}_a) \geq x \log n + \ell\right) \leq o(n^{-1}) + o(e^{-(1-\beta_*)\ell}).$$

Upper Bound: $d_{\min} = 1$

Lemma

$$\mathbb{P}\left(\tilde{T}_a(\alpha_n \wedge \tilde{l}_a) \geq x \log n + \ell\right) \leq o(n^{-1}) + o(e^{-(1-\beta_*)\ell}).$$

Lemma

For two uniformly chosen vertices a, b, and any $\varepsilon > 0$, we have

$$\mathbb{P}\left(\frac{1+\varepsilon}{1-\beta_*} \log n < \mathrm{dist}_w(a,b) < \infty,\ \mathcal{C}_a^c,\ \mathcal{C}_b^c\right) = o(n^{-2}).$$

Lower Bound

Denote by Ω_a the ball centered at a containing exactly one node (possibly in addition to a) of degree at least 3.

Lower Bound

Denote by Ω_a the ball centered at a containing exactly one node (possibly in addition to a) of degree at least 3. For two nodes a, b, define the event $\mathcal{H}_{a,b}$ as

$$\mathcal{H}_{a,b} := \left\{ \frac{1-\varepsilon}{\nu-1} \log n < \mathrm{dist}_w(\Omega_a, \Omega_b) < \infty \right\}.$$

Proposition

If $u_1^{(n)} = o(n)$,

$$\mathbb{P}\left(\mathcal{H}_{a,b}\right) = 1 - o(1).$$

Lower Bound

(i) If the minimum degree $d_{\min} \geq 3$, then there are pairs of nodes a and b of degree $d_{\min}$ such that the event $\mathcal{H}_{a,b}$ holds and in addition all the weights on the edges adjacent to a or b are at least $(1-\varepsilon)\log n/d_{\min}$ w.h.p., for all $\varepsilon > 0$.

Lower Bound

(i) If the minimum degree $d_{\min} \geq 3$, then there are pairs of nodes a and b of degree $d_{\min}$ such that the event $\mathcal{H}_{a,b}$ holds and in addition all the weights on the edges adjacent to a or b are at least $(1-\varepsilon)\log n/d_{\min}$ w.h.p., for all $\varepsilon > 0$.

(ii) If the minimum degree $d_{\min} = 2$, then there are pairs of nodes a and b of degree two such that $\mathcal{H}_{a,b}$ holds and in addition, the closest nodes to each with forward-degree at least two is at distance at least $(1-\varepsilon)\log n/(2(1-q_1))$ w.h.p., for all $\varepsilon > 0$.

Lower Bound

(*i*) If the minimum degree $d_{\min} \geq 3$, then there are pairs of nodes a and b of degree $d_{\min}$ such that the event $\mathcal{H}_{a,b}$ holds and in addition all the weights on the edges adjacent to a or b are at least $(1-\varepsilon)\log n/d_{\min}$ w.h.p., for all $\varepsilon > 0$.

(*ii*) If the minimum degree $d_{\min} = 2$, then there are pairs of nodes a and b of degree two such that $\mathcal{H}_{a,b}$ holds and in addition, the closest nodes to each with forward-degree at least two is at distance at least $(1-\varepsilon)\log n/(2(1-q_1))$ w.h.p., for all $\varepsilon > 0$.

(*iii*) If the minimum degree $d_{\min} = 1$, then there are pairs of nodes of degree one such that $\mathcal{H}_{a,b}$ holds and in addition, the closest node to each which belongs to the 2-core is at least $(1-\varepsilon)\log n/(1-\beta_*)$ away w.h.p., for all $\varepsilon > 0$.

Hopcount Diameter

Let $G \sim \mathcal{G}(n,r)$ be a random r-regular graph with n vertices.
For $a, b \in V$, $\pi(a,b)$ denotes the minimum weight path between a and b. Let

$$f(\alpha) := \alpha \log\left(\frac{r-2}{r-1}\alpha\right) - \alpha + \frac{1}{r-2}.$$

Hopcount Diameter

Let $G \sim \mathcal{G}(n,r)$ be a random r-regular graph with n vertices.
For $a,b \in V$, $\pi(a,b)$ denotes the minimum weight path between a and b. Let

$$f(\alpha) := \alpha \log\left(\frac{r-2}{r-1}\alpha\right) - \alpha + \frac{1}{r-2}.$$

Theorem (A., Peres)

$$\frac{\max_{j\in[n]} |\pi(1,j)|}{\log n} \xrightarrow{p} \alpha^*, \ \textit{and} \ \frac{\max_{i,j\in[n]} |\pi(i,j)|}{\log n} \xrightarrow{p} \widehat{\alpha},$$

where α^* *and* $\widehat{\alpha}$ *are the unique solutions to* $f(\alpha) = 0$ *and* $f(\alpha) = 1$.

Hopcount Diameter

Let $G \sim \mathcal{G}(n,r)$ be a random r-regular graph with n vertices.
For $a, b \in V$, $\pi(a,b)$ denotes the minimum weight path between a and b. Let

$$f(\alpha) := \alpha \log\left(\frac{r-2}{r-1}\alpha\right) - \alpha + \frac{1}{r-2}.$$

Theorem (A., Peres)

$$\frac{\max_{j\in[n]} |\pi(1,j)|}{\log n} \xrightarrow{p} \alpha^*, \text{ and } \frac{\max_{i,j\in[n]} |\pi(i,j)|}{\log n} \xrightarrow{p} \widehat{\alpha},$$

where α^* *and* $\widehat{\alpha}$ *are the unique solutions to* $f(\alpha) = 0$ *and* $f(\alpha) = 1$.

Addario-Berry, Broutin and Lugosi 2010, Janson 1999: Complete Graph.
Bhamidi, van der Hofstad and Hooghiemstra 2009: $\frac{|\pi(1,2)| - \gamma \log n}{\sqrt{\gamma \log n}} \xrightarrow{d} Z$, where Z has a standard normal distribution and $\gamma = \frac{r-1}{r-2}$.
Open question: Hopcount diameter for configuration model?

THANK YOU!