

Contents

1	Introduction	1
1.1	Random Interchange	1
1.2	CTRW	2
1.3	The problem	2
1.4	Why is it a nice problem?	2
1.5	Known special cases	3
1.5.1	Trees	3
1.5.2	Complete graphs and some other graphs of diverging degree	3
2	Viewing the IP as SRW on its Cayley graph	3
2.1	On the natural time parameterisation of this SRW	4
2.2	Why should we view the IP as SRW?	4
2.3	What is Aldous' conjecture?	4
2.4	What does a (uniform) spectral gap tell us about SRWs?	5
2.4.1	On the connection to expander graphs	5
2.4.2	Results for expander graphs	6
2.4.3	More on expanders	8
2.5	Using the fact that the Laplacian is dissipative	9
2.5.1	On the geometry of finite loop permutations in Γ	11
3	Schramm's coupling argument	11
4	Representation theory & Fourier analysis	12

1 Introduction

1.1 Random Interchange

Consider the random interchange process (IP) on a finite weighted graph $G = (V, E, C)$ with vertex set $V, |V| = n$, edge set E and weights $c_{xy} = c_{yx} \geq 0$ defined for $xy = \{x, y\} \in E$. The IP is a permutation ($/S_n$)-valued process with infinitesimal generator

$$\mathcal{L}^{\text{IP}} f(\eta) := \sum_{xy \in E} c_{xy} (f(\eta^{xy}) - f(\eta)), \quad \text{for } f : S_n \rightarrow \mathbb{R}, \eta \in S_n \quad (1.1) \quad \{\text{eq:infGenIP}\}$$

where $\eta^{xy} = \eta\tau_{xy}$ (τ_{xy} being the transposition of x, y). Intuitively we start with the identity permutation and put an exponential clock with parameter c_{xy} on each edge; whenever it rings we multiply by τ_{xy} .

1.2 CTRW

For a definition and a nice introduction to the CTRW see [Ang03]. Of particular importance for this problem is [Ang03, Proposition 7]:

{pro:TC}

Proposition 1.1. *In the natural coupling between the $CTRW_T$ and the η_T (the IP at time T), the $CTRW_T$ started at $x \in V$ is recurrent if and only if the orbit of η_T containing x is finite.*

1.3 The problem

We want to prove [Ang03, Conjecture 6]:

{conj:main}

Conjecture 1.2. *The function $T \mapsto \mathbf{1}\{CTRW_T \text{ is transient}\}$ is non-decreasing.*

With Proposition 1.1 this is equivalent the following conjecture:

{conj:IP}

Conjecture 1.3. *Let $T_c := \sup\{T : \text{IP at time } T \text{ has no infinite cycles a.s.}\}$ (note that T_c may be ∞). Then $T > T_c$ implies that*

$$\mathbb{P}(\text{IP at time } T \text{ has an infinite cycle}) > 0.$$

Denote by $p_{\text{esc}}(x, A; T)$ the probability that a $CTRW_T$ started at x reaches any point $a \in A \subseteq V$. Let $B_n(x) := \{y \in V : d_G(x, y) \leq n\}$ with the natural graph metric d_G . Noting that $CTRW_T$ being transient is equivalent to $p_{\text{esc}}(x, \partial B_n(x))$ being uniformly positive in n means that the following conjecture implies Conjecture 1.2:

{conj:Pesc}

Conjecture 1.4. *The function $T \mapsto p_{\text{esc}}(x, \partial B_n(x); T)$ is non-decreasing $\forall n \in \mathbb{N}$.*

1.4 Why is it a nice problem?

We know that $CTRW_\infty$ is just the SRW, so if Conjecture 1.2 was true, we would have proved that on recurrent graphs* G the IP does not have infinite cycles or, equivalently, that the $CTRW_T$ is recurrent for all $T > 0$. In particular there is no phase transition on $\mathbb{Z}^2$; this is still an open problem.

*I.e. those graphs where the SRW is recurrent.

1.5 Known special cases

1.5.1 Trees

In the special case of rooted d -regular trees with $d \geq 3$ Conjecture 1.2 was proved (indirectly) in [Ham12, Ham15]. Unfortunately the proof is unlikely to generalise since it only follows from the stronger* statement that there is a $T_c < \infty$ such that CTRW_T is recurrent for all $T < T_c$ and transient for all $T > T_c$. The proof makes heavy use of the fact that we are on trees, i.e. G has no cycles. To the best of our knowledge, this is the only “finite dimensional” setting where a phase transition is proved to occur.

1.5.2 Complete graphs and some other graphs of diverging degree

On the complete graph [Sch05] proved that once there are mesoscopic cycles in the IP, one rapidly converges to the invariant distribution (which is Poisson Dirichlet). This implies that once we have giant cycles in the IP, there is no way we go back to having only finite cycles again after some fixed finite time and hence implies Conjecture 1.2. For more on this approach, see Section 3.

2 Viewing the IP as SRW on its Cayley graph

{sec:CayleySRW}

Consider the random interchange process on a fixed weighted graph $G = (V, E, C)$ with $|V| = N$ and $c_{xy} = 1$ whenever $xy \in E$. Since the transpositions τ_{xy} that are added are independent of the past (memoryless-ness of the exponential distribution + independence of exponential clocks on different edges) the IP is a Markov process. From the way we chose the weights, it is not hard to see that the IP is just a simple random walk on its Cayley graph $\Gamma = (S_N, A)$. Here $A = \{\tau_e : e \in E\}$ is the set of generators, i.e. transpositions “supported on the edges of G ”. Due to the exponential delay between jumps the IP, viewed as SRW on Γ , is aperiodic and since Γ is connected (clearly $\langle A \rangle = S_N$) it is also irreducible. As a Cayley graph Γ is furthermore a $|A|$ -regular graph with $n := |S_N| = N!$ vertices. Putting these

*This statement is stronger since it establishes that a phase transition occurs at some *finite* T_c . Proving existence of infinite loops is likely a harder problem than proving that once they exist with positive probability, they cannot “go away” again.

facts together we see that the IP has a unique stationary distribution, given by the uniform measure on S_N .

2.1 On the natural time parameterisation of this SRW

{subsec:timeparam

A natural question to ask is which time scales we are interested in. Note that for the complete graph on N vertices we have a phase transition at $t_c = N$, on d -regular trees we have the phase transition as soon as there are $\sim 1/d$ crosses per edge of the graph, so again $t_c = N$, etc. Typically we expect phase transitions to occur at times such that the number of crosses (corresponding to transpositions) is of the same order as the number of vertices of G , i.e. $t \approx O(N)$.

For the SRW on the Cayley graph Γ , however, a time parameterisation in terms of N is not natural. A more natural parameterisation is given in terms of the number of the vertices of Γ , i.e. $n \stackrel{\text{(def)}}{=} |S_N|$:

$$t \approx O(N) = O(\log n / \log \log n) \ll \log n,$$

where we used that $\log n \stackrel{\text{(def)}}{=} \log N! = O(N \log N)$.

2.2 Why should we view the IP as SRW?

The only theorem I know that establishes a connection between the geometry of the underlying graph G and the induced Cayley graph Γ is (the recently resolved, see [CLR10]*) Aldous' conjecture and some mild(?) extension: [AK10][†]. Both results relate the eigenvalues of the Laplacians of the SRW on G to the SRW on Γ . Most importantly, [CLR10] establishes that the spectral gaps of the two SRWs coincide.

Before we look at what a (uniform) spectral gap might tell us about the SRW on Γ , we give a brief explanation of what Aldous' conjecture says:

2.3 What is Aldous' conjecture?

Keeping track only of the trajectory of one fixed label $\in V$ instead of the entire permutation η , we obtain a random walk with conductances c_{xy} . More precisely we

*See also [Ces10] for a weaker result (only on multipartite graphs), but with a good introduction and some representation theory.

[†]Comparison of other eigenvalues might be important since control on spectral gap usually gives exponential convergence to equilibrium, but we want more: Namely a phase transition. For this more precise information on the spectrum will be needed.

can see the random walk with infinitesimal generator

$$\mathcal{L}^{\text{RW}} f(x) := \sum_{y \neq x} c_{xy} (f(y) - f(x)), \quad \text{for } f : V \rightarrow \mathbb{R}, x \in V \quad (2.1) \quad \{\text{eq:infGenRW}\}$$

as sub-process of the interchange process. In particular we immediately get that $\text{spec}(-\mathcal{L}^{\text{RW}}) \subseteq \text{spec}(-\mathcal{L}^{\text{IP}})$ and hence

$$\lambda_1^{\text{IP}} \leq \lambda_1^{\text{RW}} \quad (2.2) \quad \{\text{eq:EvalIneq}\}$$

Aldous’ conjecture (proved and explained much better in [CLR10]) states that (2.2) is actually an equality.

2.4 What does a (uniform) spectral gap tell us about SRWs?

So assume the family G_N has a positive spectral gap, uniformly in N . By [CLR10] the same holds for the SRW on Γ .

Generally speaking, a uniformly bounded spectral gap ensures “fast mixing times” (i.e. exponential rate of convergence to the equilibrium distribution) and good isoperimetric inequalities (see Cheeger’s inequalities). This does not seem to be enough for our purposes.

Representation theory, i.e. analysis of invariant subspaces and more precise study of the spectrum of the Laplacian, will provide a better understanding, see Section 4. However, this does not seem feasible unless the underlying family of graphs G_N have a lot of symmetries. So for the rest of this subsection we will try to restrict ourselves to methods as robust as possible, i.e. only using the fact that we have a positive spectral gap.

2.4.1 On the connection to expander graphs

Note that for the SRW induced by the IP we cannot be in the setting of expander graphs. While G_N may be a family of expander graphs, $\Gamma_N = \Gamma(S_N, \{\tau_e : e \in E_N\})$ is certainly no expander family, since the vertex degree is $|E_N|$ which is not uniformly bounded in N . It is worth noting, however, that whenever G_N is a family of expander graphs, we have that $|E_N| = dN = O(\log n / \log \log n) = o(\log n)$, where $n = n(N)$ is the number of vertices in the Cayley graph Γ_N . So the number of edges of our Cayley graphs Γ_N are in some sense “almost” bounded, i.e. compared to the number

of vertices n they diverge very slowly (only $o(\log n)!$). Hence one might be optimistic and try to see whether some results from expander graphs carry over to this setting after all.

2.4.2 Results for expander graphs

For a proper introduction on expander graphs in general, see [Kow]. For a briefer sketch of some selected results, but with very good comments, see [these notes](#).

- [Gil98] proves one of the most famous results on mixing on expanders:

{thm:expanderWalk

Theorem 2.1 (Expander walk sampling theorem). *Let $G = (V, E)$ be a graph with $|V| = n$ and spectral gap $\gamma > 0$. Let $f : V \rightarrow [0, 1]$. Denote by $Y_1, \dots, Y_k \in V$ the (random) vertices encountered in a k -step SRW on G , where q s.t. $\mathbb{P}(Y_1 = x) = q(x)$. Then we have*

$$\mathbb{P} \left(\left| \frac{1}{k} \sum_{i=1}^k f(Y_i) - \mathbb{E}_\pi(f) \right| \geq \delta \right) \leq C_q e^{-k\delta^2\gamma/20\nu}, \quad (2.3)$$

where π is the uniform distribution on V , $\nu := \max_{x,y \in V} \pi(x)/\pi(y)$ and $C_q := 4\|q/\sqrt{\pi}\|_2$.

For regular graphs, in particular Cayley graphs, $\nu = 1$. Now if we let our SRW Y on a regular graph start at a fixed vertex, then the factor $C_q = 4n$; on the other hand, if we choose $q = \pi$, then $C_q = 4$. For our purposes we will thus have to work with

$$\mathbb{P} \left(\left| \frac{1}{k} \sum_{i=1}^k f(Y_i) - \mathbb{E}_\pi(f) \right| \geq \delta \right) \leq 4e^{-k\delta^2\gamma/20+\log n}.$$

Choosing $k \gg \frac{\log n}{\delta^2\gamma}$ we then get that the empirical average of f over $(Y_i)_{i=1}^k$ converges to the mean of f . Note, however, that our time scales go up to $t = o(\log n)$, so this theorem does not seem to be particularly helpful. Reflecting a little bit more reveals that we should not expect results like that to hold in our setting: After all the point of these phase transitions (of occurrence of giant cycles) is that they occur *before* we are close to the equilibrium distribution! (This really should be independent of the geometry of the graph, so even if G_N is a family of expanders, this doesn't help us here.)

- [BY16] proves that, starting at a uniformly chosen random vertex, the probability of hitting some set A is roughly (up to some logarithmic correction) proportional to $\pi(A) = |A|/|V|$. One might extend this result to include an error introduced by starting from some distribution q that is not the uniform distribution π – compare with [Gil98] where this introduces a factor of $\|q/\sqrt{\pi}\|_2$, hence a factor of n if q is concentrated at a point. For reasons as above, however, it is not likely to hold in our setting since there is a strong link to the equilibrium distribution.*
- MIT lecture notes prove the following lemma:

Lemma 2.2. *For any regular graph $G = (V, E)$ on $N = |V|$ vertices with spectral gap $\gamma > 0$ we have the following inequality:*

$$\text{diam}(G) \leq \frac{\ln N}{\gamma}.$$

The proof is simple enough to be given here:

Proof. Let u, v be vertices that as far apart from each other as possible, i.e. $d(u, v) = \text{diam}(G)$. Let X be a SRW on G started at u . Clearly $p_t(v) > 0$ implies $t \leq \text{diam}(G)$. Since G is regular the invariant distribution π takes values $1/N$ on every vertex. Hence $p_t(v) > 0$ is equivalent to $|p_t(v) - \pi(v)| < 1/N$. Standard spectral mixing time arguments give $|p_t(v) - \pi(v)| < (1 - \gamma)^t \sqrt{\frac{d(v)}{\min_y d(y)}} = (1 - \gamma)^t$, since G is regular. Noting that $(1 - \gamma)^{1/\gamma} < 1/e$ for all $\gamma > 0$ we get for $t = \ln N/\gamma$ that $(1 - \gamma)^t < 1/N$ which proves the claim. $\square$

From the proof we see that this is unlikely to give a very sharp bound, but asymptotically in N (not in γ !†) this is about right, up to constants.

Applying this lemma to $\Gamma_N = \Gamma(S_N, \{\tau_e : e \in E_N\})$ with the underlying graphs G_N being a family of expander graphs, we know from [CLR10] that the spectral gap γ of Γ_N is the same as for G_N (in particular it is uniformly bounded from below) and thus the diameter of Γ_N grows at most like $\ln n = O(N \log N)$.

*In particular the harmonic measure is obtained by letting a SRW run until it hits its target, this may be much longer than the time scales we want to consider.

†In fact they prove an even stronger statement, namely that $\text{diam}(G) \leq (1 + \frac{1}{\sqrt{2\gamma}}) \ln(2N)$, but for asymptotically they only differ by a factor of $\sqrt{\gamma}$.

Recall from Subsection 2.1 that for our problem we consider the SRW on Γ_N up to times $t = O(N) = o(\log n)$. Assuming this bound on the diameter captures the right asymptotic behaviour of the diameter^{*} (neglecting constants and focusing only on the logarithmic growth), this means that, even if the underlying graphs admit very fast mixing, there are some permutations (those with word-distance $\gg \log n / \log \log n^\dagger$) that our SRW on Γ_N will not visit with probability 1. In general (i.e. if the underlying family of graphs do not have a spectral gap uniformly bounded from below) we should hence expect to find even more permutations well outside the typical range of the SRW on Γ_N with our natural time scaling.

In the light of Theorem 2.1 it is also worth noting just how strong our claims about phase transitions of occurrence of giant cycles for the IP are. In the setting as in Conjecture 2.3, Theorem 2.1 with $f : S_n \rightarrow [0, 1]$ s.t. $\pi \mapsto \|\mathcal{X}(\pi)\|_\infty$ (or $\pi \mapsto \|\mathcal{X}(\pi)\|_2$) tells us that for a RW on the Cayley graph of length $k \gg N \log N \sim \log n^\ddagger$ we know that the average (over the trajectory of the that random walk on S_n) maximal cycle length is just the expected maximal cycle length under the uniform distribution on all permutations. A phase transition at time $t = O(N) = O(\log n / \log \log n)$ where we get arbitrarily close in distribution to the equilibrium distribution if we look at $t = cN$ with c sufficiently large means that we can choose $c = \log \log \log n$ to get the same result as before for $k = O(\log n \log \log \log n / \log \log n) \ll \log n$.

2.4.3 More on expanders

For more information on expanders, see Tao’s lecture notes (1), (2), (3) on expanders and various equivalent/implied properties. These notes are another good source with great comments on how to interpret some central results. Note furthermore that SRWs on expander graphs are transient (in the appropriate sense; of course expanders are typically finite). In fact we have more generally that sub-quadratic

^{*}It seems hard to find reasonably sharp lower bounds on the diameter in terms of the spectrum: <https://mathoverflow.net/q/168813/47059>.

[†]And in fact even more since a SRW doesn’t just run off to one direction, but fluctuates back and forth a bit, which should cost an additional “factor” of a square root.

[‡]Recall that N denotes the number of vertices of the underlying graph $G = (V, E)$ and $n = N!$ is the cardinality of the state space of the induced IP. The natural time scaling for the IP is expressed in terms n , but all known results, e.g. [Sch05], etc. state results in terms of N .

[growth implies recurrence](#) (this link has many other nice insights too!). Very generally [this survey on RWs](#) might be useful to relate quantities such as mixing times, spectral gap, hitting times, etc.

2.5 Using the fact that the Laplacian is dissipative

A much more elementary observation than the ones above is that the law of the IP at time t' is more “spread out” than at time $t < t'$. This will not be made more formal for now, so just recall that the Laplacian is not only the infinitesimal generator of a random walk, but also (up to normalisation) the Hamiltonian of the heat equation on general graphs and “heat spreads out in time”. Note that the fact that the Laplacian of the SRW on the Cayley graph Γ of the IP is dissipative is shown in the introduction of [CLR10].

In particular this allows us to reduce our main problem (see Conjecture 2.3 for a convenient formulation for this setting) to determining whether the density of permutations with only finite cycles decreases as the word-distance to the identity permutation increases. More precisely, to prove Conjecture 1.2 it is sufficient to show

{conj:Cayley}

Conjecture 2.3. *Consider a fixed “underlying” graph $G = (V, E)$ and its induced set of generators $A = \{\tau_e\}_{e \in E}$; clearly $\langle A \rangle = S_N$, where $N = |V|$. Let B_t be the t -ball (w.r.t. to the word-metric on the Cayley graph (S_N, A)) around id , the identity permutation and denote by ν_t the uniform measure on ∂B_t . Then*

$$t \leq t' \implies \nu_t(\|\mathcal{X}\|_{L^2}) \leq \nu_{t'}(\|\mathcal{X}\|_{L^2}),$$

where $\mathcal{X} = \mathcal{X}(\pi)$ is the vector of normalised cycle lengths.

Considering the time scales we are interested in, it suffices to restrict oneself to $t' < T$ with T arbitrarily small such that $\log n / (\log \log n) \ll T \ll \log n$ holds.

Remark 2.4 (On the choice of the observable in Conjecture 2.3). *One has a bit of wiggle room for choosing the observable. Instead of $\|\mathcal{X}\|_{L^2}$ one could have chosen $\|\mathcal{X}\|_{L^\infty}$, the size of the largest cycle (normalised by N). In this case we would expect $\nu_t(\|\mathcal{X}\|_{L^\infty}) = o(1)$ (or more precisely $O(n^{-1})$, but this would require the additional effort of ruling out the existence of mesoscopic cycles) for $t \ll \log n / \log \log n$. For $t \gg \log n / \log \log n$, however, we expect $\nu_t(\|\mathcal{X}\|_{L^\infty}) = O(1)$ whenever the underlying graph G has sufficiently high vertex degree; in particular if it is non-amenable.*

Remark 2.5 (Conjecture 2.3 is true on the complete graph). *In fact we have a slightly stronger statement, i.e. we can replace $\|\mathcal{X}\|_{L^2}$ by any functional that is non-decreasing in $\mathcal{X}$, where we equip the space of partitions with its natural partial order. In particular this includes $\|\mathcal{X}\|_{L^p}$ for all $p \in (1, \infty]$.*

Proof. The key idea is that with the complete graph being the underlying graph G our set A of generators is so large that going from ∂B_t to ∂B_{t+1} does not allow us to split cycles since this would decrease the word distance to the identity. More precisely: Every permutation $\sigma \in S_N$ can be decomposed uniquely in its cycles $\{(x_{1,1}, x_{1,2}, \dots, x_{1,\ell_1}), \dots, (x_{m,1}, \dots, x_{m,\ell_m})\}$ (here m is the number of cycles where the i -th cycle has length ℓ_i ; hence $\sum_{i=1}^m \ell_i = N$), which, if the generators are all transpositions (since the underlying graph being the complete graph) can minimally* be written as a sequence of transpositions on different edges, i.e. $\sigma = \tau_{x_{1,1}, x_{1,2}} \circ \tau_{x_{1,2}, x_{1,3}} \circ \dots \circ \tau_{x_{m,\ell_m-1}, x_{m,\ell_m}} = \prod_{i=1}^m \prod_{j=2}^{\ell_i} \tau_{x_{i,j-1}, x_{i,j}}$. $\square$

Unfortunately going from ∂B_t to ∂B_{t+1} typically does include splitting cycles. Consider, for example, as underlying graph $\{1, 2, 3\}$ with edges $\{1, 2\}, \{2, 3\}$ and the permutation $\sigma = (13)$.[†] The word distance of (13) to the identity with generators $A = \{\tau_{1,2}, \tau_{2,3}\}$ is 3 with the minimal representation being $(13) \stackrel{(\text{def})}{=} \tau_{1,3} \stackrel{!}{=} \tau_{1,2}\tau_{2,3}\tau_{1,2}$. In the first step we create a cycle of length 2, in the second there is only one cycle of length 3 and in the third step that cycle is split again into one of length 2 and one of length 1.

Note that sometimes merges (i.e. adding transpositions that make merge cycles) are not allowed either when going from ∂B_t to ∂B_{t+1} : Take the previous example and add the edge $\{3, 4\}$. Clearly $\sigma = (13)$ still has the same word distance from the identity, but adding $\tau_{2,3}$ to $\tau_{1,2}\tau_{2,3}\tau_{1,2}$ is not allowed since $\tau_{2,3}\tau_{1,2}\tau_{2,3}\tau_{1,2} = (\tau_{2,3}\tau_{1,2})^2 = (\tau_{2,3}\tau_{1,2})^{-1} = \tau_{1,2}\tau_{2,3}$; i.e. adding $\tau_{2,3}$ to $\tau_{1,2}\tau_{2,3}\tau_{1,2}$ would mean that we go from ∂B_3 to ∂B_2 .

As a matter of fact we need to control more than just the rate of splits and merges since it could be that a majority of splits are “large splits” (splits that split the largest cycle into half), while the majority of merges are only increasing the largest cycle’s size by 1. In such a scenario controlling the rates of splits and merges does not suffice. Instead one has to keep track of all the other cycles and would like to have dynamics similar to those in [Sch05]; they will be discussed in Section 3.

*I.e. using the minimal number of transpositions; this is not hard to see.

[†]In this notation τ_{ij} and (ij) are the same object. We use both notions here.

2.5.1 On the geometry of finite loop permutations in Γ

Instead of considering a fixed path and proving that more often than not it moves from an permutation with only finite cycles to one with infinite cycles, one might try to prove certain geometrical properties of sets of permutations with properties relating to cycle sizes.

Denote the set of vertices of $\Gamma = \Gamma(S_n, A)$ (permutations in S_n) with at least one cycle of length bigger than εn by A_ε . Roughly speaking estimates on the asymptotic size of A_ε , i.e. deciding whether $|A_\varepsilon|/|S_n|$ is bigger than some constant that is independent of n or not corresponds to answering the question of whether there are macroscopic loops.

3 Schramm's coupling argument

{sec:coupling}

Schramm proved the following:

Theorem 3.1 ([Sch05], Theorem 1.1). *Let $\varepsilon > 0$, $t \geq (1/2 + \varepsilon)N$ and Y be a PD(1) sample. Denote by K the size of the largest component of the induced Bernoulli percolation and by $\mathcal{X} = \mathcal{X}(\sigma)$ the cycle structure of σ , i.e. the sorted (non-increasing order) list of the lengths of the cycles. Denote furthermore by σ_t the permutation of the IP at time t . Then $\mathcal{X}(\sigma_t)/K$ converges weakly in the sup-norm to Y , as $N \rightarrow \infty$.*

To summarise: At time $N/2$ there is a phase transition. For $t \ll N/2$ there are no giant cycles, for $t \gg N/2$ giant cycles appeared and, if appropriately rescaled, they immediately converge to their equilibrium distribution: The Poisson-Dirichlet distribution.

In particular we obtain that once giant cycles exist with positive probability, they will not disappear (a.s.) at some fixed later time; hence this provides yet another proof of Conjecture 1.3 in the special case where the underlying graph is the complete graph.

Since then there have been efforts to generalise this technique to other graphs, such as the Hamming graph, see [AKM18]. The key idea is to show that once there are mesoscopic cycles, they are well-mixed, so that they essentially behave in a mean-field manner and hence techniques developed in [Sch05] apply.

Note, however, that this shows much more than we need to prove Conjecture 1.3! Here we are not interested in the exact cycle length distribution, but merely

in whether giant cycles can disappear again. In particular we don't need to know *when* giant cycles appear and we don't need to prove convergence to PD. Hence it is tempting to try to only copy the split-merge dynamics argument from [Sch05].

For general graphs it still does not seem to be reasonable to prove that large cycles (if they appear) are well-mixed, so we will not be able to work with the nice mean-field dynamics. One might try to look for general, weak isoperimetric bounds on large cycles though and see how that translates to changes of u, v (Schramm's notation): In the mean-field they are iid uniform random variables in the unit interval; the weaker the bounds, the more correlated (which might be very bad) they should be. See how far we can take it without making the associated split-merge transition kernel M degenerate in the sense that $M^k \mathcal{X}(\sigma) = (\varepsilon, \varepsilon', \dots)$ for $\varepsilon, \varepsilon', \dots$ going to zero uniformly in k .

Another problem to keep in mind is that the process induced by M is not Markovian unless we are on the complete graph. This may be accounted for by enlarging the probability space (just like in QM) to make M act on superpositions of all the possible configurations.

4 Representation theory & Fourier analysis

{sec:reprTheory}

In this section we will adopt a more group theoretic point of view. Similarly to Section 2 we shall denote the set of transpositions of adjacent vertices $\{\tau_{xy} : xy \in E\}$ by A . Equipped with the natural action we have for every connected graph $G = (V, E)$ that $\langle A \rangle = S_N$, where $N = |V|$.

We borrow some notation from additive combinatorics and write $A + B$ for^{*} $\{ab : a \in A, b \in B\}$ (and $A - B$ for $\{ab^{-1} : a \in A, b \in B\}$), whenever $A, B \subseteq S_N$. Similarly, we write mA for $A + A + \dots + A$ (m times) for $m \in \mathbb{N}$.

Recalling Subsection 2.1 it might appear as if we were interested in the set tA , for some $t = O(N)$. This could be analysed using tools from additive combinatorics and more generally group theory. While additive combinatorics[†] seems to address the issue of getting bounds on tA , group theory[‡] seems likely to provide valuable

^{*}In fact the $+$ is a bit misleading since our group operation is not commutative.

[†]See e.g. [Tao] for an accessible exposition in the commutative setting and <https://arxiv.org/pdf/math/0601431.pdf> for product set estimates in the non-commutative setting.

[‡]See e.g. the concept of quasirandom groups, i.a. explored in [Gow08] which are intimately connected to “how abelian” a group is [LT].

insights into how the geometry of G translates to good/bad mixing properties (i.e. how large/small tA is for large t).

While one might want to keep techniques from these fields in the back of one's head, it is important to notice that, a priori, the set tA is not interesting at all, since it does not take too long until $tA = S_N$ — potentially the SRW on the Cayley graph could reach every element relatively fast, but due to cancellations, etc. (just think of it as the SRW on its Cayley graph) this is not what we “typically” observe.

In fact what we are interested in is a quantitative version of “ tA ”, i.e. how often does $s \in tA$ occur when we concatenate t random elements $s_1 s_2 \dots s_t$ with $s_i \in A$, relative to other elements. Clearly this question can conveniently be rephrased in terms of (probability) measures; we put the uniform measure on A and consider its t -fold convolution with itself. At this point it should become obvious that Fourier analysis on groups is an appropriate tool to address this problem.*

One way to attack this problem is already outlined in Section 2 — clearly the t -fold convolution of the uniform measure on some set A gives the law of the SRW after t steps on $\langle A \rangle$. Hence results like [CLR10] can be used. Note, however, that having control over the spectral gap does usually only give an exponential rate of convergence to the equilibrium distribution (in that case this is the uniform distribution on $\langle A \rangle = S_N$), but no phase transition, etc. For this we will have to understand $\text{spec}(-\mathcal{L}^{\text{IP}})$ much better.

Instead of working explicitly with the Laplacian that generates the IP, however, one might as well try to use the existing theory on Fourier analysis on non-commutative groups. This approach seems to have succeeded in “easy” cases like the complete graph ($G = K_N$), see e.g. [AK10, AK13, BK12].[†] It remains to see whether these methods are robust enough to carry over to “hard” graphs (basically anything other than $G = K_N$), since the spectrum of the infinitesimal generator of the IP on the complete graph already seems to be known “quite explicitly”, i.e. in terms of other quantities from combinatorics/representation theory, see [DS81], so it is reasonable to expect that one should be able to say a lot using Fourier analysis/representation theory in that case.

*More experienced probabilists may (rightfully) insist that this “motivation” is superfluous after already having interpreted the IP as SRW on its Cayley graph.

[†]Note that [CLR10] also has some interesting remarks about simple consequences from group/representation theory in the introduction.

References

- [AK10] Gil Alon and Gady Kozma. Ordering the representations of s_n using the interchange process. *arXiv preprint arXiv:1003.1710*, 2010.
- [AK13] Gil Alon and Gady Kozma. The probability of long cycles in interchange processes. *Duke Mathematical Journal*, 162(9):1567–1585, 2013.
- [AKM18] Radosław Adamczak, Michał Kotowski, and Piotr Miłoś. Phase transition for the interchange and quantum heisenberg models on the hamming graph. *arXiv preprint arXiv:1808.08902*, 2018.
- [Ang03] Omer Angel. Random infinite permutations and the cyclic time random walk. In *DRW*, pages 9–16, 2003.
- [BK12] Nathanaël Berestycki and Gady Kozma. Cycle structure of the interchange process and representation theory. *arXiv preprint arXiv:1205.4753*, 2012.
- [BY16] Itai Benjamini and Ariel Yadin. Harmonic measure in the presence of a spectral gap. *Annales de l’Institut Henri Poincaré*, 52(3):1050–1060, 2016.
- [Ces10] Filippo Cesi. On the eigenvalues of cayley graphs on the symmetric group generated by a complete multipartite set of transpositions. *Journal of Algebraic Combinatorics*, 32(2):155–185, 2010.
- [CLR10] Pietro Caputo, Thomas Liggett, and Thomas Richthammer. Proof of aldous spectral gap conjecture. *Journal of the American Mathematical Society*, 23(3):831–851, 2010.
- [DS81] Persi Diaconis and Mehrdad Shashahani. Generating a random permutation with random transpositions. *Probability Theory and Related Fields*, 57(2):159–179, 1981.
- [Gil98] David Gillman. A chernoff bound for random walks on expander graphs. *SIAM Journal on Computing*, 27(4):1203–1220, 1998.
- [Gow08] William Timothy Gowers. Quasirandom groups. *Combinatorics, Probability and Computing*, 17(3):363–387, 2008.
- [Ham12] Alan Hammond. Infinite cycles in the random stirring model on trees. *arXiv preprint arXiv:1202.1319*, 2012.
- [Ham15] Alan Hammond. Sharp phase transition in the random stirring model on trees. *Probability Theory and Related Fields*, 161(3-4):429–448, 2015.
- [Kow] Emmanuel Kowalski. An introduction to expander graphs. <https://people.math.ethz.ch/~kowalski/expander-graphs.pdf>. [Online; accessed 18.8.2018].

- [LT] Mitchell Lee and Jonathan Tidor. Nonabelian fourier analysis. http://math.mit.edu/~csikvari/nonabelian_Fourier_analysis.pdf. [Online; accessed 21.4.2019].
- [Sch05] Oded Schramm. Compositions of random transpositions. *Israel Journal of Mathematics*, 147(1):221–243, 2005.
- [Tao] Terence Tao. Metric entropy analogues of sum-set theory. <https://terrytao.wordpress.com/2014/03/19/metric-entropy-analogues-of-sum-set-theory/>. [Online; accessed 21.4.2019].