Statement of research

Jonathan Hermon *

My research is in discrete probability theory with an emphasis on problems related to random walks, the theory of mixing times of Markov chains and the cutoff phenomenon. Markov chains are random processes that retain no memory of the past. This is an area on the interface of mathematics, statistical physics and theoretical computer science. Consequently, a discovery in this area can have impact in different disciplines. In my work I attempt to develop general theory and study large families of instances. A big part of my work is dedicated to particle systems, including some classical statistical mechanics models like the Exclusion process and the Zero Range process as well as models for a spread of infection and for a social network.

1 Brief background

The **random walk** on a graph is the process that, at each time chooses an edge uniformly at random from among those incident to its current location, independently from everything it has done previously, and then moves to the vertex at the other endpoint of that edge. A Markov chain on a countable state space V with a transition kernel P evolves as follows: when at state x it moves to a new state y with probability $P(x,y)$. A Markov chain is called **irreducible** if every state can be accessed from any other state in some finite time.

Reversiblity is the condition that the distribution of the stationary chain is invariant under reversal of time. This is equivalent to the condition that it can be represented as a weighted random walk in the above sense. This is the best understood class of Markov chains. The **spectral-gap** is the smallest non-zero eigenvalue of $I-P$.[1] For reversible chains the spectral-gap gives the asymptotic exponential rate of convergence to equilibrium.

An important class of graphs in combinatorics and theoretical computer science is that of expanders. A sequence of finite graphs is called an **expander** family if the spectral-gaps of the random walks on them are uniformly bounded away from 0. A graph is called **(vertex-)transitive** if it looks the same from any vertex (i.e., any vertex can be mapped to any other vertex by a symmetry of the graph).

Led by the pioneer works of Aldous and Diaconis and driven by applications such as Monte Carlo simulations and approximate counting problems the modern theory of mixing times of Markov chains became in the last few decades a lively and central part of modern probability theory with ties to other areas, such as e.g., particle systems, statistical mechanics[2], combinatorics and representation theory. Among the most fundamental quantities associated with a finite irreducible Markov chain is its **ε-mixing time**, which is the number of steps required for it to get within distance ε from the stationary distribution of the chain. In computer science, it is the main component in the running time of randomized algorithms for sampling, approximate counting and volume estimation [55].[3] The default choice for the 'distance' above is the **total-variation** distance $\|\mu-\nu\|_{\mathrm{TV}} := \sum_x \frac{1}{2}|\mu(x)-\nu(x)|$. For $\varepsilon<1/2$ the choice of ε may change the mixing time only by an $O(\log\frac{1}{\varepsilon})$ factor, and hence is insignificant.

We remark that the technical term *almost surely*, used extensively below, means 'with probability 1'.

2 No percolation at criticality on transitive graphs with fast heat-kernel decay

In **Bernoulli bond percolation**, each edge of a connected, locally finite graph G is either deleted or retained at random with retention probability $p \in [0,1]$, independently of all other edges. We denote the random graph obtained this way by ω_p. Connected components of ω_p are referred to as **clusters**. Percolation theorists are primarily interested in the geometry of the open clusters and how this geometry varies as the parameter p is varied. We are particularly interested in *phase transitions*, where the geometry of ω_p changes abruptly as we vary p through some special value. The first basic result about percolation, without which the model would not be nearly as interesting, is that for most infinite graphs (excluding e.g. one-dimensional counterexamples such as the infinite line graph $\mathbb{Z}$), percolation undergoes a *non-trivial phase transition*, meaning that the **critical probability**

$$p_c(G) = \inf\{p \in [0,1] : \omega_p \text{ has an infinite cluster almost surely}\}$$

is strictly between zero and one. Once we know that the phase transition is non-trivial, the next question is to determine what happens when p is exactly equal to the critical value p_c. This is a much more delicate question.

*University of Cambridge, Cambridge, UK. E-mail: E-mail: jh2129@statslab.cam.ac.uk .

[1] For continuous-time chains replace $I-P$ by $-\mathcal{L}$, where $\mathcal{L}$ is the infinitesimal generator and the time t transition-kernel is $e^{t\mathcal{L}}$.

[2] Many of the most studied Markov chains are of a statistical mechanics nature, such as the Glauber dynamics for the Ising model and the exclusion process.

[3] In other words, a rigorous justification of simulation results requires a theoretical bound on the mixing time.

Indeed, the best known open problem in percolation (and arguably in probability theory) is to prove that critical percolation on the d-dimensional hypercubic lattice $\mathbb{Z}^d$ does not contain any infinite clusters almost surely for every $d \geqslant 2$. In their highly influential paper [12], Benjamini and Schramm proposed a systematic study of percolation on general *transitive* graphs. Prominent examples of transitive graphs include Cayley graphs[4] of finitely generated groups. The following is among the most important of the many outstanding conjectures that they formulated.

Conjecture 2.1 (Benjamini and Schramm 1996)**.** *Let G be a quasi-transitive graph. If $p_c(G) < 1$ then critical Bernoulli bond percolation on G has no infinite clusters almost surely.*

Previous progress on Conjecture 2.1 beyond the Euclidean setup can briefly be summarised as follows. Benjamini, Lyons, Peres, and Schramm [7] proved that Conjecture 2.1 holds for every *unimodular, nonamenable* transitive graph. Here, *unimodularity* is a technical condition that holds for every Cayley graph and every amenable[5] transitive graph. Hutchcroft [39] verified Conjecture 2.1 for transitive graphs of *exponential growth* (i.e., the number of vertices in a ball of radius r grows exponentially in r).

In a recent work with Hutchcroft [24] we verify the conjecture for unimodular transitive graphs satisfying that the time n return probability to the origin o, $p_n(o,o)$ by the random walk is at most $O(\exp(-cn^{\gamma}))$ for some $\gamma > \frac{1}{2}$ and some constant $c > 0$ (non-amenability is equivalent to $\gamma = 1$). There are several known families of groups of intermediate growth known to satisfy this condition. Moreover, we verify Schramm's locality conjecture in this case, which says (in some precise sense) that for transitive graphs the critical parameter is determined by the local geometry of the graph. Our proof relies on Hutchcroft's recent two-ghost inequality [40] along with some random walk analysis.

Future plans: Together with Hutchcroft we intend to work on other problems in percolation. For instance, we are interested in the problem of determining for super-critical percolation the rate of decay of a cluster's size, conditioned on that cluster been finite. Establishing exponential decay in the non-amenable setup is the main missing ingredient towards establishing that the infinite clusters have a positive anchored expansion constant, which is a famous open problem. At the moment it is not even known that for $p > p_c$ the expected size of a cluster, conditioned on being finite, is finite. We have an argument which implies that in the non-amenable setup all of the moments are finite and that the distribution of the radius of a finite cluster has an exponentially decaying tail.

3 Mixing time for general symmetric Exclusion Process

The **symmetric exclusion process** $\mathrm{EX}(k)$ on a finite, connected graph $G = (V, E)$ is the following continuous-time Markov process. In a configuration, each vertex contains either a black particle or a white particle (where particles of the same colour are indistinguishable), and the number of black particles is $k < |V|$. For each edge e independently, at the times of a rate 1 Poisson process, switch the particles at the endpoints of e.

The exclusion process is among the most fundamental and best-studied processes in the literature on interacting particle systems [44, 45], with ties to card shuffling [25, 58], statistical mechanics and a variety of other processes [44]. Apart from having a rich literature on the model on infinite graphs, such as the lattices $\mathbb{Z}^d$, the exclusion process on finite graphs has been one of the major examples driving the quantitative study of finite Markov chains.

Oliveira [51] showed that the mixing time of the exclusion process on $G = (V, E)$ is at most of order of the mixing time of a single particle times $\log |V|$. He conjectured that one can bound it from above (up to a constant factor) via the mixing time of the same number k of independent particles on G (i.e., of the continuous-time Markov chain on V^k in which the coordinates are independent and each coordinate is performing a random walk on G). In a joint work with Pymar [30] we verify the conjecture when G is d-regular and the particle density is not too sparse if either of the following holds: (1) $d/\log_{|V|/k}|V|$ is bounded from below or (2) the diameter[6] of G is at least $\log^5|V|$. Moreover, we verify the conjecture without these assumptions up to a $O(\log\log|V|)$ factor. In fact, we obtain general upper and lower bounds which in many cases are sufficiently strong to imply the conjecture and determine the order of the mixing time. The **only** natural example we are aware of in which our bound overshoots the order of the mixing time is when G is an expander, in which case our bound is $O(\log|V|\log\log|V|)$ which is still an improvement over the previously best bound of $O(\log^2|V|)$ [51]. Furthermore, we determine the order of the mixing time for arbitrary k when G is: The hypercube $(\{\pm 1\}^d, \{\{x,y\} : \sum_i |x_i - y_i| = 2\})$, resolving a conjecture of Wilson [58], or is transitive (i.e., 'symmetric') satisfying a technical condition called *moderate growth* [15].

Future plans: Attempt to remove the $\log\log|V|$ factor (at least for expanders) and to extend our results to the **interchange process**, which is the analog of the exclusion process when the particles are all distinguishable.

[4]The Cayley graph of a group G w.r.t. a symmetric set of generators $S = S^{-1} := \{s^{-1} : s \in S\}$ is the graph with vertex set G in which x, y are connected if $xy^{-1} \in S$.

[5]A graph is **non-amenable** if the return probability to the origin of the walk decays exponentially in time. Otherwise it's **amenable**.

[6]Defined as the maximal graph distance $d(x,y)$ between a pair of vertices x, y, where $d(x,y)$ is the minimal number of edges along a path in the graph connecting x and y.

4 Zero-Range-Process

Introduced by Spitzer in 1970 [56], the zero-range process (**ZRP**) is one of the most widely studied models of interacting particles [44, 45]. It describes the evolution of a fixed number $m \geq 1$ of indistinguishable particles evolving on a finite set of sites V. The exit rate away from a site depends only on the number of particles currently occupying the site. When a particle jumps, it picks its next position according to some transition matrix P.

Together with Salez [32] we showed that the spectral-gap of a general zero range process is determined by the spectral-gap of a single particle the spectral-gap of the corresponding mean field ZRP. More precisely, we showed that when P has stationary distribution π the Dirichlet form corresponding to the generator of the ZRP can be compared with that of the corresponding *mean-filed* ZRP (here 'mean field' means that all of the rows of the transition matrix are equal to π) with the same exit rates, up to a factor which is precisely the spectral-gap of P (i.e. of one particle) denoted gap_P. Moreover, the ratio of the spectral-gaps of the two ZRPs is shown to be gap_P.

This is in the spirit of Aldous' famous spectral-gap conjecture for the interchange process, now resolved by Caputo et al. [14]. Our main inequality decouples the role of the geometry (defined by the jump matrix P) from that of the kinetics (specified by the exit rates). Among other consequences, the various spectral-gap estimates that were so far only available on the complete graph or the d-dimensional torus now extend effortlessly to arbitrary geometries. As an illustration, we determine the exact order of magnitude of the spectral-gap of the rate-one ZRP on any regular graph and, more generally, for any doubly-stochastic (i.e., sum of entries along each column is 1) jump matrix. We also determine the order of the mixing time when P corresponds to a random walk on an expander and the exit rates all equal one, when m is proportional to $|V|$ (namely, it is $\Theta(|V|)$).

In a recent work with Salez [31] we consider the ZRP with arbitrary bounded monotone rates on the complete graph, in the regime where the number of sites diverges while the density of particles per site converges. We determine the asymptotics of the mixing time from any initial configuration, and establish the cutoff phenomenon (see definition below). The intuitive picture is that the system separates into a slowly evolving solid phase and a quickly relaxing liquid phase: as time passes, the solid phase dissolves into the liquid phase, and the mixing time is essentially the time at which the system becomes completely liquid.

Future plans: It appears that determining the mixing time for the rate 1 ZRP is a very challenging problem. The only cases where it is understood, apart from the case of expanders mentioned above, are the cases of the cycle (due to a bijection with the exclusion process specific to this case [43]) and the mean field case [31, 48]. In both cases the model was shown to exhibit an abrupt convergence to stationarity known as cutoff (see below). Together with Salez we intend to exploit new tools we are developing in order to analyze the mixing time and prove cutoff in various natural families of instances.

5 The cutoff phenomenon, general Markov chain theory and sensitivity of mixing times

A Characterization of the cutoff phenomenon: Loosely speaking, the **cutoff phenomenon** occurs when over a negligible period of time, the distance from equilibrium drops abruptly essentially from its maximal value to 0. More precisely, a sequence of chains is said to exhibit cutoff if the ε-mixing time is asymptotically independent of ε. The study of the cutoff phenomenon is a major topic in modern probability, and despite remarkable progress over the last 30 years, there are still only relatively few examples which are completely understood. A natural problem, posed by Aldous and Diaconis in their seminal 86 paper [1] is to find general conditions for cutoff.

The **hitting time** of a set is the first time it is visited by the chain. Aldous [3] showed that for reversible chains the order of the mixing time can be understood in terms of expected hitting times of sets. This was later substantially refined [51, 52], however was still too coarse for describing the cutoff phenomenon.

In a joint work with Basu and Peres [6] we derive a necessary and sufficient condition for cutoff for reversible chain in terms of concentration of hitting times of large sets which are "worst" in some sense. Previously there was no such general characterization for cutoff. It is a consequence of the following quantitative general hit-mix relation we prove: for every ε the ε total-variation mixing time is close to a certain hitting-times parameter $\text{hit}(\varepsilon)$

We used our characterization to give a simple spectral necessary and sufficient condition for cutoff for chains that the graph supporting their transitions is a tree or is 'similar' (quasi-isometric) to an interval (a tree is a connected graph with no cycles). This extended previous results on birth and death chains [17], for which the graph is an interval. The result from [6] and the method of proof, involving a novel usage of maximal inequalities for self-adjoint contracting positive operators led to a related body of work which we now describe.

Cutoff for Ramanujan graphs: Ramanujan graphs is an important class of regular graphs in extremal graph theory with applications to constructions of error correcting codes and in quantum computing. They were defined and first constructed by Lubotzky, Phillips, and Sarnak [47]. A finite connected d-regular graph is said to be **Ramanujan** if all of the non-unit eigenvalues of the transition matrix of the random walk on it lie in the interval $[-\rho_d, \rho_d]$, where $\rho_d := \frac{2\sqrt{d-1}}{d}$ is the spectral radius of the random walk on a d-regular tree. In light of the

Alon-Boppana bound [50], Ramanujan graphs are "optimal expanders" as they have asymptotically the largest spectral-gap possible for regular graphs of that degree.

Lubetzky and Peres [46] proved that random walks on a sequence of Ramanujan graphs of diverging sizes exhibit cutoff. In [19] we used the above characterization of cutoff via hitting times to give a short alternative proof (under a mild condition known to hold in all existing examples of Ramanujan graphs, and which is always satisfied in the transitive case). Our proof has the advantage that it only requires the negative eigenvalues of the transition matrix to be bounded away from -1, rather than to lie in the "Ramanujan interval" $[-\rho_d, \rho_d]$.

A characterization of the L_∞ mixing time and of the log-Sobolev constant: A common choice for the metric w.r.t. which the distance from equilibrium π is measured is the L_∞ norm $\|\mu - \pi\|_{\infty,\pi} := \max_x |\frac{\mu(x)}{\pi(x)} - 1|$, giving rise to the L_∞ mixing time, $t_{\mathrm{mix}}^{(\infty)}$. Using similar techniques as in [6], with Peres [28] we gave the first ever characterization of $t_{\mathrm{mix}}^{(\infty)}$ (and also of the "relative-entropy mixing-time") which is sharp up to a constant factor. Our bound is probabilistic and is also expressed in terms of hitting times.

While the spectral gap is a natural and simple parameter, the *log-Sobolev constant* c_{LS}, is a more involved quantity. When one first encounters c_{LS}, it may seem like an artificial parameter that "magically" gives good bounds on $t_{\mathrm{mix}}^{(\infty)}$. We give a new extremal characterization of the log-Sobolev constant as a weighted version of the spectral-gap. The so called hypercontractivity phenomenon is characterized by c_{LS} [16]. We show that it can also be understood via maximal inequalities.

As an application we determined the order of $t_{\mathrm{mix}}^{(\infty)}$ in the case that the graph supporting the transitions is a tree, and showed that it is robust in the sense that changing the edge weights[7] by a constant factor may change it only by a constant factor. We also show that the mixing time is robust under the operation of multiplying each row i of the generator by some r_i, as long as the r_i's are uniformly bounded from above and below.

Sensitivity of mixing and of cutoff: While there are several sophisticated analytic and geometric tools for bounding $t_{\mathrm{mix}}^{(\infty)}$, none of them determines it up to a constant factor. These bounds are *robust* under small changes to the geometry. An important question is whether mixing times are robust under small changes to the geometry of the Markov chain. For instance, can bounded perturbations of the edge weights change the mixing time by more than a constant factor? Similarly, how far apart can the mixing times of the random walks on two roughly-isometric[8] graphs of bounded degree be? Different variants of this question were asked by various authors such as Pittet and Saloff-Coste [54], Kozma [42, p. 4], Diaconis and Saloff-Coste [16, p. 720] and Aldous and Fill [2] Our characterization of $t_{\mathrm{mix}}^{(\infty)}$ guided us [22] in the construction of a counter-example consisting of a sequence of bounded degree graphs, such that replacing some of their edges by two edges in series increases the order of $t_{\mathrm{mix}}^{(\infty)}$ of the corresponding random walk (by an optimal factor). This serves as a cautionary note on the possibility of developing sharp geometric bounds on mixing times.

With Peres [29] we also gave examples demonstrating that the occurrence of cutoff is itself sensitive to bounded perturbations, even when the edge-weights are changed only in an asymptotically negligible manner.

The power of averaging over two consecutive steps: The **period** of a chain with transition matrix P is the greatest common denominator of $\{n : P^n(x,x) > 0\}$ (this is independent of x when the chain is irreducible). To converge to equilibrium a discrete time chain must have period 1. Even when the period is one, when the graph supporting the transitions is nearly bipartite (i.e., it has two large sets with few edges between pairs belonging to the same set) the mixing time in discrete-time can be much larger than of its continuous-time version[9]. Reversible Markov chains can have period either 1 or 2. It is thus plausible that averaging over two consecutive time steps is enough to ensure the mixing time is not larger than that of the continuous-time version of the chain. A quantitative version of this question was conjectured by Aldous and Fill [2]. Using the maximal inequalities techniques from [6], together with Peres [27] we verified their conjecture by showing that the process which makes a single lazy step at time 0 (i.e., either stays put or moves according to P with equal probability) and then evolves according to P reaches within distance ε from equilibrium (in total variation) around the same time the continuous-time chain does (up to smaller order terms). In particular, cutoff in these two setups are equivalent. In [28] we proved a certain extension of this result to other distances from equilibrium, such as the L_∞ and relative-entropy distances.

Cutoff for random Abelian Cayley graph of diverging degree: An informal conjecture of Aldous and Diaconis [1] is that when the number of generators k of a Cayley graph (see footnote 4) of a finite group G is large, and the generators are picked uniformly at random, the random walk should exhibit cutoff around a time $t(k, |G|)$

[7] Any reversible chain on a countable state space can be represented as a weighted random walk on a graph, in which each edge is assigned a positive weight, and in each step the walk picks an edge incident to its current location with probability proportional to its weight and crosses to its other end-point.

[8] Two graphs are roughly-isometric if there is a map from one to the other which shrink or distort distances by at most some multiplicative and additive factor, and that every vertex is within some constant factor of the image of the map.

[9] The continuous-time version of the chain has mean 1 exponentially distributed waiting times between jumps.

which is independent of the algebraic structure of the group in the following sense: For some numbers $a(k,|G|)$ the ratio of the ε-mixing-time and $a(k,|G|)$ converges to 1 in distribution for all ε, provided k diverges with $|G|$. Hildebrand [36] verified this for cyclic groups when $k = \lceil \log^a |G| \rceil$ for some $a > 1$. A partial answer was given by Duo and Hildebrand [18] for general groups when $k = \lceil \log^a |G| \rceil$ for some $a > 1$. Wilson [57] proved cutoff when the group is $(\mathbb{Z}/2\mathbb{Z})^d$, conditioned on the event the generators generate the group. Hough recently proved this when G is $\mathbb{Z}/p\mathbb{Z}$ for p prime and the number of generators k satisfies $1 \ll k \ll \frac{\log p}{\log \log p}$.

Together with Thomas [33] we verify this for an arbitrary sequence of finite Abelian groups G_n admitting representations $G_n = \oplus_{i=1}^{d_n} \mathbb{Z}/m_{i,n}\mathbb{Z}$ provided that $k = k_n$ is slightly larger than d_n, that $\min_i m_{i,n} > |G|^{1/k}\sqrt{\log k}$ and that $\log k \ll \log |G_n|$. We provide a unified approach for all ranges of k and give a general description of the cutoff time as the time it takes the distribution of the simple random walk on $\mathbb{Z}^k$ to obtain $\log |G|$ entropy. As opposed to the aforementioned works, our analysis is refined enough to determine the exact profile of convergence to equilibrium (namely, a gaussian profile, under an appropriate scaling). It is interesting to note that in the spirit of the conjecture of Aldous and Diconis the cutoff time is dimension free, i.e., independent of d. Finally, we show that for some constant C the spectral-gap of the walk lies in $[\frac{1}{C}|G|^{-2/k}, C|G|^{-2/k}]$ with probability at least $1 - C2^{-k}$. This extends a classic result of Alon and Roichman [4] beyond the case $k \geq C \log |G|$.

We are currently working on showing that the diameter of the graphs agrees up to smaller order terms with a deterministic lower bound obtained by a counting argument (namely the radius of an L_1 ball in $\mathbb{Z}^k$ of volume $|G|$).

Random walk in evolving random environment Dynamical percolation on a graph G with parameter p and rate μ is the stochastic process in which the status of each edge is updated at rate μ to be open (i.e. present) (resp. closed, i.e. deleted)) with probability p (resp. $1-p$) independently of everything else. It is interesting to consider the case that μ tends to 0 as the size of the graph diverges, as a model of a slowly evolving network. This interpretation motivated several works on the mixing time of the random walk on dynamical percolation. So far the model has been studied mostly on the complete graph and on tori [53]. Together with Sousi [35] we consider dynamical percolation on arbitrary graphs. We develop a general comparison technique between the random walk on G and the chain involving both the walk and the evolving environment, which allows us to extend the mixing result of [53] to Cayley graphs of moderate growth and their hitting-times result to arbitrary graphs. Namely, we show that the hitting times for the walk co-ordinate of the process are at most $O(\frac{t_{\mathrm{hit}}}{\mu p})$, where t_{hit} is the maximal expected hitting time for the simple random walk on the graph. Finally, we obtain an upper bound of $O(\frac{d \log d}{\mu p})$ for the mixing time of the full process on the d-dim hypercube $\{0,1\}^d$. There is no known analog in the static case (i.e. for the random walk on the giant component of super-critical percolation on the hypercube).

Future plans:

- A natural question is what is the smallest change to a graph ensuring cutoff for the random walk? In a joint work with Sly and Sousi we plan to show that for an arbitrary sequence of graphs $G_n = (V_n, E_n)$ of increasing even sizes, of uniformly bounded degree, and of connected components[10] of size at least 3, adding a random perfect matching of V_n (chosen uniformly at random) to the edge set results with high probability with graphs on which the random walk exhibits cutoff (the precise statement is slightly more involved).

- In a work in preparation with Benjamini, Tessera and Tointon we show that a sequence of finite connected vertex-transitive graphs whose mixing times are comparable with the maximal expected hitting time of a state, denoted by t_{hit}, satisfies that when viewed as metric spaces with the usual graph distance, their Gromov-Hausdorff (**GH**) scaling limit (after a normalization by their diameters) is the unit cycle S^1. Key to the proof is one further insight regarding the structure of transitive graphs in line of current activity on the topic: If G_n are transitive graphs of degree D_n satisfying that $|G_n| = O(D_n\,(\mathrm{diameter}(G_n))^d)$ and the GH scaling limit of G_n is a d-dim torus, then G_n are uniformly quasi-isometric to d-dim Euclidean tori of the same diameter.

- In [23] I show that for a sequence of transitive graphs G_n, the **cover-times** $\tau_{\mathrm{cov}}(G_n)$, defined as the first time by which each vertex is visited at least once by the random walk, are concentrated iff $\mathbb{E}[\tau_{\mathrm{cov}}(G_n)] \times \mathrm{gap}(G_n)$, diverges, where $\mathrm{gap}(G_n)$ is the spectral-gap of the walk on G_n. Together with N. Berestycki we are trying to prove that under the mild condition that $\mathrm{gap}(G_n)t_{\mathrm{hit}}$ diverges, the cover times concentrate around $t_{\mathrm{hit}} \log |G_n|$.

- Among my favorite open problems two are due to Yuval Peres: showing that the mixing time of the random walk on a transitive graph G of degree d and diameter γ is at most $O(d\gamma^2)$ and proving that for a sequence of transitive expanders the random walk exhibits cutoff. I am also interested in the problem of showing that the TV mixing time of transitive graphs are robust under small changes to the geometry. This appears to be related to the famous open problem concerning stability of the Liouville property for transitive graphs.

[10] A connected component is a maximal set of vertices that have a path between them in the graph.

6 Recurrence of Markov chain traces

An irreducible Markov chain is called **transient** (resp. **recurrent**) if it visits the origin only finitely many times with probability 1 (resp. 0). Here 'the origin' is an arbitrary state. With Benjamini [11] we showed that for every transient irreducible Markov chain on a countable state space which admits a stationary measure, its trace is almost surely recurrent for the random walk. The case that the Markov chain is reversible is due to Benjamini et al. [9]. We exploit recent results in potential theory of non-reversible Markov chains in order to extend their result to the non-reversible setup. Consequently, transient graphs for random walk do not admit a nearest neighbor transient Markov chain (not necessarily a reversible one), that crosses all edges with positive probability. In particular, this is the case for the d-dimensional grid $\mathbb{Z}^d$ for $d > 2$. Conversely, we construct such a chain for the square grid $\mathbb{Z}^2$.

7 Comparison of random walk with the non-backtracking random walk

A non-backtracking random walk **NBRW** moves at each step to a random neighbor chosen uniformly at random from the collection of neighbors it did not visit in the previous step (x and y are neighbors if there is an edge connecting them). More generally, one can define a k-NBRW as a walk that attempts to avoid (if possible) the last k edges it crossed or vertices it visited. This is a natural model for a stochastic process with limited memory. A major difficulty in analyzing this model is that (even for $k = 1$) it is non-reversible. To overcome this difficulty, we show that (under some mild assumptions) viewed at some "nice" stopping times the k-NBRW is a reversible Markov chain which can be compared with the random walk via standard comparison techniques. In particular, we show [20] that for regular graphs the random walk is recurrent iff the NBRW is. Moreover, this remains valid as long as the graph is of bounded degree and there is some $R > 0$ such that every ball at radius R in the graph contains a cycle. We prove an analog of this for general k under similar assumptions. We also show that for Cayley graphs of infinite finitely generated Abelian groups no assumptions are needed (i.e., this holds for all k).

8 Social networks and spread of infection

In a sequence of papers [10, 21, 26, 8] with various collaborators (Benjamini, Fontes, Machado, Morris and Sly) we studied the intimately related *random walks* ***social network*** (**SN**) *model* and the ***frog*** *model*, a model for a spread of a rumor/epidemic. The former was studied extensively in the infinite setup [5, 37], and is part of a larger class of models [41]. In both models there are Poisson(λ) particles per site performing independent random walks. In the former when two particles collide they are declared to be acquainted and we define the **social connectivity time** SC as the first time at which there is a path of acquaintances between every pair of particles. In [26] we show that for infinite transitive graphs the occurrence of a phase transition for the SN model w.r.t. whether any pair of particles will eventually have a path of acquaintances between them is equivalent to non-amenability.

In [10] we show that for every connected n-vertex graph of average degree d, the social connectivity time is polylogarithmic in the graph's size: $c \log n \le \mathrm{SC} \le C_d \log^3 n$ (for $\lambda = 1$) with probability tending to 1 as n diverges. For regular graphs, c in the lower bound can be improved to $1 - o(1)$.

In the fog model all particles are initially inactivated, apart from the ones at the origin, where we plant an additional particle, to ensure the process does not die instantly. The particles first "pick" an infinite trajectory independently, and then upon activation perform $t \in \mathbb{R} \cup \{\infty\}$ steps of it, after which they vanish. Active particles activate all inactive particles they meet during their length t walk. We are interested in the "**susceptibility**", the minimal t sufficient for waking up all particles, and in the **cover time**, the time until this happens (when $t = \infty$). Loosely speaking, the former is the minimal lifetime t of an infected individual which results in extinction of the population, while the latter is the total time until extinction when $t = \infty$.

We obtain sharp estimates (for both models, for a wide range of λ, including $\lambda \ll 1$ and $\lambda \gg 1$) in the case the base graph is either a d-dim torus or an expander. For d-dim tori of side length L with $d \ge 2$ we in fact show that the susceptibility is concentrated around the (expected) time it takes λL^d particles performing simultaneously independent random walk, each starting from a position chosen uniformly at random, to cover the entire graph (i.e. until every vertex is visited by at least one particle; we show that the last random variable is concentrated around its mean, which we determine up to smaller order terms as a function of (λ, L, d)).

For d-ary tree of depth r (again for a wide range of λ) we determine the order of the susceptibility (it is of order $\frac{r}{\lambda} \log(\frac{r}{\lambda})$ for $\lambda \le 1$) and show that the cover time is at most exponential in $\sqrt{r}$. We conjectured that the cover time exhibits a phase transition in λ. This was verified by Hoffman et al. [38] who also showed that both the $r \log r$ and the $e^{\Theta(\sqrt{r})}$ bounds are attained for certain regimes of λ.

9 Approximate counting independent sets in hypergraphs

A hypergraph $H = (V, F)$ consists of a vertex set V and a collection F of vertex subsets, called the hyperedges. An independent set of H is a set $I \subseteq V$ such that no hyperedge $a \in F$ is a subset of I. Consider the case that H is k-uniform, i.e. $|a| = k$ for all $a \in F$. Denote the maximal degree by Δ. The natural analogue to the well-understood graph setup would predict that the threshold for approximate counting should correspond to the uniqueness threshold of the Gibbs measure for the Δ-regular tree which corresponds to $\Delta \asymp \frac{1}{2k}2^k =: \Delta_{\mathrm{uniq}}(k)$. This turns out to be false and in fact [13] showed that it is NP-hard to approximately count independent sets when $\Delta \leqslant 5 \cdot 2^{k/2}$. Conversely, when $\Delta \leqslant k$ there exists a Fully polynomial-time approximation scheme (**FPTAS**) for approximating the number of independent sets[11] [13]. Two technical difficulties are that a technical condition called strong spatial mixing fails when $\Delta \geq 6$ for all $k \geq 2$ [13] and that a standard technique for the Galuber dynamics called path coupling also breaks down for linear sized Δ.

In a joint work with Allan Sly and Yumeng Zhang [34], we show that the corresponding Glauber dynamics mixes in $O(|V| \log |V|)$ steps, provided that $\Delta \leq c2^{k/2}$. Using a standard reduction from sampling via Markov chain Monte Carlo to counting this gives rise to an FPTAS for approximate counting hypergraph independent sets, reducing the exponential gap to a constant factor. Under the condition that any pair of hyperedges share at most one vertex, we improve this bound to $\Delta \leq c2^k k^{-2}$, and when each site belongs to a bounded number of "short" cycles we improve this further to $\Delta \leq c2^k/k = \Theta(\Delta_{\mathrm{uniq}}(k))$.

We show that if coupling fails, then some disagreement at time 0 must propagate to the present time, which corresponds to a vertical crossing of some auxiliary subcritical percolation process.

References

[1] David Aldous and Persi Diaconis. Shuffling cards and stopping times. *Amer. Math. Monthly*, 93(5):333–348, 1986.

[2] David Aldous and Jim Fill. *Reversible Markov chains and random walks on graphs*. Berkeley, 2002.

[3] David J. Aldous. Some inequalities for reversible Markov chains. *J. London Math. Soc. (2)*, 25(3):564–576, 1982.

[4] Noga Alon and Yuval Roichman. Random Cayley graphs and expanders. *Random Structures Algorithms*, 5(2):271–284, 1994.

[5] O. S. M. Alves, F. P. Machado, and S. Yu. Popov. The shape theorem for the frog model. *Ann. Appl. Probab.*, 12(2):533–546, 2002.

[6] Riddhipratim Basu, Jonathan Hermon, and Yuval Peres. Characterization of cutoff for reversible Markov chains. *Ann. Probab.*, 45(3):1448–1487, 2017.

[7] I. Benjamini, R. Lyons, Y. Peres, and O. Schramm. Group-invariant percolation on graphs. *Geom. Funct. Anal.*, 9(1):29–66, 1999.

[8] Itai Benjamini, L.R. Fontes, Jonathan Hermon, and F.P. Machado. On an epidemic model on finite graphs. *Arxiv preprint arXiv:1610.04301*, 2016.

[9] Itai Benjamini, Ori Gurel-Gurevich, and Russell Lyons. Recurrence of random walk traces. *Ann. Probab.*, 35(2):732–738, 2007.

[10] Itai Benjamini and Jonathan Hermon. Rapid social connectivity. *arXiv preprint arXiv:1608.07621*, 2016.

[11] Itai Benjamini and Jonathan Hermon. Markov chain traces. *arXiv preprint arXiv:1711.03479*, 2017.

[12] Itai Benjamini and Oded Schramm. Percolation beyond $\mathbb{Z}^d$, many questions and a few answers. In *Selected works of Oded Schramm. Volume 1, 2*, Sel. Works Probab. Stat., pages 679–690. Springer, New York, 2011.

[13] Ivona Bezáková, Andreas Galanis, Leslie Ann Goldberg, Heng Guo, and Daniel Štefankovič. Approximation via correlation decay when strong spatial mixing fails. In *43rd International Colloquium on Automata, Languages, and Programming*, volume 55 of *LIPIcs. Leibniz Int. Proc. Inform.*, pages Art. No. 45, 13. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2016.

[14] Pietro Caputo, Thomas M. Liggett, and Thomas Richthammer. Proof of Aldous' spectral gap conjecture. *J. Amer. Math. Soc.*, 23(3):831–851, 2010.

[15] P. Diaconis and L. Saloff-Coste. Moderate growth and random walk on finite groups. *Geom. Funct. Anal.*, 4(1):1–36, 1994.

[16] P. Diaconis and L. Saloff-Coste. Logarithmic Sobolev inequalities for finite Markov chains. *Ann. Appl. Probab.*, 6(3):695–750, 1996.

[17] Jian Ding, Eyal Lubetzky, and Yuval Peres. Total variation cutoff in birth-and-death chains. *Probab. Theory Related Fields*, 146(1-2):61–85, 2010.

[11] Shortly before posting our paper, Moitra [49] posted a result giving a new algorithm which applies up to $\Delta \leqslant 2^{k/20}$.

[18] Carl Dou and Martin Hildebrand. Enumeration and random random walks on finite groups. *Ann. Probab.*, 24(2):987–1000, 1996.

[19] Jonathan Hermon. Cutoff for Ramanujan graphs via degree inflation. *Electron. Commun. Probab.*, 22:Paper No. 45, 10, 2017.

[20] Jonathan Hermon. Reversibility of the non-backtracking random walk. *arXiv preprint arXiv:1707.01601*, 2017.

[21] Jonathan Hermon. Frogs on trees? *Electron. J. Probab.*, 23:Paper No. 17, 40, 2018.

[22] Jonathan Hermon. On sensitivity of uniform mixing times. *Ann. Inst. Henri Poincaré Probab. Stat.*, 54(1):234–248, 2018.

[23] Jonathan Hermon. A spectral characterization for concentration of the cover time. *arXiv preprint arXiv:1809.00145*, 2018.

[24] Jonathan Hermon and Tom Hutchcroft. No percolation at criticality on certain groups of intermediate growth. *arXiv preprint arXiv:1809.11112*, 2018.

[25] Jonathan Hermon, Hubert Lacoin, and Yuval Peres. Total variation and separation cutoffs are not equivalent and neither one implies the other. *Electron. J. Probab.*, 21:Paper No. 44, 36, 2016.

[26] Jonathan Hermon, Ben Morris, Chuan Qin, and Allan Sly. The social network model on infinite graphs. *Arxiv preprint arXiv:1610.04293*, 2016.

[27] Jonathan Hermon and Yuval Peres. The power of averaging at two consecutive time steps: proof of a mixing conjecture by Aldous and Fill. *Ann. Inst. Henri Poincaré Probab. Stat.*, 53(4):2030–2042, 2017.

[28] Jonathan Hermon and Yuval Peres. A characterization of L_2 mixing and hypercontractivity via hitting times and maximal inequalities. *Probab. Theory Related Fields*, 170(3-4):769–800, 2018.

[29] Jonathan Hermon and Yuval Peres. On sensitivity of mixing times and cutoff. *Electron. J. Probab.*, 23:Paper No. 25, 34, 2018.

[30] Jonathan Hermon and Richard Pymar. The exclusion process mixes (almost) faster than independent particles. *arXiv preprint arXiv:1808.10846*, 2018.

[31] Jonathan Hermon and Justin Salez. Cutoff for the mean-field zero-range process with bounded monotone rates. *arXiv preprint arXiv:1811.03520*, 2018.

[32] Jonathan Hermon and Justin Salez. A version of aldous' spectral-gap conjecture for the zero range process. *arXiv preprint arXiv:1808.00325*, 2018.

[33] Jonathan Hermon and Thomas Sam. Cutoff for mixing times on random abelian cayley graphs. *arXiv preprint arXiv:1810.05130*, 2018.

[34] Jonathan Hermon, Allan Sly, and Yumeng Zhang. Rapid mixing of hypergraph independent set. *To appear in Random Structures and algorithms. ArXiv preprint arXiv:1610.07999*, 2016.

[35] Jonathan Hermon and Perla Sousi. Random walks on dynamical percolation on general graphs: comparison of mixing and hitting with srw. *In preparation*, 2018.

[36] Martin Hildebrand. Random walks supported on random points of $\mathbf{Z}/n\mathbf{Z}$. *Probab. Theory Related Fields*, 100(2):191–203, 1994.

[37] Christopher Hoffman, Tobias Johnson, and Matthew Junge. Recurrence and transience for the frog model on trees. *Ann. Probab.*, 45(5):2826–2854, 2017.

[38] Christopher Hoffman, Tobias Johnson, and Matthew Junge. Cover time for the frog model on trees. *arXiv preprint arXiv:1802.03428*, 2018.

[39] Tom Hutchcroft. Critical percolation on any quasi-transitive graph of exponential growth has no infinite clusters. *C. R. Math. Acad. Sci. Paris*, 354(9):944–947, 2016.

[40] Tom Hutchcroft. Locality of the critical probability for transitive graphs of exponential growth. *arXiv preprint arXiv:1808.08940*, 2018.

[41] Harry Kesten and Vladas Sidoravicius. A shape theorem for the spread of an infection. *Ann. of Math. (2)*, 167(3):701–766, 2008.

[42] Gady Kozma. On the precision of the spectral profile. *ALEA Lat. Am. J. Probab. Math. Stat.*, 3:321–329, 2007.

[43] Hubert Lacoin. The cutoff profile for the simple exclusion process on the circle. *Ann. Probab.*, 44(5):3399–3430, 2016. MR3551201.

[44] Thomas M. Liggett. *Stochastic interacting systems: contact, voter and exclusion processes*, volume 324 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 1999. MR1717346.

[45] Thomas M. Liggett. *Interacting particle systems*. Classics in Mathematics. Springer-Verlag, Berlin, 2005. Reprint of the 1985 original.

[46] Eyal Lubetzky and Yuval Peres. Cutoff on all Ramanujan graphs. *Geom. Funct. Anal.*, 26(4):1190–1216, 2016.

[47] A. Lubotzky, R. Phillips, and P. Sarnak. Ramanujan graphs. *Combinatorica*, 8(3):261–277, 1988.

[48] Mathieu Merle and Justin Salez. Cutoff for the mean-field zero-range process. *arXiv preprint arXiv:1804.04608*, 2018.

[49] Ankur Moitra. Approximate counting, the Lovász local lemma, and inference in graphical models. In *STOC'17—Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 356–369. ACM, New York, 2017.
[50] A. Nilli. On the second eigenvalue of a graph. *Discrete Math.*, 91(2):207–210, 1991.
[51] Roberto Imbuzeiro Oliveira. Mixing and hitting times for finite Markov chains. *Electron. J. Probab.*, 17:no. 70, 12, 2012.
[52] Yuval Peres and Perla Sousi. Mixing times are hitting times of large sets. *J. Theoret. Probab.*, 28(2):488–519, 2015.
[53] Yuval Peres, Alexandre Stauffer, and Jeffrey E Steif. Random walks on dynamical percolation: mixing times, mean squared displacement and hitting times. *Probability Theory and Related Fields*, 162(3-4):487–530, 2015.
[54] Ch. Pittet and L. Saloff-Coste. On the stability of the behavior of random walks on groups. *J. Geom. Anal.*, 10(4):713–737, 2000.
[55] Alistair Sinclair. *Algorithms for random generation and counting*. Progress in Theoretical Computer Science. Birkhäuser Boston, Inc., Boston, MA, 1993. A Markov chain approach.
[56] Frank Spitzer. Interaction of Markov processes. *Advances in Math.*, 5:246–290 (1970), 1970. MR0268959.
[57] David Bruce Wilson. Random random walks on $\mathbf{Z}_2^d$. *Probab. Theory Related Fields*, 108(4):441–457, 1997.
[58] David Bruce Wilson. Mixing times of Lozenge tiling and card shuffling Markov chains. *Ann. Appl. Probab.*, 14(1):274–325, 2004. MR2023023.