Dynamics on Graphs

Rick Durrett

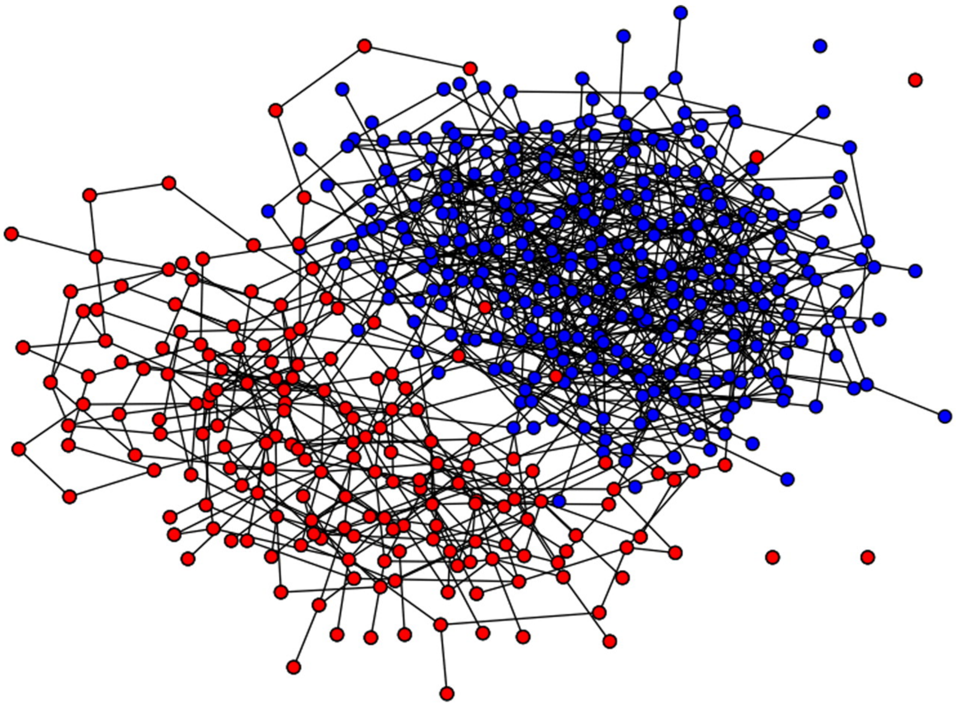

Graph fission in an evolving voter model, see section 8.1.1

https://link.springer.com/book/9783032219077

For Tom Liggett, and Harry Kesten
Mentors, colleagues, and friends, now departed

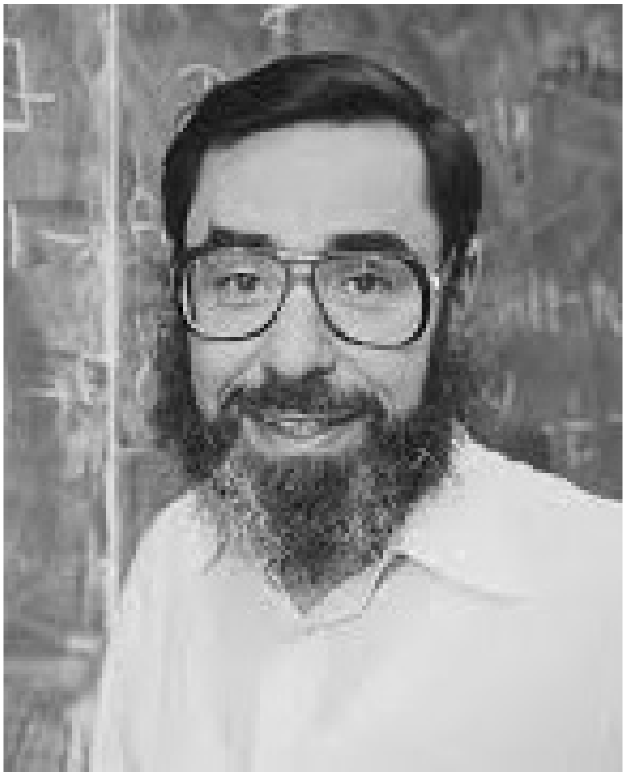

Memorial articles.

Harry Kesten (1931–2019): A Personal and Scientific Tribute. G.R. Grimmett and G.F. Lawler. *Notices of the A.M.S.* vol. 67 issue 6 (2020), pp. 822–831

The life and mathematical legacy of Thomas M. Liggett. Paul Jung, Amber Puha, et al. *Notices of the A.M.S.* vol. 68 issue 1 (2021), pp. 67–79

Dynamics on Graphs: A Quick Tour

1. Erdös-Rényi Random Graphs

In this chapter we will analyze the random graph model introduced by Erdös and Rényi in the 1950's. This example has been extensively studied. A very nice account of many of the results can be found in the classic book of Bollobás (2001), where the analysis has been done primarily by combinatorial methods. In contrast to this and many other treatments, we mainly rely on methods from probability.

To define the model, we begin with the set of vertices $V = \{1, 2, \ldots n\}$. For $1 \le x < y \le n$ let $\eta_{x,y}$ be independent, $= 1$ with probability p and $= 0$ otherwise, and then set $\eta_{y,x} = \eta_{x,y}$. If $\eta_{x,y} = 1$ there is an edge connecting x and y. Here, we will be primarily concerned with situation $p = \lambda/n$ and in particular with showing that there is a phase transition. Let $\mathcal{C}_x$ be the component containing x, i.e., the set of vertices that can be reached from x.

- When $\lambda < 1$ all of the components are small, i.e., there is a constant $a(\lambda)$ so that $P(\max_{1 \le x \le n} |\mathcal{C}_x| \ge a(\lambda) \log n) \to 0$ as $n \to \infty$
- When $\lambda > 1$ there is a **giant component** with $\sim g(\lambda)n$ vertices, and the second largest component is of size $O(\log n)$. There is also a central limit theorem for the size of the giant component.
- At the critical value, $\lambda = 1$, the largest component has size $n^{2/3}$, and there are many other components of this size
- As the reader will learn in Sections 1.7 and 1.8, the critical value is surrounded by a critical regime $\lambda = 1 + tn^{-1/3}$, $t \in (-\infty, \infty)$, in which cluster growth as a function of t can be approximated by a **multiplicative coalescent**, which gives an intuitive picture of the emergence of the giant component.

The first three results show that there is a **double jump transition**: the size of the largest component goes from $\log n$ to $n^{2/3}$ to n.

In Section 1.1 we define **branching processes** and describe results that will be useful in understanding the phase transition in the Erdős-Rényi random graph. To compute the size of $\mathcal{C}_x$ on a given graph, we set $I_0 = \{x\}$, let I_1 be the neighbors of x, and for $t \ge 2$ let I_t be the vertices y for which the distance $d(x, y) = t$.

In Section 1.2 we will show that when t is not too large $Z_t = |I_t|$, is approximately a branching process in which each individual has an average of λ children. If $\lambda < 1$ the branching process dies out exponentially fast and all components are small. When $\lambda > 1$, the branching process survives with probability $g(\lambda)$, and all sites with surviving branching processes combine to make the giant component.

In a branching process, a large number of individuals are born at one time, so in Section 1.3 we introduce the **exploration process** in which we reveal the neighbors of sites in the cluster one site at a time. In this process the number of active sites is a random walk run until it hits 0, so the large body of results about random walk leads to a detailed understanding of cluster sizes, much more precise than from the branching process viewpoint.

In Section 1.4 we compute the threshold for the graph to be connected, i.e., all vertices in one components. It is trivial that if there is a vertex with degree 0 then the graph is not connected. A result of Erdös and Rényi from 1959 states that if $p = (c_n + \log n)/n$), $H(n,p)$ is the probability that Erdős-Rényi(n,p) is connected, and $c_n \to c$ then

$$P(H(n,p)) \to \exp(-e^{-c}).$$

Most of the results in Chapter 1 are well-known, but Section 1.5 is an exception. In 1981 Ajtai, Kolmos, and Szemeredi showed that when $\lambda > 1$ not only is there a connected set of vertices of size cn but there is also a self-avoiding path with length $c'n$. This is useful for studying the contact process on Erdős-Rényi graphs because it is known that the survival time of the contact process on $[0,m]$ is $\geq \exp(cm)$ with high probability as $m \to \infty$.

Returning to well-known results, in Section 1.6 we prove a central limit theorem for the size of the giant component by using a clever argument of Martin-Löf based on the exploration process. In the first step we show that the fraction of unexplored vertices at time ns converges to the solution of the differential equation $du_s/ds = -\lambda u_s$, then use a little stochastic calculus to show that fluctuations around the limit when multiplied by $n^{1/2}$ converge to a Gaussian process.

In Section 1.7 we describe a combinatorial approach that leads to explicit formulas for the distribution of the cluster size in the critical regime. A useful consequence of this approach is a duality between the supercritical and subcritical regimes which makes it easy to prove that the second largest component has size $\sim a(\lambda)\log n$ when $\lambda > 1$.

In Section 1.8 we introduce the probabilistic approach to the critical regime. The most important results are two theorems of David Aldous that (i) gives an interesting description of the cluster sizes in terms of the excursions of a reflection Brownian with linear drift and (ii) show that on the time scale introduced in the fourth bullet point, the cluster sizes as a function of t are a multiplicative coalescent.

Finally in Section 1.9 we compute the critical exponents of percolation for the phase transition in the Erdős-Rényi random graph. These computations are relevant to percolation theory in $\mathbb{Z}^d$ since they are the mean-field values of the critical exponents, which are supposed to hold above ($>$) the critical dimension $d_c = 6$.

2. General Degree Distributions

The motivation for the study of graphs with a pre-specified degree distribution comes from the fact that in an Erdös-Rényi random graph, vertices have degrees that have asymptotically a Poisson distribution. However, in social and communication networks, the distribution of degrees is much different from the Poisson, and in many cases has a power law form, i.e., the fraction of vertices of degree k,

$$p_k \sim Ck^{-\gamma} \quad \text{as } k \to \infty.$$

Values of γ in the range (2.3) are common. These networks exhibit a "scale-free" property because their structure appears similar at different scales.

Some concrete example are: (i) the **world wide web**, i.e., the collection of web pages and the oriented links between them, $\gamma = 2.1$, (ii) the **Internet**, i.e., the physically connected network of routers that move email and files around, $\gamma = 2.3$, (iii) the **movie actor network** in which two actors are connected by an edge if they have appeared in a film together, $\gamma = 2.3$. (iv) The **collaboration graph** in a subject is a graph with an edge connecting two people if they have written a paper together, The fitted power laws had $\gamma = 2.4$ for math and $\gamma = 2.1$ for neuroscience, (v) a paper published in Nature in 2001 analyzed data gathered in a study of **sexual behavior of 4,781 Swedes**, and found that the number of partners per year had a power law of $\gamma = 3.3$ for men and $\gamma = 3.5$ for women. In all five examples the power law holds only for a range of values with exponential decay for very large values.

In Section 2.1 we will introduce the **configuration model**, which is a simple algorithm for creating a graph with a given degree distribution p_k . Generalizing our approach to Erdős-Rényi graphs, if we look at the number of vertices Z_n at distance n from x then we get a branching process, but this time the offspring distributions after the first one are different. Z_1 has distribution p_k, but the number of children in subsequent generations has the **size-biased distribution**,

$$q_{k-1} = kp_k/\mu \qquad \text{where } \mu = \textstyle\sum_k kp_k,$$

since as the cluster grows vertices of degree k are k times as likely as those of degree 1 to be chosen for connection. From this observation, we see that there is a giant component if (and only if) the mean of q_k, which we call $\nu > 1$. The fraction of vertices in the giant component can be computed almost as it was for the Erdős-Rényi but we need to use the generating function G_1 for q_k to compute the probability ρ_1 that a first generation individual starts a branching process that dies and then use the generating function for p_k to compute the probability that the **two-phase branching process dies out**.

In Section 2.2 we introduce the **limiting degree distribution approach**. It is more flexible and avoids some of the drawbacks of the configuration model, e.g., if we randomly delete edges from a configuration model graph, then the results is not one. Moilloy and Reed (1995) were the first to have the idea of proving results assuming only that the degrees distribution converged to a limit and some technical conditions held, for example to guarantee that the mean of the limiting degree distribution is the limit of the means. However, a proof under optimal conditions had to wait for Janson and Luczak (2009).

In subcritical Erdös-Rényi random graphs the largest component is $O(\log n)$, but this is no longer true for graphs with power law degree distributions $p_k \sim Ck^{-\gamma}$, for the trivial reason that the largest degree of a vertex is the graph is $O(n^{1/(\gamma-1)})$. In Section 2.3 we prove a result of Janson (2008), which shows that when the size-biased degree distribution has $P(d(x) \geq k) \leq Ck^{1-\gamma}$ the size of the largest component in a graph with n vertices $|\mathcal{C}^*| \leq An^{1/(\gamma-1)}$.

In Section 2.4 we will look at the distance between two randomly chosen vertices. When the degree distribution has finite variance the situation is much like the Erdős-Rényi case. Distances are of order $\log_\nu n$ with the fluctuations of order 1. When the distribution is a power law with $2 < \gamma < 3$, the world is very small: distances are

$$\sim (2 \log\log n) \log(1/(\gamma-2)) \qquad (\star)$$

Also in Section 2.4 we introduce the **Chung-Lu model** in which vertex i is assigned weight w_i and the probability of a connection from i to j is $(w_i w_j / \sum_k w_k) \wedge 1$. If the weights are chosen so that the degree distribution has a power law distribution, $p_k \sim Ck^{-\gamma}$ with $\gamma \in (2,3)$, then the distance between two randomly chosen vertices is given by $(\star)$. The proof is much easier for the Chung-Lu since edges are independent. As a bonus it gives some information about the structure of the graph: The set of vertices with degree $\geq n^{1/\sqrt{\log\log n}}$ form a connected set.

In Section 2.5 we tackle the subject of **first passage percolation** on random graphs, which could be a book all by itself. The topic occupies 71 pages in van der Hofstad's notes for the 2017 St. Flour Summer School. The four main papers we discuss have a total of 270 pages. Despite the fact that the graphs we consider are locally tree like, which limits the ability to optimize the passage time there are some surprises. When the degree distribution is a power law with $2 < \gamma < 3$ and passage times are exponential the shortest path between two randomly chosen vertices has length $O(\log\log n)$, but the shortest passage time between them which is $O(1)$ occurs on a path of length $\sim \alpha n$. Branching process fans will enjoy the fact that the analysis of passage times for general continuous distributions in Section 2.5.3 uses **Bellman-Harris processes** (also called age-dependent branching processes), **Malthusian parameters** and the **stable age distribution**. In some papers the very general *Crump-Mode-Jagers processes* are used but those will not play a significant role in this book.

Section 2.6 considers the analogue of Aldous' result for cluster sizes in the critical case discussed in Section 1.7. When the degree distribution has finite third moment the results are almost the same as in the Erdős-Rényi case, and can be formulated in terms of a limitingBrownian motion with drift. The analysis of the power-law case with $3 < \gamma < 4$ (which means that the size biased distribution has $2 < \gamma < 3$) leads to left continuous Lévy processes.

Finally Section 2.7 considers percolation on graphs. This is perhaps the simpliest process to consider on graphs since if we delete edges with probability $1-p$ we have another random graph for which we want the answer to the question: Is there a giant component? There are four cases to consider: finite variance degrees, and power laws with $3 < \gamma < 4$, $\gamma = 3$, and $2 < \gamma < 3$. When $\gamma = 3$ the critical value $p_c = 0$ and the percolation probability goes to 0 like $\exp(-(1 + o(1))/cp)$.

3. Inhomogeneous random graphs

In maximum generality an inhomogeneous random graph is one in which an edge bettween vertices i and j is present with probability $p(i,j)$. However, it is not possible to prove anything significant in this level of generality. Bollobás, Janson, and Riordan (2007) have found a framework that is very general and allows for a rich theory to be developed. The paper filled an entire issue of *Random Structures and Algorithms*. In the special case that we will consider, the vertices $\{1, 2, \ldots, n\}$ are embedded in $(0, 1]$ by the map $x \to x/n$. The interval is equipped with the Borel sets and Lebesgue measure, and there is a kernel κ defined on $(0, 1] \times (0, 1]$ (with properties to be specified later) so that

$$p(i,j) = \kappa(i/n, j/n)/n.$$

In Section 3.1 we will consider the situation in which there are only a finite number of different types of vertices. This is the simplest possible example, but it is important, because results for this case are a first step for proving things in general. These graphs can be called **multitype Erdős-Rényi graphs**, since the number of edges that connected type j vertices to type k has a Poisson distribution. Let $m_{j,k}$ be the mean number of vertices of type k that are neighbors of a vertex of type j. Generalizing results from the single type case, there is a giant component when the largest eigenvalue of the matrix $m_{j,k}$ is larger than 1, and the fraction of neighbors of type j in the giant component is given by the positive fixed point of the (multivariate) generating function of the offspring distribution. Finally, in the subcritical phase the largest cluser is $O(\log n)$.

In Section 3.2 we introduce a number of examples to show that the BJR framework is useful. Three examples that will be our constant companions as we develop the theory.

- **Dubins' model** has several identities. It can be a graph with connection probabilities $p(i,j) = c/(i \vee j)$ on $1, \ldots n$ or on the positive integers, $\mathbb{N}$. In the latter setting all vertices in the random graph have infinite degree, so work focuses on the question: "is the graph connected?" Dubins' model is a close relative of the **uniformly grown random graph**. On the nth step of the construction a new vertex labeled n is added, and with probability δ a new edge is added between two vertices chosen at random from the graph (so it is very unlikely that n is one of the endpoints). From this description it is not surprising that lower numbered vertices tend to have higher degrees, but the simple formula is a surprise (to me at least).

- **The square root model** is an inhomogeneous graph on $\{1, \ldots n\}$ with $p(i,j) = c/\sqrt{ij}$. It is called a mean-field version of the **preferential attachmen** (PA) model because in PA the probability of an edge connecting i and j is asumptotically $p(i,j)$. The "mean-field" in the name comes from the fact that in the PA model the presence of edges are not independent, but they are in the square-root model.

- **The Chung-Lu model** is defined by assigning a weight w_k to vertex k sand declaring that edges connecting i and j have with probability $p(i,j) = (w_i w_j / \sum_k w_k) \wedge 1$, and edges are independent. We will be most interested in the version that has a power-law degree distribution $p_k \sim Ck^{-\gamma}$. In this case $\kappa(x,y) = c\psi(x)\psi(y)$ with $\psi(x) = x^{-1/(\gamma-1}$. The square root model corresponds to $\gamma = 3$.

In Section 3.3 we will describe the set-up of the BJR model and state their main results. The key to the analysis is a linear integral operator, which in our special case is

$$(T_\kappa f)(x) = \int_0^1 \kappa(x,y) f(y)\, dy$$

with norm $\|T_\kappa\|_2 = \sup\{\|T_\kappa f\|_2 : f \geq 0, \|f\|_2 \leq 1\}$. The system is supercritcal if $\|T_\kappa\|_2 > 1$ and subcritical is $\|T_\kappa\|_2 < 1$. The survival probability is the maximum fixed point of the operator

$$\Phi_\kappa(f) = 1 - \exp(-T_\kappa f)$$

This is a generalization of the answer in the one and finite type cases but now we have to compute a fixed point that is a functionfunction.

In Section 3.4 we will introduce some theory that will allow us to compute in Section 3.5 the survival probabilities σ for our three examples. The answers are exotic in the first two cases

$$\sigma(\delta) \sim c\exp(-\pi(\delta - 1/8)^{1/2}) \qquad \text{Dubins' model}$$
$$\sigma(c) \sim 2e^{1-\gamma}\exp(-1/(2c)) \qquad \text{square root model}$$

where this time $\gamma = 0.57721566\ldots$ is Euler's constant. In the Chung-Lu model we have answers that match the percolation probabilities for the power law configuration model graphs discussed in Section 2.7.

Finally in Section 2.8, we give results for cluster sizes in the subcritical regime. The results for the Chung-Lu model are again similar to the ones for the power law case in Section 2.3.

4. Epidemics

In this chapter we will consider four epidemic models whose names are abbreviated as SI, SIR, SIS, and SIRS. At any time each vertex v is in state S = susceptible, I = infected, or R = removed (is immune to infection). Common to all four dynamics is the infection step: $S - I$ edges become $I - I$ at rate λ. Once an individual becomes infected they stay infected for a random amount of time T. These times for the various vertices are independent. We will primarily be concerned with the following situations:

- $T = \infty$ is the **SI mode**l, a depressing scenario in which individuals never recover, so we prefer to think of it as the spread of a rumor or Internet meme. This type of spread can also be modeled by first passage percolation, which is considered in Section 2.5.

- The simplest system with $T < \infty$ is the **SIR model**. If T is a fixed constant we have a simple process with connections to percolation, which is considered in Section 2.7. If T is exponential with rate γ, i.e., $P(T > t) = e^{-\gamma t}$, then the system has the Markov property.

- In the SIR system, when the infected period is over the individual enters the removed state. In the **SIS model**, which is a close relative of the contact process considered in Chapter 5, when the infected period is over the individual is again susceptible to the disease.

- Finally there is the **SIRS model**, discussed in Section 4.7.3, in which immunity in the removed state is only temporary. This is the case for the flu since the dominant strain of influenza changes from year to year.

We begin in Section 4.1 with the classical theory of epidemics in a homogeneously mixing population (or on a complete graph). We assume that the infection times are exponential so that models are Markov chains and lead in the limit to ODEs. While it is not possible to explicitly solve the differential equations, one can extract information such as the probability of a large epidemic and the fraction of individuals infected in one (which are the same in the SIR model).

In Section 4.2 we investigate epidemics in which the infection time is constant. There is a close relationship between the network of infections that occur in a homogeneously mixing population (complete graph) and the edges of an Erdős-Rényi random graph. As in the branching process, the probability of a large epidemic is the fixed point of a generating function.

In Section 4.3 we take the important step of allowing the infection times to have general distributions. If we let $I_{x,y}$ be the event that x infects neighbor y then $I_{x,y}$ and $I_{x,z}$ are independent when the infection time is constant. This breaks down in general, but since attempted infections from different sites are independent it is still possible to use branching processes to compute the probability of a large epidemic. That probability is no longer the same as the fraction of sites infected in a large epidemic. The good news is that the two probabilities are still fixed points of generating functions of branching processes.

In Section 4.4 we discuss the Miller-Volz equations, a set of differential equations that describe the spread of an epidemic on a graph generated by the configuration model. When the original paper of Volz (2008) was published, I believe that these equations were regarded as heuristics that produced good approximations. Indeed, the authors take pains to say that "calculations closely match simulations." Somewaht surprisingly, as we explain in Section 4.5, the equations turn out to be exact in the large graph limit. Several proofs of thiss have been given. We follow the approach of Janson, Luczak, and Windridge (2015) which proves the result under minimal assumptions.

In Section 4.6 we describe an epidemic model called the household model, due to Ball, Mollison, and Scalia-Tomba (1997) that incorporates the important feature that humans are organized into families. Their model maintains simplicity by assuming that the relationship between the houses is that of the complete graph. In the second half of the section, we describe an extension of their work that was used in the summer of 2020 to explore a pressing problem that was the focus of an REU project at Dule: How should the university prepare to welcome students back to campus in the Fall in order to minimize the impact of the covid epidemic? To study this question we investigated the **dorm model** which is a two-level version of the household model. Dorms have rooms, and rooms can have one or two students. Our analysis showed that epidemic spread was reduced if all students had single rooms. Duke did implement this idea for the 2020-2021 academic year, but we cannot claim credit for Duke making this decision.

Our final topic, pursued in Section 4.7, is epidemics on $\mathbb{Z}^2$. The main result is: if we start with one infected individual and the epidemic does not die out, then the epidemic grows linearly and has an asymptotic shape. The linear speed is seen in epidemics such as the spread of fox rabies in Europe and raccoon rabies in the US

At the end of the section we consider the SIRS epidemic. A result of Durrett and Neuhauser (1991) shows that if the SIR epidemic is supercritical and the $R \to S$ transition occurs at a positive rate then there is an equilibrium in which the epidemic persists. If we think of the SIRS model as a forest fire with I being burning trees, and R being burnt trees, then it is natural to assume that healthy trees (state s) replace burnt ones at a rate proporitonal to the number of neighbors in state S but for that version the Durrett-Neuhauser theorem is an open problem.

5. Contact Process

Harris (1974) introduced the contact process on $\mathbb{Z}^d$. Generalizing his definition to a graph G, the state at time t is a set $\xi_t \subset G$ of sites occupied by particles. Particles die at rate 1. Each occupied site x gives birth onto each neighbor y at rate λ. If y is vacant then it becomes occupied. If y is occupied then no change occurs. Let ξ_t^x be the contact process starting from one particle at x and let

$$\lambda_c = \inf\{\lambda : P(\xi^x \neq \emptyset \text{ for all } t) > 0\}$$

This value is independent of x on connected graphs. Interest focuses on estimating the value of λ_c and understanding the qualitative behavior of the model when $\lambda < \lambda_c$ (subcritical phase) and $\lambda > \lambda_c$ (supercritical). One is also interested in the behavior of critical systems ($\lambda = \lambda_c$) but that is a very difficult problem.

In Section 5.1, we construct the contact process from a **graphical representation** and introduce basic concepts from the theory of interacting particle systems that will be useful for its study, such as attractiveness and additivity. The most important fact is the contact process is **self-dual**. if we let ξ_t^A be the state at time t when we start from sites in A occupied at time t then

$$P(\xi_t^A \cap B \neq \emptyset) = P(A \cap \xi_T^B \neq \emptyset).$$

In Section 5.2, we calculate the predictions of **mean-field theory**, which is a typical first step for physicists (and probabilists) in determing the properties of a stochastic spatial model. These calculation can be done in two ways: (i) compute the evolution under the assumption that the states of sites are always independent, or (ii) consider the system on the complete graph. Using either viewpoint, we conclude that then the process dies out exponentially fast when $\lambda < \lambda_c$, and converges to an equilibrium with positive density when $\lambda > \lambda_c$.

Viewpoint (ii) enables us to conclude that if we start the contact process on the complete graph from all n sites occupied and $\lambda > \lambda_c$ then the system survives for time $\exp(c(\lambda)n)$. that result is the best possible for graphs with n vertices and $\leq Kn$ edges, since in that case the probability the system goes from all sites occupied to all vacant in time 1 is $\geq e^{-\alpha(\lambda)n}$

In Section 5.3, we consider the contact process on bounded degree graphs such as the random regular graph. From our study of the geometry of random graphs we know that this graph is locally tree like. So to prepare for the study of these examples we recall that Pemantle (1992) studied the contact process on the tree $\mathbb{T}^d$ in which each vertex has degree $d+1$ and found that in this situation if $d \geq 2$ (i.e., degree ≥ 3) then the contact process has two critical values:

$$\lambda_1 = \inf\{\lambda : P_0(\xi_t \neq \emptyset \text{ for all } t) > 0\}$$
$$\lambda_2 = \inf\{\lambda : \liminf_{t\to\infty} P_0(0 \in \xi_t) > 0\}$$

For many graphs with bounded degree it can be proved that if there are n vertices and we start from all sites occupied then there are constants that depend on λ so that

the process survives for a time $\leq c \log n$ when $\lambda < \lambda_1$

the process survives for time $\geq \exp(cn)$ when $\lambda > \lambda_1$

Readers with some experience with the contact process on trees may wonder why the critical value here is λ_1 instead of λ_2. The reason is that when $\lambda_1 < \lambda < \lambda_2$ particles on the infinite tree wander off to infinity while on finite graphs they go far away but eventually "wrap around" and come back to where they started.

In Section 5.4 we investigate the contact process on an Erdős-Rényi graph. The Poisson degree distribution is unbounded. but has a very thin tail. Using the existence of long paths proved in Section 1.5 and comparing with the contact processs on an interval we can prove survival for time $\exp(cn)$ for $\lambda > \lambda_c(\mathbb{Z})$. Also in Section 5.4 we describe a very clever proof of Cator and Don (2021) that lower bounds the survival time of the contact process on an Erdős-Rényi graph by comparing the number of occupied sites with a birth and death chain.

In Section 5.5 we begin the study of the contact process on graphs with a power law degree distribution, $p_k \sim Ck^{-\alpha}$. Using **degree based mean-field theory** physicists calculated that if $\rho(\lambda)$ is the density of occupied sites in equilibrium,

- If $\alpha \le 3$ then $\lambda_c = 0$.
- If $3 < \alpha < 4$, then $\lambda_c > 0$ then $\lambda_c > 0$ and $\rho(\lambda) \sim C(\lambda - \lambda_c)^\beta$ with $\beta > 1$.
- If $\alpha > 4$ then $\lambda_c > 0$ and $\rho(\lambda) \sim C(\lambda - \lambda_c)$ so the critical exponent $\beta = 1$.

In contrast, Chatterjee and Durrett (2009) proved

Theorem 5.5.1. *Consider a graph G_n with n vertices generated by the configuration model with $P(d_i = k) \sim Ck^{-\alpha}$ where $\alpha > 3$ and $P(d_i \le 2) = 0$. Let ξ_t^1, $t \ge 0$ denote the contact process on G_n starting from all sites occupied. Then for any $\lambda > 0$ there is a positive constant $p(\lambda) > 0$ so that for any $\delta > 0$*

$$\inf_{t \le \exp(n^{1-\delta})} P\left(|\xi_t^1|/n \le p(\lambda)\right) \to 1 \qquad \text{as } n \to \infty.$$

Mountford, Valesin and Yao (2013) extended the results to cover the powers $2 < \alpha \le 3$, and computed the correct values of the critical exponent β.

$$\rho(\lambda) \sim \begin{cases} \lambda^{1/(3-\alpha)} & 2 < a \le 5/2, \\ \lambda^{2\alpha-3} \log^{2-\alpha}(1/\lambda) & 5/2 < \alpha \le 3, \\ \lambda^{2\alpha-3} \log^{4-2\alpha}(1/\lambda) & 3 < \alpha. \end{cases} \tag{0.0.1}$$

Later Mountford, Mourrat, Valesin and Yao (2016) improved the survival time estimate to $\ge e^{cn}$.

In Section 5.6 and 5.7 we will describe an improved proof of Theorem 5.5.1 due to Huang and Durrett (2022). In Section 5.6 we give results for the survival time on the star graph that are the key to the proof, since they give lower bounds of how long the infection persists on vertices of large degree. In Section 5.7 we apply these results to show $\lambda_2 = 0$ when the degree distribution is subexponential, i.e., $\limsup_{k\to\infty}(1/k)\log p_k = 0$. The method of proof is due to Pemantle (1992), who, in his remarkable paper that initiated the study of the contact process on trees, proved that $\lambda_2 = 0$ for degree distributions that are **stretched exponential** $p_k \sim C\exp(-ck^\gamma)$ with $\gamma < 1$.

The goal of Huang and Durrett's (2022) paper was to show that if the degree distribution is exponentially bounded, i.e., $\sum_k p_k e^{\theta k} < \infty$ for some $\theta > 0$ then $\lambda_1 > 0$. In Section 5.8 we describe the solution of this problem due to Bhamidi, Nam, Nguyen, and Sly in 2019. In hindsight the proof is not very difficult. They add a new vertex ρ^+ above the root of the tree ρ that is always infected, and they do not allow recovery at the root until all other vertices have recovered. This modification of the graph enables one to write recursive equations, which lead to a remarkably simple proof

The paper of BNNS proves a lot, but there are still significant open probales. For example Theorem 5.8.5 shows that (i) for $\lambda < \lambda_i$ the process on a finite confiuration model graph with exponentially bounded degrees and n vertices survives for time $n^{1+o(1)}$ and (ii) when $\lambda > \lambda_{ii}$ the survival time is $e^{\Theta(n)}$. One would like to prove that $\lambda_i = \lambda_{ii} = \lambda_c$. In addition there is the question: "does the subcritical survival result give the right answer?" The supercritical behavior matches the results shown for a number of graphs above, and as remarked earleir is the best possible result for graphs with n vertices and $O(n)$ edges.

The chapter concludes in Section 5.9 with a look at the threshold-θ contact process, a discrete time system in which sites that have at least θ occupied neighbors at time n will be occupied with probability p at time $n+1$. Chatterjee and Durrett (2013) showed that if $\theta \geq 2$ and we consider the system on a random r-regular graph with $r \geq \theta + 2$ then the phaes transition is discontinuous. That is, as p is reduced from 1, the equilibirum density drops from a positive level at p_c to 0 for $p < p_c$. Identifying contact processes with attractive nonlinear birth rates and constant death rates is a very interesting research question.

6. Random Walks, Mixing Times

In this chapter we will investigate how the time for a random walk on a graph to reach equilibrium depends on the geometry of the graph. In Section 6.1 we review basic results about Markov chains and state two general results on the rate of convergence to equilibirum that are based on the size of the spectral gap. In Section 6.2 we introduce the notion of **conductance** for reversible Markov chains and prove **Cheeger's inequality** which relates the spectral gap to the "bottleneck ratio."

In Section 6.3 we begin to apply these results to bounding convergence times of random walks on random graphs. The first result is that if the minimum degree is at least three then the convergence time is $O(\log n)$. As shown in Section 6.4 when there are also vertices of degree two then convergence is slowed to $O(\log^2 n)$ because there are paths of length $\geq \delta \log n$ and if a random walk starts in the middle of one of them then time to escape from the path is $\geq c\delta^2 \log^2 n$.

Section 6.5 investigates how the rate of convergence to equilibrium on Erdős-Rényi graphs depends on the mean degree. When the graph is connected, i.e., the mean degreee is $c \log n$ with $c > 1$ then convergence occurs in time $O(\log n)$, but as the mean decreases to $\lambda > 1$ the time increases to $O(\log^2 n)$ due to the presence of long paths. It is interesting that $O(\log^2 n)$ is a worst case result but not a typical result. If the walk starts from a vertex chosen at random from an Erdős-Rényi graph with mean degree $\lambda > 1$ then the mixing time is $O(\log n)$, because it is unlikely to hit one of these long paths,

Section 6.7 introduces the **cutoff phenomenon**: which occurs when t_n is the time to converge to equilibrium and the distance from the stationary distribution goes from 1 ot 0 over a window of size $o(t_n)$. It is easy to see that this does not occur for simple random walk on the integers modulo n. But it does occur for random transpositions, riffle shuffles, random walk on the hypercube, and even for the Ehernfest chain

In Section 6.8 we describe results of Lubetzky and Sly (2010) on random d-regular graphs. They show that if $d \geq 3$ the mixing time is $\sim c_d \log_{d-1} n$ where $c_d = d/(d-2)$ and there is cutoff with a window of size $\sqrt{\log_{d-1} n}$. It is easy to see this is the right answer. By results in Chapter 1, the largest distance on the graph is $\sim \log_{d-1} n$. The random walk moves away from its starting point with probability $(d-1)/d$ and closer to it with probability $1/d$ so the drift is $(d-2)/d$. However, proving this is not so easy even when you get the hint that "the covering space of a random regular graph is a regular tree."

7. Voter Models, Coalescing RWs

The voter model was introduced independently by Clifford and Sudbury (1973) and Holley and Liggett (1975) on the d-dimensional integer lattice. It is a very simple model for the competition of two opinions. Each voter at the times of a rate one Poisson process wakes up and imitates the opinion of one of its neighbors chosen at random. While no one could think that this is an accurate model of how people form opinions, as a mathematical model, it has the desirable properties that it is simple and has interesting behavior.

The key to the analysis of the voter model ξ_t is a duality similar to the one used to study the contact process, which is introduced in Section 7.1. Here the dual, which traces the origin of the opinion at x at time t, is ζ_t^x a **coalescing random walk**. Particles in the dual move like independent random walks until they hit, at which point they coalesce to one. Writing a superscript A to indicate the initial of set of sites with opinion 1 in the voter model and a suserscript B to indicate the set of sites occupied by particles in the CRW, the duality equation is

$$P(\xi_t^A \cap B \neq \emptyset) = P(A \cap \zeta_t^B \neq \emptyset).$$

The chapter could have been named coalescing random walks (CRW) because they are the main focus of our investigations. In Section 7.2 we consider the behavior of CRW on the torus in dimensions $d \geq 3$. Let τ_N be the time it takes for the CRW to be reduced to one particle when it starts with one particle at each site. Cox (1989) has shown for the torus of side N in d-dimensions

Theorem 7.2.1. As $N \to 0$ τ_N/s_N converges im distribution to a limit we call τ where

$$s_N = \begin{cases} N^2 & d=1, \\ N^2 \log N & d=2, \\ N^d & d \geq 3. \end{cases}$$

It is clear that in one dimension τ_N has to be at least of order N^2. In $d \geq 3$ the time is of order the number of sites which we will see is the typical situation on random graphs. Two dimensions is a transition from N^2 to N^d behavior and has a logarithmic corrction factor due to the recurrence of random walks.

In Section 7.3 we begin the study of CRW on graphs by examining the hitting times for two random walks. **Aldous's Poisson Clumping Heuristic** allows us to compute the asymptotic behavior of the hitting time,including the constant. For example on a random r-regular graph, if we start random walks from two sites chosen at random then the time for them to hit is $\sim n(r-1)/(r-2)$. Another result of Aldous and Fill shows that given a sequence of graphs G_n, the hitting time T_n for walks starting from two randomly chosen starting points divided by its mean converges to an exponential distribution if we have

(H) the mixing time for the random walk on the graph is of smaller order of magnitude than the time for the two walks to hit.

A major theme in the chapter is that if (H) holds then coalescing random walk when properly rescaled converges to Kingman's coalescent.

In Section 7.4, we describe a result of Cooper et al (2012) which gives a bound on the coalescence time on general graphps with n vertices, average mean $\bar{d}$,and maximum degree $\Delta = O(n^{1-\epsilon})$. Here and in the next result $\mathbf{C}$ is the coalescence time we previously called τ_N. Particles start at each site of the graph perform independent continuous time random walks jumping at rate 1, and coalescing when they hit.

Theorem 7.4.1. Let G be a connected graph with n vertices, let $\nu = \sum_v d^2(v)/\bar{d}^2 n)$, and λ_1 is the second largest eigenvakue of the transiton matrix. Then

$$\mathbf{C} = O\left(\frac{n}{\nu(1-\lambda_1)}\right).$$

Section 7.5 is devoted to work of Oliveira (2012, 2013) on CRW on general graphs that confirms a conjecture from Aldous and Fill's (2002) book on Markov chains.

Let T_v be the hitting time of v and let $T^G_{hit} = \max_{u,v} E_u T_v$. Prove that there is a constant K so that

$$E\mathbf{C} \leq K T^G_{hit}$$

In addition, he proved convergence of the CRW o Kingman's coalescent.

In Section 7.6 we turn our attention to the short time behavior of CRW starting from one particle at each site. The combined work of a half dozen people from 1977-1980 showed that on $\mathbb{Z}^d$ the density of occupied sites p_t obeyed

$$p_t \sim \begin{cases} c_1/t^{1/2} & d=1 \\ c_2(\log t)/t & d=2 \\ c_d/t & d\geq 3 \end{cases}$$

In the second half of Section 7.6 we describe work of Hermon, Li, Yao, and Zhang (2022) who show that for random graphs generated by the configuration model the decay of the density of CRW is like that on the torus in $d \geq 3$, i.e., it decays like C/t.

In Section 7.7, the focus returns to the voter model. We consider three classes of finite graphs: the complete graph, the d-dimensional torus in $d \geq 3$, and graphs generated by the configuration model. On a finite graph the voter model will eventually enter an absorbing state that represents consensus. Results for CRW tell us how long it takes for this to occur: it is the same as time s_N for the coalescing random walk starting from one particle at each site to reduce to one particle . Results of Cox with Greven in 1990 and with Chen and Choi in 2016 describe how the behavior of the voter model starting from product measure with density p evolves. At times that are large but $o(s_N)$ the system looks like the voter model equilibrium ν_p. At times $O(s_N)$ it is a mixture of these equilibrium states where the mixture is given by the distribution of a Wright-Fisher diffusion., which has generator $(1/2)\gamma y(1-y)f''(y)$.

8. Coevolving systems

This chapter is concerned with models in which the states of the vertices and the connections between them coevolve. This topic has been studied in the physics literature for many years. However, there are only a small number of rigorous results, which will be our focus here. We consider three types of systems.

8.1. Evolving voter models

During 2010-2011 program on Complex Networks at SAMSI, an NSF funded institute in Reseach Triangle Park, Rick Durrett, James Gleeson, Alun Lloyd, Peter Mucha, Feng Shi, David Sivakoff, Josh Socloar, and Chris Varghese were part of a working group on *Dyanmics On and Of Graphs*. Inspired by work of Holme and Newman (2006), they formulated a model in which on each step an oriented edge (x, y) is picked at random. If the two connected individuals have the same opinion nothing happens. If they hold different opinions then: with probability $1 - \alpha$ vertex x imitates the opinion of y, otherwise, i.e., with probability α, the link between them is broken and x makes a new connection to an individual chosen at random (i) from those with the same opinion, or (ii) from the network as a whole.

The two versions of the model are called (i) **rewire to same** and (ii) **rewire to random**. The evolution of the system stops at the first time τ when there are no longer any discordant edges (that connect individuals with different opinions). At time τ in each connected component of the graph all the individuals have the same opinion. For several months in 2010, participants in the SAMSI group thought that the two versions had the same behavior, but then they did some simulations.

Figure 1 gives simulation results starting from Erdös-Rényi graphs with $N = 100,000$ nodes and average degree $\lambda = 4$. Opinions are initially assigned randomly with the probability of opinion 1 given by $u = 0.5$, 0.25, 0.1, and 0.05. The figure shows the final fraction ρ of voters with the minority opinion from five realizations for each u. In case (i), there is a critical value α_c which does not depend on u, with $\rho \approx u$ for $\alpha > \alpha_c$ and $\rho \approx 0$ for $\alpha < \alpha_c$. In case (ii), the transition point $\alpha_c(u)$ depends on the initial density u. For $\alpha > \alpha_c(u)$, $\rho \approx u$, while for $\alpha < \alpha_c(u)$ we have $\rho(\alpha, u) = \rho(\alpha, 0.5)$. Since all of the $\rho(\alpha, u)$ agree with $\rho(\alpha, 0.5)$ when they are $< u$, we call the graph of $\rho(\alpha, 0.5)$ on $[0, \alpha_c(0.5)]$ the **universal curve**.

In Section 8.1.1, we will give heuristic arguments and numerical results from Durrett et al (2012) in support of these conjectures. In Section 8.1.2, we will describe work of Basu and Sly (2017) that rigorously prove the existence of a phase transition between rapid disconnection and prolonged persistence (see Theorems 8.1.4 and 8.1.5). While the proofs of these results required an incredible amount of ingenuity they only work for dense Erdős-Rényi graphs and do not provide much insight into the exotic behavior shown by rewire to random.

8.2. Evolving SIS epdiemics

Gross, D'Lima and Blasius (2006) introduced the following model: "we consider a network with a constant number of nodes N, and bidirectional links K. The nodes represent individuals that can be either susceptible (S) or infected (I).

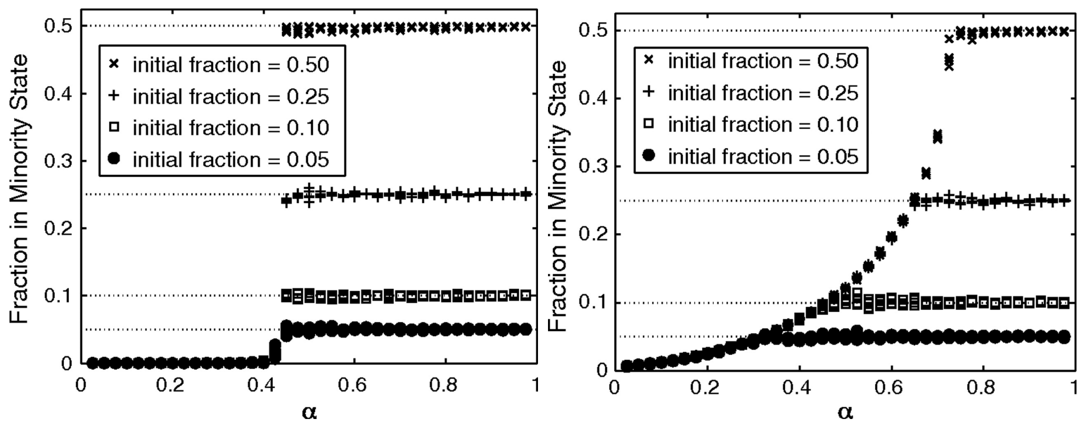

Figure 1: Simulation results for rewire to same model (left panel) and rewire to random (right panel)

- On every time step and for every link connecting an infected with a susceptible (an SI link) the susceptible becomes infected with a fixed probability p, and the infected recovers from the infection with probability γ, becoming susceptible again.

- In addition we allow susceptible indviduals to rewire their links. With probability w for every SI link, the susceptible breaks its link to the infected and rewires to another randomly chosen susceptible. Double connections and self-connections are not allowed to form in this way.'

Simulations suggested the phase diagram drawn in Figure 2

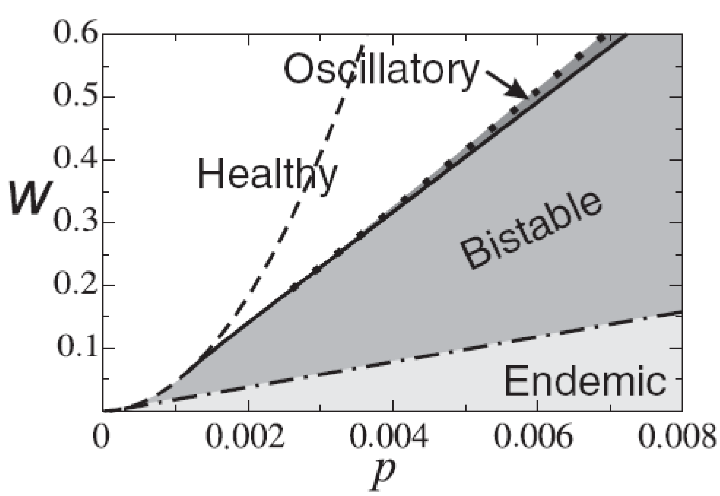

Figure 2: Simulation with $N = 10^5$ sites, $K = 10^6$ edges (average degree 20), and $\gamma = 0.002$. Dotted line is predicted critical value $p^*(w)$.

In the **Healthy** region the infection dies out. In the **Endemic** region there is positive probability of the infection surviving from a single infected individual and there is a unique equilbrium state, while in the **Bistable** region the epidemic always dies out starting from a single infected but there is a nontrivial stationary distribution. The authors observe that the **Oscillations** can only be observed in a relatively small parameter region so perhaps they are a finite size artifact. The most striking observation is that the equilibrium density jumps from 0 to a positive level when increasing p moves the parameters from the **Healthy** to the **Bistable** region.

Note that $p^*(w)$ in an increasing function of the rewiring probability. Though this may seem obvious there is no mathematical proof. As usual physicists, who are free from the burdens of rigor. Gross et al (2006) argue that

$$p^*(w) = \frac{w}{\langle k \rangle [1 - \exp(-w/\gamma)]}.$$

Note that this corresponds to $p^*(0) = \gamma/\langle k \rangle$, but $p^* = w/\langle k \rangle$ for $w \gg \gamma$."

In Section 8.2.1, we give results obtained by Zanette (2007) using mean-field theory and the pair approximation. As far as we know there are no rigorous results for the original model. da Silva, Oliveira, and Valesin (2021) and Schapira and Vlesin (2023) have consider contact processed on dynamical d-regular graph where the graph evolves by degree conserving Volz-Meyers edge swap dynamics, described in Section 8.3.1, scaled so that in the limit as the number of nodes $n \to \infty$, each edge is involve in exchanges at rate v.

In Section 8.2.2, we describe rigorous work of Chatterjee, Sivakoff, and Wascher (2012) on an SIS epdemic with model with **link inactivation**, i.e., susceptibles may drop their connections to infectious neighbors. This version of the model is much simplerto analyze since the graph does not change over time. This version might be more relevant for applications, since it seems a little harsh not totalk to one of your friends for the rest of your life just because they got sick.

8.3. Evolving SIR epidemics

In the summer of 2018, postdoc Matt Junge, graduate student Zoe Huang, and I worked with four students in the Duke math departments 8 week summer REU: Duke Opportunities in Math or DOMath. Three of the students continued to work with Matt and I in the Fall. The object of study was evoSIR, an epidemic model in which each infection lasts for exactly time 1 with the new feature that susceptibles who are neighbors of an infected break their connections at rate ρ and connect to a randomly chosen neighbor.

In order for the infection to be transmitted along an edge, it must occur before any rewiring and before time 1, so the transmission probability is

$$\tau_{f,r} = \frac{\lambda}{\lambda + \rho}(1 - e^{-(\lambda+\rho)}). \qquad (0.0.2)$$

where f is for fixed time infections, and r is for rewiring. They also considered a much easier model delSIR in which connections were deleted instead. Their first result shows that evoSIR has the same critical value as delSIR.

Theorem 8.3.1. *The critical value for the fixed time epidemic with rewiring is given by* $\mu_c = 1/\tau_{f,r}$. *Moreover, in the subcritical regime the ratio of the expected epidemic size in delSIR to the size in evoSIR converges to 1.*

One can also compute the distribution of the number of secondary infections caused by an infected vertex early in the epidemic and hence the probability of a large epidemic. Since the delSIR model is essentially percolation, this probability is the same as the fraction of individuals infected in a large epidemic, and hence goes to 0 at the critical value $\mu_c = 1/\tau_r^f$. In evoSIR, as the simulations in Figures 3 suggest the fraction of individuals infected is discontinuous at the critical value.

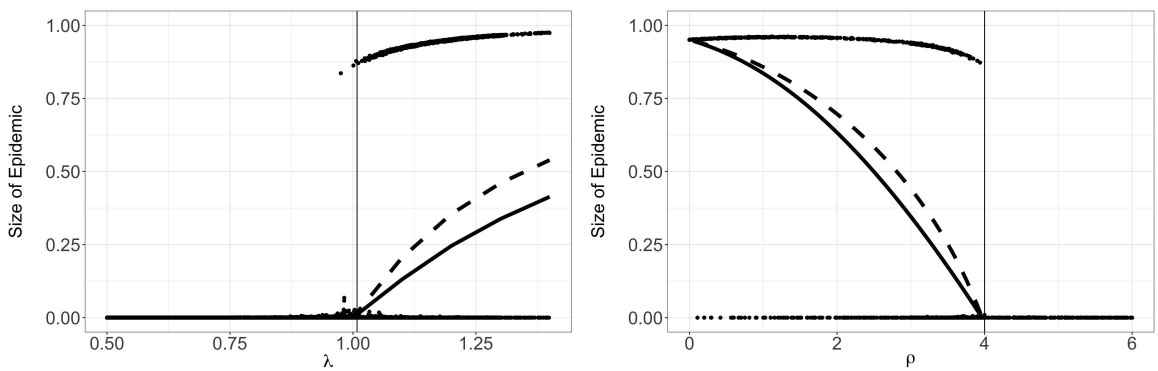

Figure 3: Simulation of the fixed time evoSIR on an Erdős-Renyi graph with $\mu = 5$. In the left panel, the rewiring rate $\rho = 4$ and λ varies with $\lambda_c \approx 1.0084$ in agreement with Theorem 8.3.1. In the right panel, $\lambda = 1$ and ρ varies with $\rho_c \approx 4$. In both panels the top curve is the final size of the evoSIR epidemic, and solid curve is the final size of the delSIR epidemic with the same parameters. The dashed line is an approximation developed in Jiang et al (2019) that did not do a good job of predicting the final size in evoSIR.

As the authors of Jiang et al (2019) were finishing up the writing of their paper they learned of work of Britton, Juher, and Saldana (2016) who studied a fmaily of models called SIR-ω that interpolated between evoSIR and delSIR. See Section **??**. Work of Leung, Ball, Sirl, and Britton (2016) demonstrated the paradoxical fact that individual preventative measures (i.e., dropping connections to infected sites) may increase the final size of the epidemic. It is interesting to note that their simulations did not show a discontinuous phase transition.

Meanwhile back in Durham, Durrett and his graduate student Dong Yao used ideas of Janson, Luczak, and Windridge (2015) described in Section 4.5 to analyze the evolving SIR model on graphs generated by the configuration model. The trick is to grow the graph at the same time as you run the epidemic. Unfortunately their construction had an error: the algorithm resulted in rewiring connections between $I - I$ pairs, which should not occur. All was not lost: the formula was an upper bound on the true extent of the epidemic and a modification of the proof worked for the evoSI epidemic.

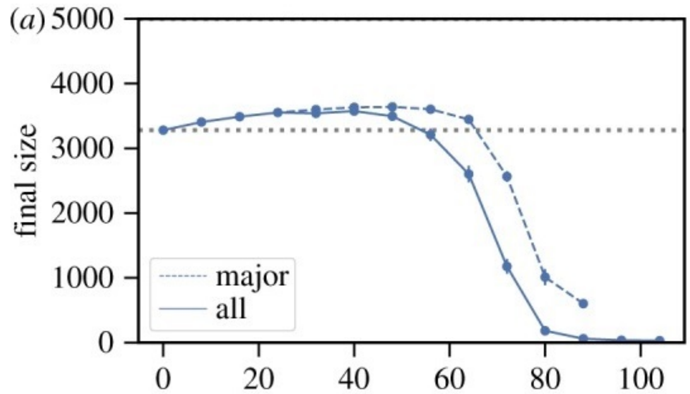

Figure 4: Social distancing can lead to an increase in the final epidemic size in the configuration model. The x-axis indicates the rewiring rate. The horizontal line is the final size when $\omega = 0$.

While the paper of Durrett and Yao (2020) was being refereed, work of Ball and Brittion posted on the arXiv showed that their predictions of DY(2020) were not correct. BB(2022) studied the SIR-ω model on Erdös-Renyi graphs. This is tractable because one can derive a limiting differential equation for S, I, i_E, and w, the fraction of susceptibles infecteds, $I - S$ edges, and $S - S$ edges produced by rewiring. They were able to completely analyze the model except for one small gap, which was later filled by W. Chen, Y. Hou, and D. Yao (2022).

Despite the large amount of work on this topic and the existence of a complete solution on Erdős-Rényi random graphs, much work remains to be done. The biggesst open problem is to understand the behavior of evoSIR on graphs generated by the configuration model. While the proof of Durrett and Yao (2022) was not correct the answer may very well be correct. It gives the right answer in the Erdős-Rényi case, see Theorem 8.3.4.