

Cutoff for conjugacy-invariant random walks on the permutation group

Nathanaël Berestycki¹ · Batı Şengül¹

Received: 14 June 2017 / Revised: 9 February 2018 / Published online: 14 May 2018
© The Author(s) 2018

Abstract We prove a conjecture raised by the work of Diaconis and Shahshahani (*Z Wahrscheinlichkeitstheorie Verwandte Geb* 57(2):159–179, 1981) about the mixing time of random walks on the permutation group induced by a given conjugacy class. To do this we exploit a connection with coalescence and fragmentation processes and control the Kantorovich distance by using a variant of a coupling due to Oded Schramm as well as contractivity of the distance. Recasting our proof in the language of Ricci curvature, our proof establishes the occurrence of a phase transition, which takes the following form in the case of random transpositions: at time $cn/2$, the curvature is asymptotically zero for $c \leq 1$ and is strictly positive for $c > 1$.

Keywords Random walks · Symmetric group · Mixing times · Random transpositions · Conjugacy classes · Coalescence and fragmentation · Coarse Ricci curvature

Mathematics Subject Classification 60B15 · 60G50 · 60C05

Contents

1 Introduction	1198
1.1 Main results	1198
1.2 Relation to previous works on the geometry of random transpositions	1201

Nathanaël Berestycki’s research was supported in part by EPSRC Grants EP/GO55068/1, EP/I03372X/1 and EP/L018896/1. Batı Şengül’s research was carried out in part at the University of Cambridge by EPSRC Grant EP/H023348/1 and at the University of Bath by EPSRC Grant EP/L002442/1.

✉ Nathanaël Berestycki
beresty@statslab.cam.ac.uk

¹ Statistical Laboratory, University of Cambridge, Cambridge, UK

- 1.3 Relation to previous works on mixing times 1201
- 1.4 Organisation of the paper 1203
- 2 Curvature and mixing 1203
 - 2.1 Curvature theorem 1203
 - 2.2 Curvature implies mixing 1206
 - 2.3 Stochastic commutativity 1208
- 3 Preliminaries on random hypergraphs 1208
 - 3.1 Hypergraphs 1208
 - 3.2 Giant component of the hypergraph 1209
 - 3.3 Poisson–Dirichlet structure 1220
- 4 Proof of curvature theorem 1220
 - 4.1 Proof of the upper bound on curvature 1220
 - 4.2 Proof of lower bound on curvature 1222
 - 4.2.1 Coupling for $[0, s_1]$ 1224
 - 4.2.2 Coupling for $(s_1, s_2]$ 1225
 - 4.2.3 Coupling for $(s_2, s_3]$ 1233
- A Lower bound on mixing 1234
- B Proof of Theorem 3.6 1236
- References 1241

1 Introduction

1.1 Main results

Let \mathcal{S}_n denote the multiplicative group of permutations of $\{1, \dots, n\}$. Let $\Gamma \subset \mathcal{S}_n$ be a fixed conjugacy class in \mathcal{S}_n , i.e., $\Gamma = \{g\gamma g^{-1} : g \in \mathcal{S}_n\}$ for some fixed permutation $\gamma \in \mathcal{S}_n$. Alternatively, Γ is the set of permutation in \mathcal{S}_n having the same cycle structure as γ . Let $X^\sigma = (X_0, X_1, \dots)$ be discrete-time random walk on \mathcal{S}_n induced by Γ , started in the permutation $\sigma \in \mathcal{S}_n$, and let Y^σ be the associated continuous time random walk. These are the processes defined by

$$\begin{aligned} X_t^\sigma &= \sigma \circ \gamma_1 \circ \dots \circ \gamma_t; & t = 0, 1, \dots \\ Y_t^\sigma &= X_{N_t}^\sigma; & t \in [0, \infty) \end{aligned} \tag{1}$$

where $\gamma_1, \gamma_2, \dots$ are i.i.d. random variables which are distributed uniformly in Γ ; and $(N_t, t \geq 0)$ is an independent Poisson process with rate 1. Then Y is a Markov chain which converges to an invariant measure μ as $t \rightarrow \infty$. If $\Gamma \subset \mathcal{A}_n$ (where \mathcal{A}_n denotes the alternating group) then μ is uniformly distributed on \mathcal{A}_n and otherwise μ is uniformly distributed on \mathcal{S}_n . The simplest and most well known example of a conjugacy class is the set T of all transpositions, or more generally of all cyclic permutations of length $k \geq 2$. This set will play an important role in the rest of the paper. Note that Γ depends on n but we do not indicate this dependence in our notation.

The main goal of this paper is to study the cut-off phenomenon for the random walk X . More precisely, recall that the total variation distance $\|X - Y\|_{TV}$ between two random variables X, Y taking values in a set S is given by

$$\|X - Y\|_{TV} = \sup_{A \subset S} |\mathbb{P}(X \in A) - \mathbb{P}(Y \in A)|. \tag{2}$$

For $0 < \delta < 1$, the mixing time $t_{\text{mix}}(\delta)$ is by definition given by

$$t_{\text{mix}}(\delta) = \inf\{t \geq 0 : d_{TV}(t) \leq \delta\}$$

where

$$d_{TV}(t) = \sup_{\sigma} \|Y_t^{\sigma} - \mu\|_{TV} \quad (3)$$

and μ is the invariant measure defined above.

In the case where $\Gamma = T$ is the set of transpositions, a famous result of Diaconis and Shahshahani [10] is that the cut-off phenomenon takes place at time $(1/2)n \log n$ asymptotically as $n \rightarrow \infty$. That is, $t_{\text{mix}}(\delta)$ is asymptotic to $(1/2)n \log n$ for any fixed value of $0 < \delta < 1$. It has long been conjectured that for a general conjugacy class such that $|\Gamma| = o(n)$ (where here and in the rest of the paper, $|\Gamma|$ denotes the number of non fixed points of any permutation $\gamma \in \Gamma$), a similar result should hold at a time $(1/|\Gamma|)n \log n$. This has been verified for k -cycles with a fixed $k \geq 2$ by Berestycki et al. [6]. This is a problem with a substantial history which will be detailed below. The primary purpose of this paper is to verify this conjecture. Hence our main result is as follows.

Theorem 1.1 *Let $\Gamma \subset \mathcal{S}_n$ be a conjugacy class and suppose that $|\Gamma| = o(n)$. Define*

$$t_{\text{mix}} := \frac{1}{|\Gamma|} n \log n. \quad (4)$$

Then for any $\epsilon > 0$,

$$\lim_{n \rightarrow \infty} d_{TV}((1 - \epsilon)t_{\text{mix}}) = 1 \quad \text{and} \quad \lim_{n \rightarrow \infty} d_{TV}((1 + \epsilon)t_{\text{mix}}) = 0. \quad (5)$$

The first limit of (5) is proved in Appendix A. The rest of the paper focuses on the second limit. Our main tool for this result is the notion of discrete Ricci curvature as introduced by Ollivier [18], for which we obtain results of independent interest. We briefly discuss this notion here; however we point out that this turns out to be equivalent to the more well-known *path coupling method* and transportation metric introduced by Bubley and Dyer [8] and Jerrum [14] (see for instance Chapter 14 of the book [16] for an overview). However we will cast our results in the language of Ricci curvature because we find it more intuitive. Recall first that the definition of the L^1 -Kantorovich distance (sometimes also called Wasserstein or transportation metric) between two random variables X, Y taking values in a metric space (S, d) is given by

$$W_1(X, Y) := \inf \mathbb{E}[d(\hat{X}, \hat{Y})] \quad (6)$$

where the infimum is taken over all couplings (\hat{X}, \hat{Y}) which are distributed marginally as X and Y respectively. Ollivier's definition of Ricci curvature of a Markov chain $(X_t, t \geq 0)$ on a metric space (S, d) is as follows:

Definition 1.1 Let $t > 0$. The curvature between two points $x, x' \in S$ with $x \neq x'$ is given by

$$\kappa_t(x, x') := 1 - \frac{W_1(X_t^x, X_t^{x'})}{d(x, x')} \quad (7)$$

where X_t^x and $X_t^{x'}$ denote Markov chains started from x and x' respectively. The curvature of X is by definition equal to

$$\kappa_t := \inf_{x \neq x'} \kappa_t(x, x').$$

In the terminology of Ollivier [18], this is in fact the curvature of the discrete-time random walk whose transition kernel is given by $m_x(\cdot) = \mathbb{P}(X_t = \cdot | X_0 = x)$. We refer the reader to [18] for an account of the elegant theory which can be developed using this notion of curvature, and point out that a number of classical properties of curvature generalise to this discrete setup.

For our results it will turn out to be convenient to view the symmetric group as a metric space equipped with the metric d which is the word metric induced by the set T of transpositions (we will do so even when the random walk is not induced by T but by a general conjugacy class Γ). That is, the distance $d(\sigma, \sigma')$ between $\sigma, \sigma' \in \mathcal{S}_n$ is the minimal number of transpositions one must apply to get from one element to the other (one can check that this number is independent of whether right-multiplications or left-multiplications are used).

For simplicity we focus in this introduction on the case where the random walk is induced by the set of transpositions T . (A more general result will be stated later on the paper). For $c > 0$ and $\sigma \neq \sigma'$, let

$$\kappa_c(\sigma, \sigma') = 1 - \frac{W_1\left(X_{\lfloor cn/2 \rfloor}^\sigma, X_{\lfloor cn/2 \rfloor}^{\sigma'}\right)}{d(\sigma, \sigma')} \tag{8}$$

and define $\kappa_c(\sigma, \sigma) = 1$. That is, $\kappa_c(\sigma, \sigma') = \kappa_{\lfloor cn/2 \rfloor}(\sigma, \sigma')$ with our notation from (7). In particular, κ_c depends on n but this dependency does not appear explicitly in the notation. It is not hard to see that $\kappa_c(\sigma, \sigma') \geq 0$ (apply the same transpositions to both walks X^σ and $X^{\sigma'}$). For parity reasons it is obvious that $\kappa_c(\sigma, \sigma') = 0$ if σ and σ' do not have the same signature. Thus we only consider the curvature between elements of even distance. For $c > 0$ define

$$\kappa_c = \inf \kappa_c(\sigma, \sigma'),$$

where the infimum is taken over all $\sigma, \sigma' \in \mathcal{S}_n$ such that $d(\sigma, \sigma')$ is even. Our main result states that κ_c experiences a phase transition at $c = 1$. More precisely, the curvature κ_c is asymptotically zero for $c \leq 1$ but for $c > 1$ the curvature is strictly positive asymptotically. In order to state our result, we introduce the quantity $\theta(c)$, which is the largest solution in $[0, 1]$ to the equation

$$\theta(c) = 1 - e^{-c\theta(c)}. \tag{9}$$

It is easy to see that $\theta(c) = 0$ for $c \leq 1$ and $\theta(c) > 0$ for $c > 1$. In fact, $\theta(c)$ is nothing else but the survival probability of a Galton–Watson tree with Poisson offspring distribution with mean c .

Theorem 1.2 *For any $c > 0$, we have:*

$$\theta(c)^4 \leq \liminf_{n \rightarrow \infty} \kappa_c \leq \limsup_{n \rightarrow \infty} \kappa_c \leq \theta(c)^2 \tag{10}$$

In particular, $\lim_{n \rightarrow \infty} \kappa_c = 0$ if and only if $c \leq 1$, while $\liminf_{n \rightarrow \infty} \kappa_c > 0$ otherwise.

A more general version of this theorem will be presented later on, which gives results for the curvature of a random walk induced by a general conjugacy class Γ . This will be stated as Theorem 2.3.

We believe that the upper bound is the sharp one here, and thus make the following conjecture.

Conjecture 1.3 *For $c > 0$,*

$$\lim_{n \rightarrow \infty} \kappa_c = \theta(c)^2.$$

Of course the conjecture is already established for $c \leq 1$ and so is only interesting for $c > 1$.

1.2 Relation to previous works on the geometry of random transpositions

The transition described by Theorem 1.2 says that the discrete Ricci curvature increases abruptly (asymptotically) from zero to a positive quantity as c increases past the critical value $c = 1$, and so as we consider longer portions of the random walk. It is related to a result proved by the first author in [2]. There it was shown that the triangle formed by the identity and two independent samples X_t and X'_t from the random walk run for time $t = cn/4$, is thin (in the sense of Gromov hyperbolicity) if and only if $c < 1$. Note that by reversibility, the path running from X_t to X'_t (via the identity) is a random walk run for time $cn/2$. In other words, the result from [2] implies that the permutation group appears Gromov hyperbolic from the point of view of a random walker so long as it takes fewer than $cn/2$ steps with $c < 1$.

Hence, in both Theorem 1.2 and [2], there is a change of geometry (as perceived by a random walker) from low to high curvature after running for exactly $t = cn/2$ steps with $c = 1$. At this point, we do not know of a formal way to relate these two observations, so they simply seem analogous. In a private conversation with the first author in 2005, Gromov had suggested that the hyperbolicity transition of [2] could be translated more canonically into the language of Ricci curvature and was an effect of the global positive curvature of \mathcal{S}_n rather than a breakdown in hyperbolicity. In a sense, Theorem 1.2 can be seen as a formalisation and justification of his prediction.

1.3 Relation to previous works on mixing times

Mixing times of Markov chains were initiated independently by Aldous [1] and by Diaconis and Shahshahani [10]. In particular, as already mentioned, Diaconis and Shahshahani proved Theorem 1.1 in the case where Γ is the set T of transpositions. Their proof relies on some deep connections with the representation theory of \mathcal{S}_n and bounds on so-called character ratios. The conjecture about the general case appears to have first been made formally in print by Roichman [20] but it has no doubt been asked privately before then. We shall see that the lower bound $t_{\text{mix}}(\delta) \geq (1 + o(1))(1/|\Gamma|)n \log n$ is fairly straightforward (it is carried out in

Appendix A and is as usual based on a coupon-collector type argument); the difficult part is the corresponding upper bound.

Flatto et al. [13] built on the earlier work of Vershik and Kerov [25] to obtain that $t_{\text{mix}}(\delta) \leq (1/2 + o(1))n \log n$ when $|\Gamma|$ is bounded (as is noted in [9, pp. 44–45]). This was done using character ratios and this method was extended further by Roichman [20, 21] to show an upper bound on $t_{\text{mix}}(\delta)$ which is sharp up to a constant when $|\Gamma| = o(n)$ (and in fact, more generally when $|\Gamma|$ is allowed to grow to infinity as fast as $(1 - \delta)n$ for any $\delta \in (0, 1)$). Again using character ratios Lulov and Pak [17] showed the cut-off phenomenon as well as $t_{\text{mix}} = (1/|\Gamma|)n \log n$ in the case when $|\Gamma| \geq n/2$. Roussel [22, 23] obtains the correct value of the mixing time and establishes the cut-off phenomenon for the case $|\Gamma| \leq 6$.

Finally, let us discuss two more recent papers to which this work is most closely related to. Berestycki et al. [6], show using coupling arguments and a connection to coalescence–fragmentation processes that the cutoff phenomenon occurs at $t_{\text{mix}} = (1/k)n \log n$ in the case when Γ consists only of cycles of length k for any $k \geq 2$ fixed.

Shortly after, Bormashenko [7] devised a path coupling argument for the coagulation-fragmentation process associated to random transpositions to obtain a new proof of a slightly weaker version of the Diaconis–Shahshahani result: her argument implies that the mixing time of random transpositions is $O(n \log n)$ (unfortunately the implicit multiplicative constant is not sharp, so this is not sufficient to obtain cutoff). See also [19] for another discussion of her results together with a reformulation in the language of coarse Ricci curvature. In a way her approach is very similar to ours, to the point that it can be considered a precursor to our work, since our method is also based on a certain path coupling for the coagulation-fragmentation process which exploits certain remarkable properties of Schramm’s coupling [6, 24].

Comparison with [6] The authors in Berestycki et al. [6] remark that their proof can be extended to cover the case when Γ is a fixed conjugacy class and indicate that their methods can probably be pushed to cover the case when $|\Gamma| = o(n^{1/2})$, but it is clear that new ideas are needed if $|\Gamma|$ is larger. Indeed, their argument uses very delicate estimates about the behaviour of small cycles, together with a variant of a coupling due to Schramm [24] to deal with large cycles. The most technical part of their argument is to analyse the distribution of small cycles, using delicate couplings and carefully bounding the error made in these couplings.

However, when $k = |\Gamma|$ is larger than $n^{1/2}$, we can no longer think of the points in the conjugacy class as being sampled independently (with replacement) from $\{1, \dots, n\}$, by the birthday problem. This introduces many more ways in which errors in the above coupling arguments could occur. These seem quite hard to control, and hence new ideas are required for the general case.

The proof in this paper relies on similar observations as [6], and in particular the connection with coalescence–fragmentation process as well as Schramm’s coupling argument play a crucial role. The key new idea however, is to try to prove mixing not just in the total variation sense but in the stronger sense of the L^1 -Kantorovich distance (Ricci curvature) and to estimate it at a time well before the mixing time, roughly $O(n/k)$ instead of $O(n(\log n)/k)$. This may seem counterintuitive initially, however studying the random walk at this time scale allows us to make precise comparisons

between the random walk and an associated random graph process. It turns out the random graph at these time scales can be described rather precisely. Furthermore, due to the contraction properties of the Kantorovich distance, somehow (and rather miraculously, we find), the estimate we obtain can be bootstrapped with sufficient precision to yield mixing exactly at the time $t_{\text{mix}} = (1/k)n \log n$.

In particular, since the heart of the proof consists in studying the situation at a time well before mixing, and purely to take advantage of the giant component at such times, we *never* have to study the distribution of small cycles. This is really quite surprising, given that the small cycles (in particular, the fixed points) are responsible for the occurrence of the cutoff at time t_{mix} .

1.4 Organisation of the paper

We stress that compared to [6], the main arguments are quite elementary. The heart of the proof is contained in Sects. 4.2 and 2. Readers who are familiar with [6] are encouraged to concentrate on these two short sections.

The paper is organised as follows. In Sect. 2 we state and discuss Theorem 2.3, which is a general curvature theorem (of which Theorem 1.2 is the prototype). We also discuss why this implies the main theorem (Theorem 1.1). In Sect. 3.1 we study the associated random hypergraph process. The main result in that section is Theorem 3.1, which proves the existence and uniqueness of the giant component. Curiously this is the most technical aspect of the paper, and really the only place where the myriad of ways in which the conjugacy class Γ might be really big plays a role and needs to be controlled. Section 4 contains a proof of the main curvature theorem (Theorem 2.3), starting with the easy upper bound on curvature (Sect. 4.1) and following up with the slightly more complex lower bound (Sect. 4.2), which really is the heart of the proof. The two appendices contain respectively a proof of the lower bound on the mixing time (certainly known in the folklore, essentially a version of the coupon collector lemma); and an adaptation of Schramm's argument [24] for the Poisson–Dirichlet structure of cycles inside the giant component, which is needed in the proof.

2 Curvature and mixing

2.1 Curvature theorem

As discussed above, the lower bound (5) is relatively easy and is probably known in the folklore; we give a proof in Appendix A. We now start the proof of the main results of this paper, which is the upper bound (the right hand side) of (5). In this section, we first state the more general version of Theorem 1.2 discussed in the introduction, and we will then show how this implies the desired result for the upper bound on $t_{\text{mix}}(\delta)$. To begin, we define the cycle structure (k_2, k_3, \dots) of Γ to be a vector such that for each $j \geq 2$, there are k_j cycles of length j in the cycle decomposition of any $\gamma \in \Gamma$ (note that this does not depend on $\tau \in \Gamma$). Then $k_j = 0$ for all $j > n$ and we have that $k := |\Gamma| = \sum_{j=2}^{\infty} jk_j$.

In the case for the transposition random walk the quantity $\theta(c)$ which appears in the bounds is the survival probability of a Galton–Watson process with offspring distribution given by a Poisson random variable with mean c . Our first task is to generalise $\theta(c)$. We do so via a fixed point equation, which is more complex here. Define

$$\alpha_j = \frac{jk_j}{k},$$

and note that $\alpha_j \in [0, 1]$ (α_j is the proportion of the mass in cycles of size j for any $\gamma \in \Gamma$). Thus $(\alpha_j)_{j \geq 2}$ is compact in the product topology (the topology of pointwise convergence). Suppose that the limit

$$(\bar{\alpha}_2, \bar{\alpha}_3, \dots) := \lim_{n \rightarrow \infty} (\alpha_2, \alpha_3, \dots) \tag{C}$$

exists, where the limit is taken to be pointwise. It follows that for each $j \geq 2$, $\bar{\alpha}_j \in [0, 1]$ and $\sum_{j=2}^{\infty} \bar{\alpha}_j \leq 1$ by Fatou’s lemma. Note that the sum is strictly less than 1 when a positive fraction of the mass of conjugacy class Γ comes from cycles whose size tends to ∞ . This will be an important distinction in what follows. For $x \in [0, 1]$ and $c > 0$ define

$$\psi(x, c) = \exp \left\{ -c \left(1 - \sum_{j=2}^{\infty} \bar{\alpha}_j (1-x)^{j-1} \right) \right\}. \tag{11}$$

Note that for each $c > 0$, $x \mapsto \psi(x, c)$ is convex on $[0, 1]$. Moreover, the function $x \mapsto \psi(1-x, c)$ is the generating function of a random variable whose law depends on c and is degenerate if $\sum_{j \geq 2} \bar{\alpha}_j < 1$. Note that in the case $\Gamma = T$ of transpositions, $\psi(x, c) = e^{-cx}$ so that random variable is simply Poisson (c).

Lemma 2.1 *Define*

$$c_\Gamma := \begin{cases} \left(\sum_{j=2}^{\infty} (j-1)\bar{\alpha}_j \right)^{-1} & \text{if } \sum_{j=2}^{\infty} \bar{\alpha}_j = 1 \\ 0 & \text{if } \sum_{j=2}^{\infty} \bar{\alpha}_j < 1. \end{cases} \tag{12}$$

Then for $c > c_\Gamma$ there exists a unique $\theta(c) \in (0, 1)$ such that

$$\theta(c) = 1 - \psi(\theta(c), c).$$

For $c > c_\Gamma$, $c \mapsto \theta(c)$ is increasing, continuous and differentiable. Further $\lim_{c \downarrow c_\Gamma} \theta(c) = 0$ and $\lim_{c \uparrow \infty} \theta(c) = 1$.

Proof For $x \in [0, 1]$ and $c > 0$ define $f_c(x) := 1 - \psi(x, c) - x$. There are two cases to consider. First suppose that $z = \sum_{j=2}^{\infty} \bar{\alpha}_j < 1$. Then we have that

$$f_c(0) = 1 - e^{-c(1-z)} > 0 \quad \text{and} \quad f_c(1) = -e^{-c} < 0.$$

As $x \mapsto f_c(x)$ is concave on $[0, 1]$ it follows that there exists a unique $\theta(c) \in (0, 1)$ such that $f_c(\theta(c)) = 0$.

Next suppose that $\sum_{j=2}^{\infty} \bar{\alpha}_j = 1$, then

$$f_c(0) = 0 \quad \text{and} \quad f_c(1) = -e^{-c} < 0$$

Moreover we have that

$$\frac{d}{dx} f_c(x)|_{x=0} = c \sum_{j=2}^{\infty} (j-1) \bar{\alpha}_j - 1.$$

Hence for $c > c_{\Gamma}$ we have that $\frac{d}{dx} f_c(x)|_{x=0} > 0$ and again by concavity it follows that there exists a unique $\theta(c) \in (0, 1)$ such that $f_c(\theta(c)) = 0$.

For the rest of the statements suppose that $c > c_{\Gamma}$. The fact that $c \mapsto \theta(c)$ is increasing follows from the definition of $\psi(x, c)$ and the fact that $\theta(c) = \psi(\theta(c), c)$. Continuity and differentiability for $c > c_{\Gamma}$ is a straightforward application of the inverse function theorem.

Notice that $\theta(c) \in [0, 1]$ and is monotone, hence $\theta(c)$ converges as $c \downarrow c_{\Gamma}$ to a limit L . Then it follows that L solves the equation $L = 1 - \psi(L, c_{\Gamma})$. This equation has only a zero solution and thus $L = 0$ and hence $\lim_{c \downarrow c_{\Gamma}} \theta(c) = 0$. The limit as $c \uparrow \infty$ follows from a similar argument. □

Remark 2.2 In the case when $\Gamma = T$ is the set of transpositions we have that $k'_2 = 1$ and $\bar{\alpha}_j = 0$ for $j \geq 3$, hence $\psi(x, c) = e^{-cx}$ and thus the definition of $\theta(c)$ above agrees with the definition given in the introduction.

Having introduced $\theta(c)$ we now introduce the notion of Ricci curvature we will use in the general case. For $c > 0$ and $\sigma \neq \sigma'$, let

$$\kappa_c(\sigma, \sigma') = 1 - \frac{W_1 \left(X_{[cn/k]}^{\sigma}, X_{[cn/k]}^{\sigma'} \right)}{d(\sigma, \sigma')} \tag{13}$$

where d is the graph distance associated with transpositions (even in the case $\Gamma \neq T$). Define $\kappa_c(\sigma, \sigma) = 1$. Then let

$$\kappa_c = \inf \kappa_c(\sigma, \sigma'),$$

where the infimum is taken over all $\sigma, \sigma' \in \mathcal{S}_n$ such that $d(\sigma, \sigma')$ is even. That is, $\kappa_c(\sigma, \sigma') = \kappa_{[cn/k]}(\sigma, \sigma')$ with our notation from (7). We now state a more general form of Theorem 1.2 which in particular covers the case of Theorem 1.2.

Theorem 2.3 *Let $\Gamma \subset \mathcal{S}_n$ be a conjugacy class such that $k = |\Gamma| = o(n)$ and the convergence in (C) holds. Recall the definition of c_{Γ} from (12). Then for $c \leq c_{\Gamma}$,*

$$\lim_{n \rightarrow \infty} \kappa_c = 0. \tag{14}$$

On the other hand, for $c > c_{\Gamma}$

$$\theta(c)^4 \leq \liminf_{n \rightarrow \infty} \kappa_c \leq \limsup_{n \rightarrow \infty} \kappa_c \leq \theta(c)^2 \quad (15)$$

where $\theta(c)$ is the unique solution in $(0, 1)$ of

$$\theta(c) = 1 - \psi(\theta(c), c). \quad (16)$$

where ψ is given by (11).

2.2 Curvature implies mixing

We now show how Theorem 2.3 implies the second limit in Theorem 1.1. First suppose that $\Gamma = \Gamma(n)$ is a sequence of conjugacy classes for which the limit (C) holds and $|\Gamma| = o(n)$. Again fix $\epsilon > 0$ and define $t = (1 + 2\epsilon)(1/k)n \log n$ and let $t' = \lfloor (1 + \epsilon)(1/k)n \log n \rfloor$ where $k = |\Gamma|$. We are left to prove that $d_{TV}(t) \rightarrow 0$ as $n \rightarrow \infty$. For $s \geq 0$ let

$$\bar{d}_{TV}(s) := \sup_{\sigma, \sigma'} \|X_s^\sigma - X_s^{\sigma'}\|_{TV},$$

where the sup is taken over all permutations at even distances. We first claim that it suffices to prove that

$$\bar{d}_{TV}(t') \rightarrow 0 \text{ as } n \rightarrow \infty. \quad (17)$$

Indeed, assume that $\bar{d}_{TV}(t') \rightarrow 0$ as $n \rightarrow \infty$. Then there are two cases to consider. Assume that $\Gamma \subset \mathcal{A}_n$. Then $X_s \in \mathcal{A}_n$ for all $s \geq 1$ and μ is uniform on \mathcal{A}_n . Then by Lemma 4.11 in [16],

$$\sup_{\sigma \in \mathcal{A}_n} \|X_{t'}^\sigma - \mu\|_{TV} \leq 2\bar{d}_{TV}(t').$$

Hence Theorem 1.1 (or more precisely the second limit in that theorem) follows from (17) in this case. In the second case, $\Gamma \subset \mathcal{A}_n^c$. In this case $X_s \in \mathcal{A}_n$ for s even, and $X_s \in \mathcal{A}_n^c$ for s odd. Using the same lemma, we deduce that if $s \geq t'$ is even,

$$\|X_s^{\text{id}} - \mu_1\|_{TV} \leq 2\bar{d}_{TV}(s)$$

where μ_1 is uniform on \mathcal{A}_n . However, if $s \geq t'$ is odd,

$$\|X_s^{\text{id}} - \mu_2\|_{TV} \leq 2\bar{d}_{TV}(s)$$

where this time μ_2 is uniform on \mathcal{A}_n^c . Let $N = (N_s : s \geq 0)$ be the Poisson clock of the random walk Y . Then $\mathbb{P}(N_s \text{ even}) \rightarrow 1/2$ as $s \rightarrow \infty$, $\mu = (1/2)(\mu_1 + \mu_2)$, and $\mathbb{P}(N_t \geq t') \rightarrow 1$ as $n \rightarrow \infty$. Thus we deduce that

$$\|Y_t^{\text{id}} - \mu\|_{TV} \rightarrow 0.$$

Again, the second limit in Theorem 1.1 follows. Hence it suffices to prove (17).

Note that for any two random variables X, Y on a metric space (S, d) we have the obvious inequality $\|X - Y\|_{TV} \leq W_1(X, Y)$ provided that $x \neq y$ implies $d(x, y) \geq 1$ on S . This is in particular the case when $S = \mathcal{S}_n$ and d is the word metric induced by the set T of transpositions. In other words it suffices to prove mixing in the L^1 -Kantorovich distance.

Note that by definition of κ_c , if σ, σ' are at an even distance then

$$W_1\left(X_{\lfloor cn/k \rfloor}^\sigma, X_{\lfloor cn/k \rfloor}^{\sigma'}\right) \leq (1 - \kappa_c)d(\sigma, \sigma'),$$

so that, iterating as in Corollary 21 of [18] (and noting that the distance between $X_{\lfloor cn/k \rfloor}^\sigma$ and $X_{\lfloor cn/k \rfloor}^{\sigma'}$ is again even), we have for each $s \geq 1$,

$$\sup_{d(\sigma, \sigma') \text{ even}} W_1\left(X_{s\lfloor cn/k \rfloor}^\sigma, X_{s\lfloor cn/k \rfloor}^{\sigma'}\right) \leq (1 - \kappa_c)^s \sup_{d(\sigma, \sigma') \text{ even}} d(\sigma, \sigma') \leq n(1 - \kappa_c)^s \tag{18}$$

since the diameter of \mathcal{S}_n is equal to $n - 1$. Solving $n(1 - \kappa_c)^s \leq \delta$ we get that

$$s \geq \frac{\log n - \log \delta}{-\log(1 - \kappa_c)} \tag{19}$$

Thus if $u = scn/k \geq s\lfloor cn/k \rfloor$, it suffices that

$$u \geq \frac{1}{k - \log(1 - \kappa_c)} c n (\log n - \log \delta). \tag{20}$$

Now, Theorem 2.3 gives

$$\liminf_{n \rightarrow \infty} -\log(1 - \kappa_c) \geq -\log(1 - \theta(c)^4).$$

Lemma 2.4 *We have that*

$$\lim_{c \rightarrow \infty} \frac{c}{\log(1 - \theta(c)^4)} = -1.$$

Proof Using L'Hopital's rule twice we have that

$$\lim_{\theta \uparrow 1} \frac{\log(1 - \theta)}{\log(1 - \theta^4)} = \lim_{\theta \uparrow 1} \frac{1 - \theta^4}{(1 - \theta)4\theta^3} = 1.$$

Next we have that $\lim_{c \rightarrow \infty} \theta(c) = 1$ and hence

$$\begin{aligned} \lim_{c \rightarrow \infty} \frac{c}{\log(1 - \theta(c)^4)} &= \lim_{c \rightarrow \infty} \frac{c}{\log(1 - \theta(c))} = \lim_{c \rightarrow \infty} \frac{c}{\log \psi(\theta(c), c)} \\ &= \lim_{c \rightarrow \infty} \frac{1}{1 - \sum_{j=2}^{\infty} \bar{\alpha}_j (1 - \theta(c))^{j-1}} = -1. \end{aligned}$$

□

Consequently we have that for $u \geq t' = \lfloor (1 + \epsilon)(1/k)n \log n \rfloor$ u satisfies (20) for some sufficiently large $c > c_\Gamma$. Hence $\limsup_{n \rightarrow \infty} \bar{d}_{TV}(t') \rightarrow 0$ and thus (17) holds, which shows Theorem 1.1 for conjugacy classes such that the limit in (C) exists and $|\Gamma| = o(n)$.

Now suppose that Γ is a conjugacy class such that $|\Gamma| = o(n)$. Let $t' = \lfloor (1 + \epsilon)(1/|\Gamma|)n \log n \rfloor$ and notice that $d_{TV}(t')$ is bounded. Along any subsequence $\{n_i\}_{i \geq 1}$ such that $\lim_{n_i \rightarrow \infty} d_{TV}(t')$ exists, we can extract a further sub-sequence $\{n_{i_j}\}_{j \geq 1}$ such that (C) holds since $(\alpha_j)_{j \geq 2} \in [0, 1]^\infty$ which is compact under the product topology. Then we see that $\lim_{n_{i_j} \rightarrow \infty} d_{TV}(t') = 0$ and consequently $\lim_{n_i \rightarrow \infty} d_{TV}(t') = 0$. Since $d_{TV}(t')$ is bounded and converges to 0 along any convergent subsequence, we conclude that $\lim_{n \rightarrow \infty} d_{TV}(t') = 0$, thus concluding the proof.

2.3 Stochastic commutativity

To conclude this section on curvature, we state a simple but useful lemma. Roughly, this says that the random walk is “stochastically commutative”. This can be used to show that the L^1 -Kantorovich distance is decreasing under the application of the heat kernel. In other words, initial discrepancies for the Kantorovich metric between two permutations are only smoothed out by the application of random walk.

Lemma 2.5 *Let σ be a random permutation with distribution invariant by conjugacy. Let σ_0 be a fixed permutation. Then $\sigma_0 \circ \sigma$ has the same distribution as $\sigma \circ \sigma_0$.*

Proof Define $\sigma' = \sigma_0 \circ \sigma \circ \sigma_0^{-1}$. Then since σ is invariant under conjugacy, the law of σ' is the same as the law of σ . Furthermore, we have $\sigma_0 \circ \sigma = \sigma' \circ \sigma_0$ so the result is proved. \square

This lemma will be used repeatedly in our proof, as it allows us to concentrate on events of high probability for our coupling.

3 Preliminaries on random hypergraphs

For the proof of Theorem 1.1 we rely on properties of certain random hypergraph processes. The reader who is only interested in a first instance in the case of random transpositions, and is familiar with Erdős–Renyi random graphs and with the result of Schramm [24] may safely skip this section.

3.1 Hypergraphs

In this section we present some preliminaries which will be used in the proof of Theorem 2.3. Throughout we let $\Gamma \subset \mathcal{S}_n$ be a conjugacy class and let (k_2, k_3, \dots) denote the cycle structure of Γ . Thus Γ consists of permutations such that in their cycle decomposition they have k_2 many transpositions, k_3 many 3-cycles and so on. Note that we have suppressed the dependence of Γ and (k_2, k_3, \dots) on n . We assume that (C) is satisfied so that for each $j \geq 2$, $jk_j/|\Gamma| \rightarrow \bar{\alpha}_j$ as $n \rightarrow \infty$. We also let $k = |\Gamma|$ so that $k = \sum_{j \geq 2} jk_j$, as usual.

Definition 3.1 A hypergraph $H = (V, E)$ is given by a set V of vertices and $E \subset \mathcal{P}(V)$ of edges, where $\mathcal{P}(V)$ denotes the set of all subsets of V . An element $e \in E$ is called a hyperedge and we call it a j -hyperedge if $|e| = j$.

Consider the random walk $X = (X_t : t = 0, 1, \dots)$ on \mathcal{S}_n where $X_t = X_t^{\text{id}}$ with our notations from the introduction. Hence

$$X_t = \gamma_1 \circ \dots \circ \gamma_t$$

where the sequence $(\gamma_i)_{i \geq 1}$ is i.i.d. uniform on Γ . A given step of the random walk, say γ_s , can be broken down into cycles, say $\gamma_{s,1} \circ \dots \circ \gamma_{s,r}$ where $r = \sum_j k_j$. We will say that a given cyclic permutation γ has been applied to X before time t if $\gamma = \gamma_{s,i}$ for some $s \leq t$ and $1 \leq i \leq r$.

To X we associate a certain hypergraph process $H = (H_t : t = 0, 1, \dots)$ defined as follows. For $t = 0, 1, \dots$, H_t is a hypergraph on $\{1, \dots, n\}$ where a hyperedge $\{x_1, \dots, x_j\}$ is present if and only if a cyclic permutation consisting of the points x_1, \dots, x_j in some arbitrary order has been applied to the random walk X prior to time t as part of one of the γ_i 's for some $i \leq t$. Thus at every step, we add to H_t k_j hyperedges of size j sampled uniformly at random without replacement, and these edges are independent from step to step. However, note that the presence of hyperedges themselves are not in general independent.

3.2 Giant component of the hypergraph

In the case $\Gamma = T$, the set of transpositions, the hypergraph H_s is a realisation of an Erdős–Renyi graph. Analogous to Erdős–Renyi graphs, we first present a result about the size of the components of the hypergraph process $H = (H_t : t = 0, 1, \dots)$ (where by size, we mean the number of vertices in this component). For the next result recall the definition of $\psi(x, c)$ in (11). Recall that for $c > c_\Gamma$, where c_Γ is given by (12), there exists a unique root $\theta(c) \in (0, 1)$ of the equation $\theta(c) = 1 - \psi(\theta(c), c)$.

Theorem 3.1 Consider the random hypergraph H_s and suppose that $s = s(n)$ is such that $sk/n \rightarrow c$ as $n \rightarrow \infty$ for some $c > c_\Gamma$. Then there is a universal constant $D > 0$ such that with probability tending to one all components but the largest have size at most $Dn^{2/3}(\log(n))^3$. Furthermore, the size of the largest component, normalised by n , converges to $\theta(c)$ in probability as $n \rightarrow \infty$.

Of course, this is the standard Erdős–Renyi theorem in the case where $\Gamma = T$ is the set of transpositions. See for instance [12], in particular Theorem 2.3.2 for a proof. In the case of k -cycles with k fixed and finite, this is the case of random regular hypergraphs analysed by Karoński and Łuczak [15]. For the slightly more general case of bounded conjugacy classes, this was proved by Berestycki [4].

Discussion Note that the behaviour of H_s in Theorem 3.1 can deviate markedly from that of Erdős–Renyi graphs. The most obvious difference is that H_s can contain mesoscopic components, something which has of course negligible probability for Erdős–Renyi graphs. For example, suppose Γ consists of $n^{1/2}$ transpositions and one

cycle of length $n^{1/3}$. Then the giant component appears at time $n^{1/2}/2$ with a phase transition (i.e., $c_\Gamma > 0$, because in this case $\sum \bar{\alpha}_j = 1$, as most of the mass comes from microscopic cycles). Yet even at the first step there is a component of size $n^{1/3}$. Nevertheless we will see that once there is a giant component there is a limit to how big can the nongiant component be (we show this is less than $O(n^{2/3})$ up to logarithmic terms; this is certainly not optimal).

From a technical point of view this has nontrivial consequences, as proofs of the existence of a giant component are usually based on the dichotomy between microscopic components and giant components. Furthermore, when the conjugacy class is large and consists of many small or mesoscopic cycles, the hyperedges have a strong dependence, which makes the proof very delicate.

In effect, perhaps surprisingly this will be the only place of the proof where all the possible ways in which the conjugacy class Γ might be big (potentially of size very close to n), needs to be handled. The difficulty of the proof below is to find an argument which works no matter how Γ is made up, so long as $k = |\Gamma| = o(n)$. This is of course also the problem in the original question of studying the mixing time of the random walk induced by Γ . However, what we have gained here compared to this original question, is the monotonicity of component sizes when hyperedges are added to H_s .

Preliminaries: exploration Suppose that $s = s(n)$ is such that $sk/n \rightarrow c$ for some $c > 0$ as $n \rightarrow \infty$ for some $c \geq 0$. We reveal the vertices of the component containing a fixed vertex $v \in \{1, \dots, n\}$ using breadth-first search exploration, as follows. There are three states that each vertex can be: unexplored, removed or active. Initially v is active and all the other vertices are unexplored. At each step of the iteration we select an active vertex w according to some prescribed rule among the active vertices at this stage (say with the smallest label). The vertex w becomes removed and every unexplored vertex which is joined to w by a hyperedge becomes active. We repeat this exploration procedure until there are no more active vertices. At stage $i = 0, 1, \dots$ of this exploration process, we let A_i , R_i and U_i denote the set of active, removed and unexplored vertices respectively. Thus initially $A_0 = \{v\}$, $U_0 = \{1, \dots, n\} \setminus \{v\}$ and $R_0 = \emptyset$. We will let $a_i = |A_i|$, $u_i = |U_i|$, $r_i = |R_i|$.

For $t = 1, \dots, s$ we call the hyperedges which are associated with the permutation γ_t the t -th packet of hyperedges. Thus note that each packet consists of k_j hyperedges of size j , $j \geq 2$, which are sampled uniformly at random without replacement from $\{1, \dots, n\}$. In particular, within a given packet, hyperedges are not independent. However, crucially, hyperedges from different packets are independent. We will need to keep track of the hyperedges we reveal and where they “came from” (i.e., which packet they were part of), in order to deal with these dependencies. More precisely, as we explore the hypergraph H_s , we discover various hyperedges of various sizes in H_s and this may affect the likelihood of other types of hyperedges in subsequent steps of the exploration process. To account for this, we introduce for $t = 1, \dots, s$ and for $j \geq 2$, the random subset of $\{1, \dots, n\}$, $Y_j^{(t)}(i)$, which is defined to be the hyperedges of size j in the t -th packet that were revealed in the exploration process prior to step i . We let $y_j^{(t)}(i) = |Y_j^{(t)}(i)|$ denote the number of such hyperedges.

Additional notations Let $i \geq 0$ and let \mathcal{H}_i denote the filtration generated by the exploration process up to stage i , including the information of the number of hyperedges of each size in each packet that were revealed up to step i of the exploration process. That is,

$$\mathcal{H}_i = \sigma \left(A_1, \dots, A_i, Y_j^{(t)}(1), \dots, Y_j^{(t)}(i) : 1 \leq t \leq s, j \geq 2 \right).$$

Our first goal will be to give uniform stochastic bounds on the distribution of $a_{i+1} - a_i$, so long as i is not too large. We will thus fix i (a step in the exploration process) and in order to ease notations we will often suppress the dependence on i , in $Y_j^{(t)}(i)$: we will thus simply write $Y_j^{(t)}$ and $y_j^{(t)}$. Note that by definition, for each $t = 1, \dots, s$ and $j \geq 2$, $Y_j^{(t)} \leq k_j$ and

$$\sum_{t=1}^s \sum_{j \geq 2} j y_j^{(t)} \geq n - u_i = a_i + i, \tag{21}$$

where the right hand side counts the total number of vertices explored by stage i , while the left hand side counts the sum of the sizes of all hyperedges revealed by stage i , so the \geq sign accounts for possible intersections between the hyperedges.

Let w be the vertex being explored for stage $i + 1$. For $t = 1, \dots, s$ let M_t be the indicator that w is part of an (unrevealed) hyperedge in the t -th packet. Thus, $(M_t)_{1 \leq t \leq s}$ are independent conditionally given \mathcal{H}_i , and

$$\mathbb{P}(M_t = 1 | \mathcal{H}_i) = \sum_{j \geq 2} \frac{j \binom{k_j - y_j^{(t)}}{|U_i|}}, \tag{22}$$

since $k_j - y_j^{(t)}$ counts the number of hyperedges of size j still unrevealed in the t -th packet. If w is part of a hyperedge in the t -th packet, let V_t be the size of the (unique) hyperedge of that packet containing it. Then

$$\mathbb{P}(V_t = j | \mathcal{H}_i, M_t = 1) = \frac{j \binom{k_j - y_j^{(t)}}{\sum_{m \geq 2} m \binom{k_m - y_m^{(t)}}{}}}{\sum_{m \geq 2} m \binom{k_m - y_m^{(t)}}{}} \tag{23}$$

Note that when $M_t = 1$ it implies that the denominator above is non-zero and thus (23) is well defined. When $M_t = 0$ we simply put $V_t = 1$ by convention. Then we have the following almost sure inequality:

$$a_{i+1} - a_i \leq -1 + \sum_{t=1}^s M_t (V_t - 1). \tag{24}$$

This would be an equality if it were not for possible self-intersections, as hyperedges connected to w coming from different packets may share several vertices in common. In order to get a bound in the other direction, we simply truncate the $a_{i+1} - a_i$ at $n^{1/4}$. Let I_i be the indicator that among the first $n^{1/4}$ vertices to which w is connected, no self-intersection or intersection with the past occurs. Note that $\mathbb{E}(I_i) \geq p_n = 1 - n^{-1/2}$, by straightforward bounds on the birthday problem. We then have

$$(a_{i+1} - a_i) \wedge n^{1/4} \geq -1 + I_i \left(\sum_{t=1}^s M_t(V_t - 1) \wedge n^{1/4} \right). \tag{25}$$

Organisation of proof of Theorem 3.1 We will stop the exploration process once we have discovered enough vertices, or if the active set dies out, whichever comes first. We aim to show that starting from a given vertex v , with probability approximately $\theta(c)$ the cluster of v contains about order n vertices. However, we proceed in stages as different arguments are needed in order to reach so many vertices. In Step 1, we first show that the cluster contains about $(\log n)^2$ vertices with probability approximately $\theta(c)$. Then in Step 2, given that the exploration of the cluster has discovered $(\log n)^2$ vertices, we show that with high probability the exploration will in fact discover $n^{2/3}$ vertices. Finally, in Step 3 we show using the sprinkling technique that any two clusters that reach a size of about $n^{2/3}$ can be connected using only very few additional edges, which implies the result.

Main quantitative lemma We define

$$T^\downarrow := \inf\{i \geq 1 : a_i = 0\} \tag{26}$$

$$T^\uparrow := \inf\{i \geq 1 : a_i > n^{2/3}\} \tag{27}$$

We set $T = T^\uparrow \wedge T^\downarrow$. Hence our first goal (which we will show at the end of Step 2) will be to show that $T = T^\uparrow$ with probability $\theta(c)$: in fact we will show that T^\downarrow occurs before T^\uparrow or $n^{2/3}$ with probability approximately $1 - \theta(c)$. Either way, this means that the component is greater than $n^{2/3}$ with probability approximately $\theta(c)$. To do this we need to study the distribution of $a_{i+1} - a_i$; the next lemma shows that these random variables converge in distribution to a sequence of i.i.d. (possibly degenerate) random variables, uniformly for $i < T$: the limit is improper if $\sum_j \tilde{\alpha}_j < 1$.

Equivalently, the active process $|A_i|$ converges (at least for finite dimensional marginals) to the exploration process of a Galton–Watson tree whose offspring distribution is given by the limit of $a_{i+1} - a_i + 1$ and thus has a moment generating function given by $\psi(1 - \cdot, c)$.

It is perhaps surprising that the lemma below is sufficient for the proof of Theorem 3.1: the lemma below essentially only records whether a cycle is microscopic (finite) or “more than microscopic”; in particular, whether the mass of Γ comes from many small mesoscopic or fewer big cycles makes no difference.

Lemma 3.2 *For each $q_0 \in [0, 1)$, there exists some deterministic function $w : \mathbb{N} \rightarrow \mathbb{R}$ such that $w(n) \rightarrow 0$ as $n \rightarrow \infty$ with the following property:*

$$\sup_{i \geq 1} \sup_{q \in [0, q_0]} \left| \mathbb{E}[q^{a_{i+1} - a_i} | \mathcal{H}_i] - \frac{\psi(1 - q, c)}{q} \right| 1_{\{T > i\}} \leq w(n)$$

almost surely.

Proof Suppose $T > i$. From (24) we have that

$$\begin{aligned}
 q\mathbb{E}[q^{a_{i+1}-a_i}|\mathcal{H}_i] &\geq \mathbb{E}\left[q^{\sum_{t=1}^s M_t(V_t-1)}|\mathcal{H}_i\right] = \prod_{t=1}^s \mathbb{E}\left(q^{M_t(V_t-1)}|\mathcal{H}_i\right) \\
 &= \prod_{t=1}^s \left[1 - \mathbb{P}(M_t = 1|\mathcal{H}_i) \left(1 - \mathbb{E}(q^{V_t-1}|\mathcal{H}_i, M_t = 1)\right)\right].
 \end{aligned}$$

Recall from (23) that

$$\mathbb{E}(q^{V_t-1}|\mathcal{H}_i, M_t = 1) = \sum_{j \geq 2} q^{j-1} \frac{j \binom{k_j - y_j^{(t)}}{\sum_{m \geq 2} m(k_m - y_m^{(t)})}}{\sum_{m \geq 2} m(k_m - y_m^{(t)})} \geq \sum_{j \geq 2} q^{j-1} \frac{j \binom{k_j - y_j^{(t)}}{k}}{\sum_{m \geq 2} m(k_m - y_m^{(t)})}$$

and from (22) that

$$\mathbb{P}(M_t = 1|\mathcal{H}_i) = \sum_{m \geq 2} \frac{m(k_m - y_m^{(t)})}{|U_i|} \leq \sum_{m \geq 2} \frac{mk_m}{n-2n^{2/3}} \leq \frac{k}{n}(1 + 3n^{-1/3})$$

by definition of T^\uparrow . Therefore, using $1 - x \geq e^{-x-x^2}$ for all x sufficiently small,

$$\begin{aligned}
 q\mathbb{E}[q^{a_{i+1}-a_i}|\mathcal{H}_i] &\geq \prod_{t=1}^s \left[1 - \frac{k}{n}(1 + 3n^{-1/3}) \left(1 - \sum_{j \geq 2} q^{j-1} \frac{j \binom{k_j - y_j^{(t)}}{k}}{\sum_{m \geq 2} m(k_m - y_m^{(t)})}\right)\right] \\
 &\geq \prod_{t=1}^s \left\{1 - \frac{k}{n} \left(1 + 3n^{-1/3} - \sum_{j \geq 2} q^{j-1} \alpha_j\right) - \frac{k}{n}(1 + 3n^{-1/3}) \sum_{j \geq 2} q^{j-1} \frac{j y_j^{(t)}}{k}\right\} \\
 &\geq \prod_{t=1}^s \left\{1 - \frac{k}{n} \left(1 - \sum_{j \geq 2} q^{j-1} \alpha_j\right) - 3n^{-1/3} \frac{k}{n} - \frac{1}{n}(1 + 3n^{-1/3}) \sum_{j \geq 2} j y_j^{(t)}\right\} \\
 &\geq \exp \left\{-s \frac{k}{n} \left(1 - \sum_{j \geq 2} q^{j-1} \alpha_j\right) - O(n^{-1/3}) \frac{sk}{n} - O\left(s \frac{k^2}{n^2}\right)\right\}. \tag{28}
 \end{aligned}$$

Hence, since $sk/n = O(1)$ in the regime we are concerned with,

$$q\mathbb{E}[q^{a_{i+1}-a_i}|\mathcal{H}_i] \geq \psi(1 - q, c)(1 + o(1)) \geq \psi(1 - q, c) + o(1)$$

where the $o(1)$ term is non random and independent of i , and for the last inequality we have used that

$$\exp\left(-c \sum_j q^{j-1} \alpha_j\right) \rightarrow \exp\left(-c \sum_j q^{j-1} \bar{\alpha}_j\right) \tag{29}$$

which follows from the fact that $q \leq q_0 < 1$ and the dominated convergence theorem, as jk_j/k is uniformly bounded by 1. Note that the above estimate is uniform in $i \geq 1$.

For the upper bound, we use (25). Let $\epsilon_n \rightarrow 0$ sufficiently slowly that $\epsilon_n n^{1/3} \rightarrow \infty$. For concreteness take $\epsilon_n = n^{-1/6}$. Define $G := \{t \in \{1, \dots, s\} : \sum_{m \geq 2} m y_m^{(t)} \leq \epsilon_n k\}$, and let $I = G^c$. Packets $t \in I$ are the bad packets for which a significant fraction of the mass corresponding to that packet (at least ϵ_n) was already discovered at step i ; by contrast packets $t \in G$ are those for which a fraction at least $(1 - \epsilon_n)$ remains to be discovered in the exploration. In the case where the conjugacy class contains only one type of cycles, say k -cycles, then I coincides with the set of hyperedges already revealed. At the other end of the spectrum, when the conjugacy class Γ is broken down into many small cycles, then I is likely to be empty. But in all cases, $|I|$ satisfies the trivial bound $|I| \leq \frac{2n^{2/3}}{\epsilon_n k}$ by definition of T^\uparrow , and in particular

$$\frac{k|I|}{n} \leq \frac{2}{\epsilon_n n^{1/3}} \leq 2n^{-1/6} \rightarrow 0. \tag{30}$$

This turns out to be enough for our purposes.

Note that $\mathbb{E}(q^{\sum_{i=1}^s M_i(V_i-1)})$ and $\mathbb{E}(q^{n^{1/4} \wedge \sum_{i=1}^s M_i(V_i-1)})$ can only differ by at most $q^{n^{1/4}}$, which is exponentially small in $n^{1/4}$ for a fixed $q \leq q_0 < 1$, so we can neglect this difference. Then we may write, counting only hyper edges from good packets, using the fact that $1 - x \leq e^{-x}$ for all $x \in \mathbb{R}$, and (30) (recalling that I_i is the indicator of the event that no self-intersection occurs among the first $n^{1/4}$ vertices connected to w):

$$\begin{aligned} q \mathbb{E}[q^{a_{i+1}-a_i} | \mathcal{H}_i] &\leq 1 - \mathbb{E}(I_i) + \mathbb{E}(I_i) \left(q^{n^{1/4}} + \prod_{t=1}^s \left[1 - \frac{k - \sum_{m \geq 2} m y_m^{(t)}}{n} \right. \right. \\ &\quad \left. \left. \times \left(1 - \sum_{j \geq 2} q^{j-1} \frac{j(k_j - y_j^{(t)})}{k - \sum_{m \geq 2} m y_m^{(t)}} \right) \right] \right) \\ &\leq 2n^{-1/2} + q^{n^{1/4}} + \prod_{t \in G} \left[1 - \frac{k}{n} (1 - \epsilon_n) \right. \\ &\quad \left. \times \left(1 - \sum_{j \geq 2} q^{j-1} \frac{j k_j}{k(1 - \epsilon_n)} \right) \right] \\ &\leq o(1) + \exp \left\{ -s \frac{k}{n} (1 - \epsilon_n) + \frac{k}{n} |I| (1 - \epsilon_n) + s \frac{k}{n} \sum_{j \geq 2} q^{j-1} \alpha_j \right\} \\ &= o(1) + \exp \left\{ -s \frac{k}{n} + \frac{s k}{n} \sum_{j \geq 2} q^{j-1} \alpha_j \right\} (1 + o(1)) \\ &\leq \psi(1 - q, c) + o(1) \end{aligned} \tag{31}$$

where the $o(1)$ term again is non random and uniform in $i \geq 1$, but might depend on q (the last inequality again from comes from (29)). The proof is complete. \square

Lemma 3.2 above tells us that, at the level of generating functions, the distribution of $a_{i+1} - a_i$ behaves very much like a sequence of i.i.d. random variables with distribution determined by ψ , even if we don't ignore self-intersections. It is thus easy to build martingales from quantities of the form q^{a_i} , which behave as if the increments of a_i were i.i.d., at least until we reach size $n^{2/3}$. Hence this will allow us to reach a size of $n^{2/3}$ for a_i almost as if there were no self-intersections, and so with probability approximately $\theta(c)$. Fundamentally, this is because even if self-intersections do occur, they are rare and do not cause a significant loss of mass. Technically, it is easier to have a separate argument for bringing the cluster to a polylogarithmic size before using this information to show that the cluster reaches size $n^{2/3}$ with essentially the same probability. This is what we achieve in Step 1, which we are now ready for.

Step 1. We show that the cluster containing a given vertex v is at least logarithmically large with probability approximately $\theta(c)$, and furthermore the number of vertices for which this occurs is approximately $n\theta(c)$ in the sense of convergence in probability.

Lemma 3.3 *Let C_v denote the component containing v . We have that*

$$\lim_{n \rightarrow \infty} \mathbb{P}(|C_v| > (\log n)^2) = \theta(c). \tag{32}$$

Proof We start with the upper bound of (32), for which we simply make a comparison with a Galton–Watson process: to reach size $\log n$ the exploration process must survive more than a finite number of steps. More precisely, we make the following observation. Let $m \geq 1$ be some arbitrary fixed large integer, and observe $\mathbb{P}(|C_v| > (\log n)^2) \leq \mathbb{P}(|C_v| \geq m)$ trivially. Now, whether $|C_v|$ reaches size m is something that can be decided by performing the breadth-first search exploration of the cluster on a finite (at most m) number of steps: i.e., if we let $X_{i+1} = |A_{i+1} \setminus A_i|$, then a direct and crude consequence of Lemma 3.2 is that (X_1, \dots, X_m) converge to i.i.d. random variables $(\bar{X}_1, \dots, \bar{X}_m)$ (which are possibly improper, if $\sum \bar{\alpha}_j < 1$) having as generating function $\mathbb{E}(q^{\bar{X}}) = \psi(1 - q, c)$. Formally, the \bar{X}_i have the same distribution as

$$\bar{X} = \left(\sum_j (j - 1) \text{Poisson}(c\bar{\alpha}_j) \right) + \infty \cdot \text{Poisson} \left(c \left(1 - \sum_j \bar{\alpha}_j \right) \right)$$

where the Poisson random variables are independent. Consequently, if W is the total progeny of a Galton–Watson branching process with offspring distribution \bar{X}_i (note in particular that $W = \infty$ as soon as one nodes in the tree has offspring $\bar{X}_i = \infty$). We conclude that $\mathbb{P}(|C_v| \geq m) \rightarrow \mathbb{P}(W \geq m)$, and hence, taking the limsup and letting $m \rightarrow \infty$,

$$\limsup_{n \rightarrow \infty} \mathbb{P}(T^\downarrow \geq (\log n)^2) \leq \mathbb{P}(W = \infty) = \theta(c).$$

This proves the upper bound in (32).

We now discuss the lower bound to (32), which is essentially the same argument, together with the observation that self-intersections are unlikely to occur before

$(\log n)^2$ vertices have been explored. For this we can assume without loss of generality that $\theta(c) > 0$, otherwise there is nothing to prove. Let

$$T_1 = \inf \left\{ i \geq 1 : a_i > (\log n)^2 \right\};$$

we will prove the slightly stronger result that $\liminf_{n \rightarrow \infty} \mathbb{P}(T_1 < T_\downarrow) \geq \theta(c)$. (This is slightly stronger, because $|C_v|$ could in principle be greater than $(\log n)^2$ without the active set ever reaching that size). Let X_i be i.i.d. random variables with generating function given by

$$\psi_n(q) = \mathbb{E}(q^{X_i}) = \prod_{t=1}^s \left(1 - \frac{k}{n} \left(1 - \sum_j q^{j-1} \alpha_j \right) \right), \tag{33}$$

so that, by (23), $a_1 - a_0$ has the same distribution as X_1 when $A_0 = \{v\}$ (see e.g. (28) where a similar calculation is carried). We can use the random variables X_i to generate the breadth first exploration of C_v until we find a self-intersection. Thus let \tilde{Y}_i be a collection of randomly chosen vertices of $\{1, \dots, n\}$ of size X_i , and at each time step, add to the active set \tilde{A}_{i+1} the set \tilde{Y}_i and remove the currently explored vertex. Then we can couple A_i and \tilde{A}_i so that $A_i = \tilde{A}_i$ until the first time T_{inter} such that $\tilde{Y}_i \cap (\tilde{Y}_j \cup \{v\}) \neq \emptyset$ for some $i \neq j \leq T_{\text{inter}}$. Furthermore, until T_{inter} , \tilde{A}_i is the breadth-first search exploration of a branching process with offspring distribution (33). It becomes extinct with a probability q_n , and we claim that q_n satisfies $q_n \rightarrow 1 - \theta(c)$ as $n \rightarrow \infty$ by (29). Indeed, ψ_n clearly converges uniformly to $\psi(\cdot, c)$ on $[0, x_0]$ for $x_0 < 1$ by Lemma 3.2 and this is the regime we are interested in since by assumption $\theta(c) > 0$.

Hence, it is clear that if W_n is the total progeny of this branching process, then $\mathbb{P}(W_n \geq (\log n)^2) \geq \mathbb{P}(W_n = \infty) = 1 - q_n \rightarrow \theta(c)$, and combining with the argument in the upper bound on (32) we deduce that $\mathbb{P}(W_n \geq (\log n)^2) \rightarrow \theta(c)$. On the other hand, $T_1 < T_{\text{inter}}$ with probability tending to 1 as $n \rightarrow \infty$ by the birthday problem, and so in fact $\mathbb{P}(T_1 < T_\downarrow) = \mathbb{P}(W_n \geq (\log n)^2) + o(1)$, so we are done. \square

It is important to note that self-intersections may occur at the very step that a_i exceeds $(\log n)^2$ (for instance, think about the case when the conjugacy class has some of its mass coming from cycles larger than $n^{1/2}$: discovering such a cycle would immediately produce a self-intersection). Even so, the active set reaches size $(\log n)^2$ before such a self-intersection is discovered.

As announced at the beginning of Step 1, we complement this with a law of large numbers:

Lemma 3.4

$$\frac{1}{n} |\{v : |C_v| \geq (\log n)^2\}| \rightarrow \theta(c) \tag{34}$$

in probability as $n \rightarrow \infty$.

Proof Let $Z = \sum_{v=1}^n \mathbf{1}_{\{|C_v| \geq (\log n)^2\}}$, so by the previous lemma we know that $\mathbb{E}(Z)/n \rightarrow \theta$ by (32). Hence if we show that $\text{Var}(Z) \leq \varepsilon n^2$ for any $\varepsilon > 0$ and

any n sufficiently large, then (34) follows by Chebyshev’s inequality. In particular, it suffices to show that for $v \neq w \in \{1, \dots, n\}$,

$$\limsup_{n \rightarrow \infty} \text{Cov} \left(1_{\{|C_v| \geq (\log n)^2\}}, 1_{\{|C_w| \geq (\log n)^2\}} \right) \leq 0$$

or equivalently,

$$\limsup_{n \rightarrow \infty} \mathbb{P} \left(|C_v| \geq (\log n)^2, |C_w| \geq (\log n)^2 \right) \leq \theta(c)^2. \tag{35}$$

On the other hand, (35) can be proved in exactly the same way as the upper bound of (32) above: for both $|C_v|$ and $|C_w|$ to be larger than $(\log n)^2$, both must be greater than m where $m \geq 1$ is fixed. This is an event which depends on a finite number of steps (at most $2m$) in the explorations of C_v and C_w , and so can be approximated by Lemma 3.2 by the same event for two independent branching processes. Letting $m \rightarrow \infty$ finishes the proof. \square

For the rest of the proof we now assume that $c > c_\Gamma$ so that $\theta(c) > 0$. Hence fix $q \in [0, 1)$ such that $\psi(1 - q, c)/q < 1$, and note that using Lemma 3.2, we can suppose that, for some fixed $\epsilon > 0$, n is large enough so that

$$\mathbb{E} \left[q^{a_{i+1} - a_i} | \mathcal{H}_i \right] \leq (1 + \epsilon)^{-1} \tag{36}$$

almost surely on $\{T > i\}$.

Step 2. We now extrapolate the information obtained in the previous step to show that, still with probability approximately $\theta(c)$, the active set of C_v can reach a size of at least $O(n^{2/3})$. To do so we suppose our exploration from Step 1 yields an active set of size at least $(\log n)^2$ (which, as discussed, occurs with probability $\theta(c) + o(1)$). We will restart the exploration from that point on, calling this time $i = 0$ again. Hence the setup is the same as before, except that at time $i = 0$ we have $a_0 = \lfloor (\log n)^2 \rfloor$: we only keep the first $(\log n)^2$ of the active vertices discovered at time T_1 , and declare all further active vertices at time T_1 to be removed at time $i = 0$ in the exploration of Step 2.

Recall our notations for T^\downarrow and T^\uparrow in (26) and (27). Our goal in this step is to show the following control:

Lemma 3.5 *Suppose that given \mathcal{H}_0 , it is a.s. the case that $a_0 = \lfloor (\log n)^2 \rfloor$, and $r_0 \leq n^{2/3}$. Then*

$$\mathbb{P}(T^\downarrow < n^{2/3} \wedge T^\uparrow | \mathcal{H}_0) = O \left(q^{(\log n)^2} \right) = o(n^{-1}). \tag{37}$$

Proof Set $S = n^{2/3} \wedge T^\uparrow \wedge T^\downarrow$ and for $i \geq 0$, let

$$M_i := q^{a_{i \wedge S}} (1 + \epsilon)^{i \wedge S},$$

so $M = (M_i : i = 0, \dots)$ is a supermartingale in the filtration $(\mathcal{H}_0, \mathcal{H}_1, \dots)$. Observe that $S \leq n^{2/3}$ so M is bounded. Note that on the event $\{S = T^\downarrow\}$,

$$M_T = (1 + \varepsilon)^{T^\downarrow} \geq 1_{\{S=T^\downarrow\}}$$

hence by the optional stopping theorem (since M is bounded), given \mathcal{H}_0 and under the assumptions of the lemma on \mathcal{H}_0 ,

$$\begin{aligned} \mathbb{P}(S = T^\downarrow | \mathcal{H}_0) &\leq \mathbb{E}(M_S 1_{\{S=T^\downarrow\}} | \mathcal{H}_0) \\ &\leq M_0 = q^{a_0} \leq q^{(\log n)^2 - 1}, \end{aligned}$$

as desired. □

Consequently, since the error bound in Eq. (37) is $o(n^{-1})$, we deduce that if

$$\mathcal{G} = \left\{ v : |\mathcal{C}_v(s)| > n^{2/3} \right\}, \quad \tilde{\mathcal{G}} = \left\{ v : |\mathcal{C}_v(s)| > (\log n)^2 \right\}, \tag{38}$$

then $\mathcal{G} = \tilde{\mathcal{G}}$ with high probability, and hence in particular

$$\frac{|\mathcal{G}|}{n} \rightarrow \theta(c) \tag{39}$$

in probability as $n \rightarrow \infty$.

Step 3. We now show that if v and v' are two vertices such that $\mathcal{C}_v = \mathcal{C}_v(s)$ and $\mathcal{C}_{v'} = \mathcal{C}_{v'}(s)$ are both larger at time s than $n^{2/3}$ then they are highly likely to be connected at some slightly later time $s + s'$. This follows from a so-called ‘‘sprinkling’’ argument, as follows. That is, suppose we add s' packets, with

$$s' = \left\lceil \frac{Dn^{2/3} \log n}{k} \right\rceil$$

for some $D > 0$ to be chosen later on. Note that $s'k/n \rightarrow 0$ so that $(s + s')k/n \rightarrow c$. Since $s = s(n)$ is an arbitrary sequence such that $sk/n \rightarrow c$ it suffices to show that v and v' are then connected at time $s + s'$. In fact we will check that the two clusters can be connected using smaller edges than the hyperedges making each packet, as follows. For each hyperedge of size j we will only reveal a subset of $\lfloor j/2 \rfloor$ edges (of size 2) with disjoint support. Since $\lfloor j/2 \rfloor \geq j/3$ for any $j \geq 2$, this gives us at least $k/3$ edges for each packet; these are sampled uniformly at random without replacement from $\{1, \dots, n\}$. We will check that a connection occurs between the two clusters within these $s'k/3$ edges, with high probability.

Call the two clusters A and A' for simplicity; these are two arbitrary sets of size $\lfloor n^{2/3} \rfloor$ which we can assume to be disjoint otherwise there is nothing to prove. Call a packet of edges good if their intersections with each of A and A' contains at most $\lfloor n^{2/3}/2 \rfloor$ vertices, and call it bad otherwise. We reveal the edges in a given packet one

by one, sampling without replacement. Note that so long as packet of edges has not been observed to be bad, the probability that the next edge connects A and A' is at least $n^{4/3}/(16(n - k)^2) \geq n^{-2/3}/32$. (Note that if $k \leq n^{2/3}/2$ then every packet is necessarily good). Hence the probability that no connection between A and A' occurs for a good packet is at most

$$(1 - n^{-2/3}/32)^{k/3} \leq \exp\left(-\frac{kn^{-2/3}}{96}\right).$$

Now, each packet is bad independently of each other, with probability tending to 0 by Markov's inequality (since the expected intersection of a pack of edges with A is at most $|A|k/n = o(|A|)$) and hence less than $1/2$ say. So by standard Chernoff bounds on Binomial random variables, with probability at least $1 - \exp(-h \times s')$ (where $h > 0 = (1/4) \log(1/4) + (3/4) \log(3/4) + \log 2$ is a universal constant), at least $s'/4$ packs are good. Putting together these two observations, we deduce that the probability that there are no connections between A and A' after s' packs of edges have been added is at most

$$\exp\left(-\frac{kn^{-2/3}}{96} \frac{s'}{4}\right) + \exp(-hs') \leq \exp\left(-\frac{D}{400} \log n\right) + \exp(-hs'), \tag{40}$$

By choosing $D = 1201$, this is $o(n^{-3})$ at least if $k \geq n^{2/3}/2$ (so that $s' \geq 2D \log n$). However, if $k \leq n^{2/3}$, then every packet is good, and so (40) holds without the second term on the right hand side. Either way,

$$\mathbb{P}(\mathcal{C}_v(s + s') \cap \mathcal{C}_{v'}(s + s') = \emptyset) = o(n^{-3}). \tag{41}$$

Proof of Theorem 3.1 We are now ready to conclude that vertices are either in small component at time s or connected at time $s + s'$. Recall our notation $\mathcal{G} = \{v : |\mathcal{C}_v(s)| > n^{2/3}\}$. Then by (39), we know that $|\mathcal{G}|/n \rightarrow \theta(c)$ in probability as $n \rightarrow \infty$. We now aim to show that \mathcal{G} is connected at time $s + s'$, with high probability. For $v, v' \in \{1, \dots, n\}$, write $v \leftrightarrow v'$ to indicate that v is connected to v' . Then by Step 3 (more specifically, (41)),

$$\mathbb{P}\left(\bigcup_{v, v' \in \mathcal{G}} \{v \leftrightarrow v' \text{ at time } s + s'\}^c \middle| H_s\right) \leq |\mathcal{G}|^2 o(n^{-3}) \leq o(n^{-1}).$$

Hence \mathcal{G} is entirely connected at time $s + s'$ with probability tending to 1. This proves that $H_{s+s'}$ contains a component of relative size converging to $\theta(c)$ in probability. Let us now check that every other component at time $s + s'$ is small. Note that since $\mathcal{G} = \tilde{\mathcal{G}}$ with probability tending to one (where $\tilde{\mathcal{G}}$ is defined in (38)), any component disjoint from \mathcal{G} at time $s + s'$ must have been smaller than $(\log n)^2$ at time s . Since at most $s'k$ connections are added, this means that, on the event $\mathcal{G} = \tilde{\mathcal{G}}$, the maximal size of a component at time $s + s'$ disjoint from \mathcal{G} is smaller than $s'k(\log n)^2 \leq Dn^{2/3}(\log n)^3$. This shows that every other component is $O(n^{2/3}(\log n)^3)$ on an event of high probability.

The proof of Theorem 3.1 is complete, since $s + s'$ in an arbitrary sequence such that $(s + s')k/n \rightarrow c$. □

3.3 Poisson–Dirichlet structure

The renormalised cycle lengths $\mathfrak{X}(\sigma)$ of a permutation $\sigma \in \mathcal{S}_n$ is the cycle lengths of σ divided by n , written in decreasing order. In particular we have that $\mathfrak{X}(\sigma)$ takes values in

$$\Omega_\infty := \left\{ (x_1 \geq x_2 \geq \dots) : x_i \in [0, 1] \text{ for each } i \geq 1 \text{ and } \sum_{i=1}^\infty x_i = 1 \right\}. \tag{42}$$

We equip Ω_∞ with the topology of pointwise convergence. If σ_n is uniformly distributed in \mathcal{S}_n then $\mathfrak{X}(\sigma_n) \rightarrow Z$ in distribution as $n \rightarrow \infty$ where Z is known as a Poisson–Dirichlet random variable. It can be constructed as follows. Let U_1, U_2, \dots be i.i.d. uniform random variables on $[0, 1]$. Let $Z_1^* = U_1$ and inductively for $i \geq 2$ set $Z_i^* = U_i(1 - \sum_{j=1}^{i-1} Z_j^*)$. Then (Z_1^*, Z_2^*, \dots) can be ordered in decreasing size and the random variable Z has the same law as (Z_1^*, Z_2^*, \dots) ordered by decreasing size.

The next result is a generalisation of Theorem 1.1 in [24] to the case of general conjugacy classes. The proof is a simple adaptation of the proof of Schramm and we provide the details in an appendix.

Theorem 3.6 *Suppose $s = s(n)$ is such that $sk/n \rightarrow c$ as $n \rightarrow \infty$ for some $c > c_\Gamma$. Then we have that for any $m \in \mathbb{N}$*

$$\left(\frac{\mathfrak{X}_1(X_s)}{\theta(c)}, \dots, \frac{\mathfrak{X}_m(X_s)}{\theta(c)} \right) \rightarrow (Z_1, \dots, Z_m)$$

in distribution as $n \rightarrow \infty$ where $Z = (Z_1, Z_2, \dots)$ is a Poisson–Dirichlet random variable.

4 Proof of curvature theorem

4.1 Proof of the upper bound on curvature

We claim that it is enough to show the upper bound for $c > c_\Gamma$ in (15). Indeed, notice that $c \mapsto \kappa_c$ is nondecreasing. Hence let $c \leq c_\Gamma$ and suppose we know that $\limsup_{n \rightarrow \infty} \kappa_{c'} \leq \theta(c')^2$ holds for all $c' > c_\Gamma$. Then we have that $\limsup_{n \rightarrow \infty} \kappa_c \leq \theta(c')^2$ for each $c' > c_\Gamma$. Taking $c' \downarrow c_\Gamma$ and using the fact that $\lim_{c' \downarrow c_\Gamma} \theta(c') = 0$ shows that $\lim_{n \rightarrow \infty} \kappa_c = 0$.

Fix $c > c_\Gamma$ and let $t := \lfloor cn/k \rfloor$. We need to show the upper bound in (15). In other words, we wish to prove that for some $\sigma, \sigma' \in \mathcal{S}_n$

$$\liminf_{n \rightarrow \infty} \frac{W_1(X_t^\sigma, X_t^{\sigma'})}{d(\sigma, \sigma')} \geq 1 - \theta(c)^2.$$

We will choose $\sigma = \text{id}$ and $\sigma' = \tau_1 \circ \tau_2$, where τ_1, τ_2 are independent uniformly chosen transpositions. To prove the lower bound on the Kantorovich distance we use the dual representation of the distance $W_1(X, Y)$ between two random variables X, Y :

$$W_1(X, Y) = \sup\{\mathbb{E}[f(X)] - \mathbb{E}[f(Y)] : f \text{ is Lipschitz with Lipschitz constant } 1\}. \tag{43}$$

Let $f(\sigma) = d(\text{id}, \sigma)$ be the distance to the identity (using only transpositions, as usual). Then observe that f is 1-Lipschitz. It suffices to show

$$\liminf_{n \rightarrow \infty} \mathbb{E}[f(X_t^{\tau_1 \circ \tau_2})] - \mathbb{E}\left[f\left(X_t^{\text{id}}\right)\right] \geq 2(1 - \theta(c)^2). \tag{44}$$

We will now show (44) by a coupling argument. Construct the two walks $X_t^{\tau_1 \circ \tau_2}$ and X_t^{id} as follows. Let $\gamma_1, \gamma_2, \dots$ be a sequence of i.i.d. random variables uniformly distributed on Γ , independent of (τ_1, τ_2) . Using Lemma 2.5 with $\sigma_0 = \tau_1 \circ \tau_2$, which is independent of X_t^{id} , we can construct $X_t^{\tau_1 \circ \tau_2}$ as

$$X_t^{\tau_1 \circ \tau_2} = \gamma_1 \circ \dots \circ \gamma_t \circ \tau_1 \circ \tau_2.$$

Next we couple X_t^{id} by constructing it as

$$X_t^{\text{id}} = \gamma_1 \circ \dots \circ \gamma_t.$$

Thus under this coupling we have that $X_t^{\tau_1 \circ \tau_2} = X_t^{\text{id}} \circ \tau_1 \circ \tau_2$. Let $X = X_t^{\text{id}}$, then from (44) the problem reduces to showing

$$\liminf_{n \rightarrow \infty} \mathbb{E}[d(\text{id}, X_t \circ \tau_1 \circ \tau_2) - d(\text{id}, X_t)] \geq 2(1 - \theta(c)^2). \tag{45}$$

We recall that a transposition can either induce a fragmentation or a coalescence of the cycles. Indeed, a transposition involving elements from the same cycle generates a fragmentation of that cycle, and one involving elements from different cycles results in the cycles being merged. (This property is the basic tool used in the probabilistic analysis of random transpositions, see e.g. [5] or [24]). Hence either τ_1 fragments a cycle of X_t or τ_1 coagulates two cycles of X_t . In the first case, $d(\text{id}, X_t \circ \tau_1) = d(\text{id}, X_t) - 1$, and in the second case we have $d(\text{id}, X_t \circ \tau_1) = d(\text{id}, X_t) + 1$. Let F denote the event that τ_1 causes a fragmentation. Then

$$\mathbb{E}[d(\text{id}, X_t \circ \tau_1) - d(\text{id}, X_t)] = 1 - 2\mathbb{P}(F).$$

Using the Poisson–Dirichlet structure described in Theorem 3.6 it is not hard to show that $\mathbb{P}(F) \rightarrow \theta(c)^2/2$ (see, e.g., Lemma 8 in [6]). Applying the same reasoning to $X_t \circ \tau_1 \circ \tau_2$ and $X_t \circ \tau_1$ we deduce that

$$\lim_{n \rightarrow \infty} \mathbb{E}[d(\text{id}, X_t \circ \tau_1 \circ \tau_2) - d(\text{id}, X_t)] = 2(1 - \theta(c)^2)$$

from which the lower bound (45) and in turn the upper bound in (10) follow readily.

4.2 Proof of lower bound on curvature

We now assume that $c > c_\Gamma$ and turn our attention to the lower bound on the Ricci curvature, which is the heart of the proof. Throughout we let $k = |\Gamma|$ and $t = \lfloor cn/k \rfloor$. With this notation in mind we wish to prove that

$$\limsup_{n \rightarrow \infty} \sup_{\sigma, \sigma'} \frac{\mathbb{E}d(X_t^\sigma, X_t^{\sigma'})}{d(\sigma, \sigma')} \leq \alpha := 1 - \theta(c)^4$$

for some appropriate coupling of X^σ and $X^{\sigma'}$, where the supremum is taken over all σ, σ' with even distance. Note that we can make several reductions: first, by vertex transitivity we can assume $\sigma = \text{id}$ is the identity permutation. Also, by the triangle inequality (since W_1 is a distance), we can assume that $\sigma' = (i, j) \circ (\ell, m)$ is the product of two distinct transpositions. There are two cases to consider: either the supports of the transpositions are disjoint, or they overlap on one vertex. We will focus in this proof on the first case where the support of the transpositions are disjoint; that is, i, j, ℓ, m are pairwise distinct. The other case is dealt with very much in the same way (and is in fact a bit easier).

Clearly by symmetry $\mathbb{E}d(X_t^{\text{id}}, X_t^{(i,j) \circ (\ell,m)})$ is independent of i, j, ℓ and m , so long as they are pairwise distinct. Hence it is also equal to $\mathbb{E}d(X_t^{\text{id}}, X_t^{\tau_1 \circ \tau_2})$ conditioned on the event A that τ_1, τ_2 having disjoint support, where τ_1 and τ_2 are independent uniform random transpositions. This event has an overwhelming probability for large n , thus it suffices to construct a coupling between X^{id} and $X^{\tau_1 \circ \tau_2}$ such that

$$\limsup_{n \rightarrow \infty} \mathbb{E}d \left(X_t^{\text{id}}, X_t^{\tau_1 \circ \tau_2} \right) \leq 2(1 - \theta(c)^4). \tag{46}$$

Indeed, it then immediately follows from stochastic commutativity (Lemma 2.5) that the same is true with the expectation replaced by the conditional expectation given A , since the distance is bounded by two.

Next, let X be a random walk on S_n which is the composition of i.i.d. uniform elements of the conjugacy class Γ . We decompose the random walk X into a walk \tilde{X} which evolves by applying transpositions at each step as follows. For $t = 0, 1, \dots$, write out

$$X_t = \gamma_1 \circ \dots \circ \gamma_t$$

where $\gamma_1, \gamma_2, \dots$ are i.i.d. uniformly distributed in Γ . As before we decompose each step γ_s of the walk into a product of cyclic permutations, say

$$\gamma_s = \gamma_{s,1} \circ \dots \circ \gamma_{s,r} \tag{47}$$

where $r = \sum_{j \geq 2} k_j$. The order of this decomposition is irrelevant and can be chosen arbitrarily. For concreteness, we decide that we start from the cycles of smaller sizes and progressively increase to cycles of larger sizes. We will further decompose each

of these cyclic permutation into a product of transpositions, as follows: for a cycle $c = (x_1, \dots, x_j)$, write

$$c = (x_1, x_2) \circ \dots \circ (x_{j-1}, x_j).$$

This allows to break any step γ_s of the random walk X into a number

$$\rho := \sum_j (j - 1)k_j$$

of elementary transpositions, and hence we can write

$$\gamma_s = \tau_s^{(1)} \circ \dots \circ \tau_s^{(\rho)} \tag{48}$$

where $\tau_s^{(j)}$ are transpositions. Note that the vectors $(\tau_s^{(i)}; 1 \leq i \leq \rho)$ in (48) are independent and identically distributed for $s = 1, 2, \dots$ and for a fixed s and $1 \leq i \leq \rho$, $\tau_s^{(i)}$ is a uniform transposition, by symmetry. However it is important to observe that the transpositions $\tau_s^{(i)}$ are *not* independent. Nevertheless, they obey a crucial conditional uniformity which we explain now. First we have to differentiate between the set of times when a new cycle starts and the set of times when we are continuing an old cycle.

Definition 4.1 (*Refreshment Times*) We call a time s a refreshment time if s is of the form $s = \rho\ell + \sum_{j=2}^m (j - 1)k_j$ for some $\ell \in \mathbb{N} \cup \{0\}$ and $m \in \mathbb{N} \setminus \{1\}$.

We see that s is a refreshment time if the transposition being applied to \tilde{X} at time s is the start of a new cycle. Using this we can describe the law of the transpositions being applied to \tilde{X} .

Proposition 4.1 (Conditional Uniformity) *For $s \in \mathbb{N}$ and $i \leq \rho$, the conditional distribution of $\tau_s^{(i)}$ given $\tau_s^{(1)}, \dots, \tau_s^{(i-1)}$ can be described as follows. We write $\tau_s^{(i)} = (x, y)$ and we will distinguish between the first marker x and the second marker y . There are two cases to consider:*

- (i) $s\rho + i$ is a refreshment time and thus $\tau_s^{(i)}$ corresponds to the start of a new cycle
- (ii) $s\rho + i$ is not a refreshment time and so $\tau_s^{(i)}$ is the continuation of a cycle.

In case (i) x is uniformly distributed on $S_i := \{1, \dots, n\} \setminus \text{Supp}(\tau_s^{(1)} \circ \dots \circ \tau_s^{(i-1)})$ and y is uniformly distributed on $S_i \setminus \{x\}$. In case (ii) x is equal to the second marker of $\tau_s^{(i-1)}$ and y is uniformly distributed in S_i .

Note that in either case, the second marker y is conditionally uniformly distributed among the vertices which have not been used so far. This conditional independence property is completely crucial, and allows us to make use of methods (such as that of Schramm [24]) developed initially for random transpositions) for general conjugacy classes, so long as $|\Gamma| = o(n)$. Indeed in that case the second marker y itself is not very different from a uniform random variable on $\{1, \dots, n\}$.

We will study this random walk using this new *transposition time scale*. We thus define a process $\tilde{X} = (\tilde{X}_u : u = 0, 1, \dots)$ as follows. Let $u \in \{0, 1, \dots\}$ and write $u = s\rho + i$ where s, i are nonnegative integers and $i < \rho$. Then define

$$\tilde{X}_u := X_s \circ \tau_{s+1}^{(1)} \circ \dots \circ \tau_{s+1}^{(i)}. \tag{49}$$

Thus it follows that for any $s \geq 0$, $\tilde{X}_{s\rho} = X_s$. Notice that \tilde{X} evolves by applying successively transpositions with the above mentioned conditional uniformity rules.

Now consider our two random walks, X^{id} and $X^{\tau_1 \circ \tau_2}$ respectively, started respectively from id and $\tau_1 \circ \tau_2$, and let \tilde{X}^{id} and $\tilde{X}^{\tau_1 \circ \tau_2}$ be the associated processes constructed using (49), on the transposition time scale. Thus to prove (46) it suffices to construct an appropriate coupling between $\tilde{X}_{t\rho}^{\text{id}}$ and $\tilde{X}_{t\rho}^{\tau_1 \circ \tau_2}$. Next, recall that for a permutation $\sigma \in \mathcal{S}_n$, $\mathfrak{X}(\sigma)$ denotes the renormalised cycle lengths of σ , taking values in Ω_∞ defined in (42). The walks \tilde{X}^{id} and $\tilde{X}^{\tau_1 \circ \tau_2}$ are invariant by conjugacy and hence both are distributed uniformly on their conjugacy class. Thus ultimately it will suffice to couple $\mathfrak{X}(\tilde{X}_{t\rho}^{\text{id}})$ and $\mathfrak{X}(\tilde{X}_{t\rho}^{\tau_1 \circ \tau_2})$.

Fix $\delta > 0$ and let $\Delta = \lceil \delta^{-9} \rceil$. Define

$$\begin{aligned} s_1 &= \lfloor (cn - \delta^{-9})/k \rfloor \rho \\ s_2 &= s_1 + \Delta \\ s_3 &= t\rho. \end{aligned}$$

Our coupling consists of three intervals $[0, s_1]$, $(s_1, s_2]$ and $(s_2, s_3]$.

Let us informally describe the coupling before we give the details. In what follows we will couple the random walks \tilde{X}^{id} and $\tilde{X}^{\tau_1 \circ \tau_2}$ such that they keep their distance constant during the time intervals $[0, s_1]$ and $(s_2, s_3]$. In particular we will see that at time s_1 , the walks \tilde{X}^{id} and $\tilde{X}^{\tau_1 \circ \tau_2}$ will differ by two independently uniformly chosen transpositions. Thus at time s_1 most of the cycles of \tilde{X}^{id} and $\tilde{X}^{\tau_1 \circ \tau_2}$ are identical but some cycles may be different. We will show that given that the cycles that differ at time s_1 are all reasonably large, then we can reduce the distance between the two walks to zero during the time interval $(s_1, s_2]$. Otherwise if one of the differing cycles is not reasonably large, then we couple the two walks to keep their distance constant during the time interval $[0, s_1]$, $(s_1, s_2]$ and $(s_2, s_3]$.

More generally, our coupling has the property that $d(X_t^{\text{id}}, X_t^{\tau_1 \circ \tau_2})$ is uniformly bounded, so that it will suffice to concentrate on events of high probability in order to get a bound on the L^1 -Kantorovich distance $W(X_t^{\text{id}}, X_t^{\tau_1 \circ \tau_2})$.

4.2.1 Coupling for $[0, s_1]$

First we describe the coupling during the time interval $[0, s_1]$. Let $\tilde{X} = (\tilde{X}_s : s \geq 0)$ be a walk with the same distribution as \tilde{X}^{id} , independent of the two uniform transpositions τ_1 and τ_2 . Then we have that by Lemma 2.5 for any $s \geq 0$, $\tilde{X}_s^{\tau_1 \circ \tau_2}$ has the same distribution as $\tilde{X}_s \circ \tau_1 \circ \tau_2$. Thus we can couple $\mathfrak{X}(\tilde{X}_{s_1}^{\text{id}})$ and $\mathfrak{X}(\tilde{X}_{s_1}^{\tau_1 \circ \tau_2})$ such that

$$\begin{aligned} \mathfrak{X}(\tilde{X}_{s_1}^{\text{id}}) &= \mathfrak{X}(\tilde{X}_{s_1}) \\ \mathfrak{X}(\tilde{X}_{s_1}^{\tau_1 \circ \tau_2}) &= \mathfrak{X}(\tilde{X}_{s_1} \circ \tau_1 \circ \tau_2). \end{aligned} \tag{50}$$

4.2.2 Coupling for $(s_1, s_2]$

For $s \geq 0$ define $\bar{X}_s = \mathfrak{X}(\tilde{X}_{s+s_1}^{\text{id}})$ and $\bar{Y}_s = \mathfrak{X}(\tilde{X}_{s+s_1}^{\tau_1 \circ \tau_2})$. Here we will couple \bar{X}_s and \bar{Y}_s for $s = 0, \dots, \Delta$. During this time we aim to show that the discrepancies between \bar{X}_0 and \bar{Y}_0 resulting from performing the transpositions τ_1 and τ_2 at the end of the previous phase can be resolved. Our main tool for doing this will be a variant of a coupling of Schramm [24], which was already used in [6].

At each step s we try to create a matching between \bar{X}_s and \bar{Y}_s by matching an element of \bar{X}_s to at most one element of \bar{Y}_s of the same size. At any time s there may be several entries that cannot be matched. By parity the combined number of unmatched entries is an even number, and observe that this number cannot be equal to two. Now $\tilde{X}_{s_1}^{\text{id}}$ and $\tilde{X}_{s_1}^{\tau_1 \circ \tau_2}$ differ by two transpositions as can be seen from (50). This implies that in particular initially (i.e., at the beginning of $(s_1, s_2]$), there are four, six or zero unmatched entries between \bar{X}_0 and \bar{Y}_0 .

Fix $\delta > 0$ and let $A(\delta)$ denote the event that the smallest unmatched entry between \bar{X}_0 and \bar{Y}_0 has size greater than $\delta > 0$. We will show that on the event $A(\delta)$ we can couple the walks such that $\bar{X}_\Delta = \bar{Y}_\Delta$ with high probability. On the complementary event $A(\delta)^c$, couple the walks so that their distance remains $O(1)$ during the time interval $(s_1, s_2]$, similar to the coupling during $[0, s_1]$.

It remains to define the coupling during the time interval $(s_1, s_2]$ on the event $A(\delta)$. We begin by estimating the probability of $A(\delta)$.

Lemma 4.2 *For any $c > 1$ and $\delta > 0$,*

$$\liminf_{n \rightarrow \infty} \mathbb{P}(A(\delta)) \geq [\theta(c)(1 - p(\delta))]^4.$$

where $p(\delta) \rightarrow 0$ as $\delta \rightarrow 0$.

Proof Recall that by construction \bar{X}_0 and \bar{Y}_0 only differ because of the two transpositions τ_1 and τ_2 appearing in (50).

Recall the hypergraph $H_{s_1/\rho}$ on $\{1, \dots, n\}$ defined in the beginning of Sect. 3.1. Since $c > c_\Gamma$, by Theorem 3.1, $H_{s_1/\rho}$ has a (unique) giant component with high probability. Let A_1 be the event that the four points composing the transpositions τ_1, τ_2 fall within the largest component of the associated hypergraph $H_{s_1/\rho}$. Since the relative size of the giant component converges in probability $\theta(c)$ by Lemma 3.1, note that $\mathbb{P}(A_1) \rightarrow \theta(c)^4$.

Furthermore, it follows from Theorem 3.6 that conditionally on the event A_1 , the asymptotic relative size of the cycles containing the four points making the transpositions τ_1, τ_2 can be thought of as the size of four independent samples from a Poisson-Dirichlet distribution, multiplied by $\theta(c)$. Hence the lemma is proved with $p(\delta)$ being the probability that one of the four samples has a size smaller than $\delta/\theta(c)$. Clearly $p(\delta) \rightarrow 0$ so the result is proved. \square

Recall that the transpositions which make up the walks \tilde{X}^{id} and $\tilde{X}^{\tau_1 \circ \tau_2}$ obey what we called conditional uniformity in Proposition 4.1. For the duration of $(s_1, s_2]$ we will assume the *relaxed conditional uniformity* assumption, which we describe now.

Definition 4.2 (*Relaxed Conditional Uniformity*) For $s = s_1 + 1, \dots, s_2$ suppose we apply the transposition (x, y) at time s . Then

- (i) if s is a refreshment time then x is chosen uniformly in $\{1, \dots, n\}$,
- (ii) if s is not a refreshment time then x is taken to be the second marker of the transposition applied at time $s - 1$.

In both cases we take y to be uniformly distributed on $\{1, \dots, n\} \setminus \{x\}$.

In making the relaxed conditional uniformity assumption we are disregarding the constraints on (x, y) given in Proposition 4.1. However the probability we violate this constraint at any point during the interval $(s_1, s_2]$ is at most $2(s_2 - s_1)\rho/n = 2\Delta k/n \rightarrow 0$; and on the event that this constraint is violated the distance between the random walks can increase by at most $(s_2 - s_1) = \Delta$. Hence we can without a loss of generality assume that during the interval $(s_1, s_2]$ both \tilde{X}^{id} and $\tilde{X}^{\tau_1 \circ \tau_2}$ satisfy the relaxed conditional uniformity assumption.

Now we show that on the event $A(\delta)$ we can couple the walks such that $\bar{X}_\Delta = \bar{Y}_\Delta$ with high probability. The argument uses a coupling of Berestycki et al. [6], itself a variant of a beautiful coupling introduced by Schramm [24]. We first introduce some notation. Let

$$\Omega_n := \left\{ (x_1 \geq \dots \geq x_n) : x_i \in \{0/n, 1/n, \dots, n/n\} \text{ for each } i \leq n \text{ and } \sum_{i \leq n} x_i = 1 \right\}.$$

Notice that the walks \bar{X} and \bar{Y} both take values in Ω_n .

Marginal evolution Let us describe the evolution of the random walk $\bar{X} = (\bar{X}_s : s = 0, 1, \dots)$. Suppose that $s \geq 0$ and $\bar{X}_s = (x_1, \dots, x_n)$. Now imagine the interval $(0, 1]$ tiled using the intervals $(0, x_1], \dots, (0, x_n]$ (the specific tiling rule does not matter). Initially for $s = 0$, and more generally if s is a refreshment time, we select $u \in \{1/n, \dots, n/n\}$ uniformly at random and then call the tile that contains u the **marked tile**. If $s \geq 1$ is not a refreshment time then the marked tile is the one containing the second marker y of Proposition 4.1 from the previous step. Either way, we have a distinguished tile (the tile containing the ‘first marker’ at the beginning of each step $s = 0, 1, \dots$

We now describe the marginal evolution of this tiling for one step. In fact this evolution takes as an input a tiling \bar{X}_s and a marked tile I . The output will be another tiling \bar{X}_{s+1} and a new marked tile for the next step. Let I be the tile containing the first marker at the beginning of the step, and place I first from left. (I represents the cycle containing the first marker u and we imagine that u is the leftmost point of that tile, i.e., in position $1/n$). Select $v \in \{2/n, \dots, n/n\}$ uniformly at random and let I' be the tile that v falls into. Then there are two possibilities:

- if $I' \neq I$ then we merge the tiles I and I' . The new tile we created is now marked for the next step.
- If $I = I'$ then we split I into two fragments, corresponding to where v falls. Thus, one of size $v - 1/n$ and the other of size $|I| - (v - 1/n)$. The rightmost one of these two tiles, containing v , is now marked for the next step. Now \bar{X}_{s+1} is the sizes of the tiles in the new tiling we have created (with additional reordering of tiles in decreasing order).

This defines a transformation $T(\bar{X}_s, I, v)$. The evolution described above has the law of the projection onto Ω_n of \bar{X} . Indeed, suppose we apply the transposition (x, y) to \bar{X}_s in order to obtain \bar{X}_{s+1} . The marked tile at time s corresponds to the cycle of \bar{X}_s containing x : if s is a refreshment time then $x \in \{1, \dots, n\}$ is chosen uniformly, otherwise x is the second marker from the previous step.

Coupling We now recall the coupling of [6]. Let $s \geq 0$. Suppose that $\bar{X}_s = \bar{X} = (x_1, \dots, x_n)$ and $\bar{Y}_s = \bar{Y} = (y_1, \dots, y_n)$. Then we can differentiate between the entries that are *matched* and those that are *unmatched*: we say that two entries from \bar{X} and \bar{Y} are matched if they are of identical size. Our goal will be to create as many matched parts as possible and as quickly as possible. When putting down the tilings \bar{X} and \bar{Y} , associated with \bar{X} and \bar{Y} respectively, we will do so in such a way that all matched parts are to the right of the interval $(0, 1]$ and the unmatched parts occupy the left part of the interval.

Let $I_{\bar{X}}$ and $I_{\bar{Y}}$ be the respective marked tiles of the tilings \bar{X} and \bar{Y} at some step $s \geq 0$, and let \hat{X}, \hat{Y} be the tiling which is the reordering of \bar{X}, \bar{Y} in which $I_{\bar{X}}$ and $I_{\bar{Y}}$ have been put to the left of the interval $(0, 1]$. We assume that at the start of the step, either $I_{\bar{X}}$ and $I_{\bar{Y}}$ are both matched to each other, or they are both unmatched. (We will then verify that this property is preserved by the coupling). Let $a = |I_{\bar{X}}|$ and let $b = |I_{\bar{Y}}|$ be the respective lengths of the marked tiles, and assume without loss of generality that $a < b$. Let $v \in \{2/n, \dots, n/n\}$ be chosen uniformly. We will apply $T(\bar{X}, I_{\hat{X}}, v)$ to \hat{X} as we did before, and obtain \bar{X}_{s+1} . To obtain \bar{Y}_{s+1} we will also apply the transformation T to it, but with an other uniform random variable $v' \in \{2/n, \dots, n/n\}$ which may differ from v . To construct v' we proceed as follows.

If $I_{\hat{X}}$ is matched (so that $I_{\hat{Y}}$ is matched to it by assumption) then we take $v' = v$, as in the coupling of Schramm [24]. In the case when $I_{\hat{X}}$ is unmatched (which also implies that $I_{\hat{Y}}$ is unmatched), we apply to v a measure-preserving map Φ , defined as follows: for $w \in \{2/n, \dots, n/n\}$ consider the map

$$\Phi(w) = \begin{cases} w & \text{if } w > b \text{ or if } 2/n \leq w \leq \gamma_n + 1/n, \\ w - \gamma_n & \text{if } a < w \leq b, \\ w + b - a & \text{if } \gamma_n + 1/n < w \leq a, \end{cases} \tag{51}$$

where $\gamma_n := \lceil an/2 - 1 \rceil/n$. (This is contrast with Schramm’s original coupling, where $v' = v$ no matter what). See Fig. 1 (top right corner) for an illustration of Φ , from which it should be clear in particular that Φ is a bijection and hence measure-preserving; this is easy to check. Thus letting $v' = \Phi(v)$ we have that v' has the correct marginal

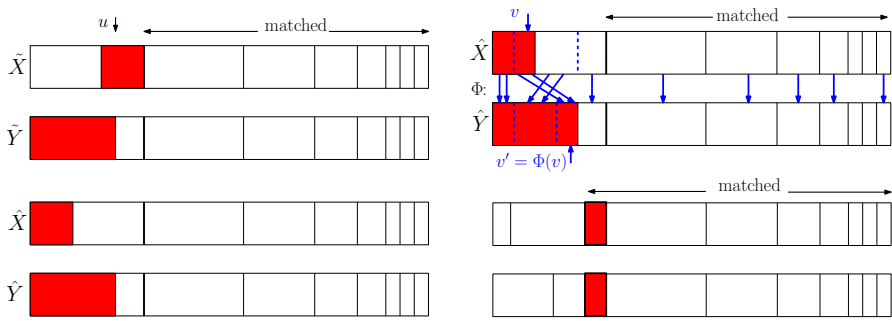


Fig. 1 One step of the evolution under the coupling between \tilde{X} and \tilde{Y} . The red entries represent the marked tiles. Left: the tilings are rearranged so that both marked tiles (which are here both unmatched) are to the left. Right: the second marker v falls, and defines a marker $v' = \Phi(v)$. In this case we get a fragmentation in both copies. The new marked tile is the rightmost fragment thus created; here they are both matched to each other, verifying in this case the consistency property of the coupling (Lemma 4.3). The total number of unmatched entries has been left unchanged, though. And while the smallest unmatched entry has decreased in size, this has not been by more than a factor of two. This is the content of the key Lemma 4.4 (colour figure online)

distribution and thus so does $\tilde{Y}_{s+1} = T(\tilde{Y}_s, I_{\tilde{Y}}, v')$. The rest of Fig. 1 illustrates the various steps in the coupling as well as the content of Lemmas 4.3 and 4.4, as explained in the caption.

Before checking that the coupling is well defined (in the sense that our assumption on the marked tiles, which are needed for the definition of the coupling, remains true throughout), we briefly add a few words of motivation for this definition.

Motivation for the coupling The coupling defined above is, as already mentioned above, the same as the one used in [6], which is a modification of a coupling due to Schramm [24]. In Schramm’s original coupling, the map Φ was taken to be the identity, which is natural enough. However this leads to the undesirable property that it is possible for very small unmatched pieces to appear; once these small unmatched pieces appear they remain in the system for a very long time which could prevent coupling. The reason for introducing the map Φ here and in [6] (where it was one of the main innovations) is that it prevents the occurrence of small unmatched pieces: as we will see in Lemma 4.4, the crucial property is that the worst thing that can happen is for the smallest unmatched piece to become smaller by a factor of two, and this only happens with small probability. This means all unmatched pieces remain relatively large, and so they disappear quite quickly (leading in turn to a coupling of the two copies). We start by a proof that the coupling is well defined:

Lemma 4.3 *At the end of a step the two marked tiles are either matched to each other or both unmatched.*

Proof The proof consists of examining several cases. If the first marker u was in a matched tile, then whether v falls in the matched or unmatched part, the property holds (if v is unmatched then we attach two unmatched tiles to two matched tiles, so they

both become unmatched. If it falls in the matched part, either two tiles of the same size are being attached, or the marked tile splits into two tiles of same size, and the rightmost piece which is the new marked tile matches in both copies).

A similar analysis can be done if the first marked tile was unmatched. The only case which requires an observation is if v falls in the same tile as $I_{\bar{x}}$ (where we assume, as in the figure, that this is the sorter of the two unmatched pieces), then if v falls in the first half of the tile this results in two matched pieces which are unmarked for the next step and two unmatched pieces which are both marked. If however v falls in the second half of $I_{\bar{x}}$ then this results in two matched pieces which are marked, and two unmatched pieces which are not marked. (Recall that the marked tile at the next step is the one containing the rightmost fragment). \square

This coupling has several remarkable deterministic properties, as already observed in [6]. Chief among those is the fact that **the number of unmatched entries can only decrease**. Unmatched entries disappear when they are coalesced. In particular they disappear quickly when their size is reasonably large. Hence it is particularly desirable to have a coupling in which unmatched components stay large. The second crucial property of this coupling is that it does not create arbitrarily small unmatched entries: even when unmatched entry is fragmented, the size of the smallest unmatched entry cannot decrease by more than a factor of two. (As these properties hold deterministically given the marked tiles, they do not need to be proved again). A direct consequence of these properties is the following lemma, which is Lemma 19 from [6].

Lemma 4.4 *Let U be the size of the smallest unmatched entry in two partitions $\bar{x}, \bar{y} \in \Omega_n$, let \bar{x}', \bar{y}' be the corresponding partitions after one transposition of the coupling, and let U' be the size of the smallest unmatched entry in \bar{x}', \bar{y}' . Assume that $2^j \leq U < 2^{j+1}$ for some $j \geq 0$. Then it is always the case that $U' \geq (1/n)\lfloor nU/2 \rfloor$, and moreover,*

$$\mathbb{P}(U' \leq 2^j) \leq 2^{j+2}/n.$$

Finally, the combined number of unmatched parts may only decrease.

Remark 4.5 In particular, it holds that $U' \geq 2^{j-1}/n$.

We now explain our strategy. On $A(\delta)$ we will expect that the unmatched components will remain of a size roughly of order at least δ for a while. In fact we will show that they will stay at least as big as $O(\delta^2)$ for a long time. Unmatched entries disappear when they are merged together. If all unmatched entries are of size at least δ^2 , we will see that with probability at least δ^8 , we have a chance to reduce the number of unmatched entries in every 3 steps. Then a simple argument shows that after time $\Delta = \lceil \delta^{-9} \rceil$, \bar{X}_Δ and \bar{Y}_Δ are perfectly matched with a probability tending to one as $\delta \rightarrow 0$.

Lemma 4.6 *There is δ_0 such that if $\delta < \delta_0$, during $[0, \Delta]$, both \bar{X}_s and \bar{Y}_s always have an entry of size greater than $\delta\theta(c)$ with probability at least $1 - 2\delta^{1/2}$ for all n sufficiently large.*

Proof Let $\delta_0 > 0$ be such that $(1 - \delta)^{9!} \geq \delta^{1/2}$ for all $\delta \leq \delta_0$; we may assume without loss of generality that $\delta \leq \delta_0$. Let $Z = (Z_1, \dots)$ be a Poisson-Dirichlet random variable on Ω_∞ and let (Z_1^*, \dots) denote the size biased ordering of Z . Recall that Z_1^* is uniformly distributed over $[0, 1]$, Z_2^* is uniformly distributed on $[0, 1 - Z_1^*]$, and so on. For the event $\{Z_1 \leq \delta\}$ to occur it is necessary that $Z_1^* \leq \delta, Z_2^* \leq \delta/(1 - \delta), \dots, Z_{10}^* \leq \delta/(1 - \delta)^9$. This has probability at most $\delta^{10}/(1 - \delta)^{9!}$. Thus

$$\mathbb{P}(Z_1 \leq \delta) \leq \frac{\delta^{10}}{(1 - \delta)^{9!}} \leq \delta^{9+1/2}.$$

Summing over $\Delta = \lceil \delta^{-9} \rceil$ steps we see that the expected number of times during the interval $[0, \Delta]$ such that \bar{X}_s or \bar{Y}_s don't have a component of size at least $\theta(c)\delta n$ is less than $\delta^{1/2}$ as $n \rightarrow \infty$ and is thus less than $2\delta^{1/2}$ for n sufficiently large, by Theorem 3.6 (note that we can apply the result because this calculation involves only a finite number of components). The result follows. \square

We now check that all unmatched components really do stay greater than δ^2 during $[0, \Delta]$. Let T_δ denote the first time s that either \bar{X}_s or \bar{Y}_s have no cycles greater than $\delta\theta(c)n$ (suppose without loss of generality that δ is small enough that $\delta^2 \leq \delta\theta(c)$).

Lemma 4.7 *On $A(\delta)$, for all $s \leq T_\delta \wedge \Delta$, all unmatched components stay greater than δ^2 with probability at least $1 - O(\delta)$, where the constant implied in $O(\delta)$ can depend on c but not on δ .*

Proof Say that a number $x \in [0, 1]$ is in scale j if $2^j/n \leq x < 2^{j+1}/n$. For $s \geq 0$, let $U(s)$ denote the scale of the smallest unmatched entry of \bar{X}_s, \bar{Y}_s . Let j_0 be the scale of δ , and let j_1 be the integer immediately above the scale of δ^2 .

Suppose for some time $s \leq T_\delta$, we have $U(s) = j$ with $j_1 \leq j \leq j_0$, and the marked tile at time s corresponds to the smallest unmatched entry. Then after this transposition we have $U(s + 1) \geq j - 1$ by the properties of the coupling (Lemma 4.4). Moreover, $U(s + 1) = j - 1$ with probability at most $r_j = 2^{j+2}/n$. Furthermore, since $s \leq T_\delta$, we have that this marked tile merges with a tile of size at least $\theta(c)\delta$ with probability at least $\theta(c)\delta$ after the transposition. We call the first occurrence a *failure* and the second a *mild success*.

Once a mild success has occurred, there may still be a few other unmatched entries in scale j , but no more than five since the total number of unmatched entries is decreasing, and there were at most six initially. And therefore if six mild successes occur before a failure, we are guaranteed that $U(s + 1) \geq j + 1$. We call such an event a *good success*, and note that the probability of a good success, given that $U(s)$ changes scale, is at least $p_j = 1 - 6r_j/(r_j + \theta(c)\delta)$. We call $q_j = 1 - p_j$.

Let $\{q_i\}_{i \geq 0}$ be the times at which the smallest unmatched entry changes scale, with q_0 being the first time the smallest unmatched entry is of scale j_0 . Let $\{U_i\}$ denote the scale of the smallest unmatched entry at time q_i . Introduce a birth–death chain on the integers, denoted v_n , such that $v_0 = j_0$ and

$$\mathbb{P}(v_{n+1} = j - 1 | v_n = j) = \begin{cases} 1 & \text{if } j = j_0 \\ 0 & \text{if } j = j_1 \\ q_j & \text{otherwise,} \end{cases} \tag{52}$$

and

$$\mathbb{P}(v_{n+1} = j + 1 | v_n = j) = \begin{cases} p_j, & j > j_1 \\ 0, & j = j_1. \end{cases} \tag{53}$$

Then it is a consequence of the above observations that $(U_i, i \geq 1)$ is stochastically dominating $(v_i, i \geq 1)$ for $s \leq T_\delta$. Set $\tau_j = \min\{n > 0 : v_n = j\}$. An analysis of the birth–death chain defined by (52), (53) gives that

$$\mathbb{P}^{j_0}(\tau_{j_1} < \tau_{j_0}) = \frac{1}{\sum_{j=j_1+1}^{j_0} \prod_{m=j}^{j_0-1} \frac{p_m}{q_m}} \leq \prod_{j=j_1+1}^{j_0-1} \frac{q_j}{p_j}$$

(see, e.g., Theorem (3.7) in Chapter 5 of [11]). Note also that $q_j \leq 6r_j/(\delta\theta(c))$, so $q_j/p_j \lesssim r_j/(\delta\theta(c))$. Moreover all the r_j in this product satisfy $r_j \lesssim \delta$. Thus, by considering the 10 terms with lowest index in the product above (and note that for $\delta > 0$ small enough, there are at least 10 terms in this product), we deduce that $\mathbb{P}^{j_0}(\tau_{j_1} < \tau_{j_0})$ decays faster than $O(\delta)^{10}$. Since $T_\delta \wedge \Delta \leq \Delta = \lceil \delta^{-9} \rceil$ we conclude that the probability that $U(s) = j_1$ before $T_\delta \wedge \Delta$ is at most $O(\delta)$. \square

We are now going to prove that on the event $A(\delta)$, after time Δ there are no unmatched entries with probability tending to one as $n \rightarrow \infty$ and $\delta \rightarrow 0$. The basic idea is again to exploit that there are initially at most six unmatched parts, and this number cannot increase. We need a few preparatory lemmas which construct a scenario which lead to a coagulation of two unmatched entries in each copy, in three steps. Let T'_δ be the first time one of the unmatched entries is smaller than δ^2 . Let \mathcal{F}_s denote the filtration generated by $(\bar{X}_1, \dots, \bar{X}_s)$ including the marked tiles at the end of each step up to time s . Let \mathcal{K}_s be the event that step s results in two unmatched entries being merged in both copies, so our first goal (achieved in Lemma 4.10) will be to get a lower bound on the probability of \mathcal{K}_s .

Step 1. We show that with good probability both marked tiles are unmatched at the end of a step (and thus also at the beginning of the next step).

Lemma 4.8 *Let \mathcal{M}_s be the event that at the end of step s , both marked tiles are unmatched. Then*

$$\mathbb{P}(\mathcal{M}_s \cup \mathcal{K}_s | \mathcal{F}_{s-1}; T'_\delta \geq s) \geq \delta^2.$$

Proof Let u, v be the two markers for step s . If the tile containing u was matched, then it suffices for v to fall in an unmatched tile (then \mathcal{M}_s occurs), which occurs again with probability at least δ^2 . If however the tile containing u was unmatched, the copy which contains the smallest of these two unmatched tiles necessarily contains at least another unmatched tile. It then suffices for v to fall in that tile. Indeed if we are very lucky and the other copy also just happen to have two unmatched entries this might lead to a reduction in the number of unmatched entries, in which case \mathcal{K}_s has occurred. Otherwise we have simply shuffled the unmatched entries and v is now in an unmatched entry, so \mathcal{M}_s holds. Either way, the conditional probability is at least δ^2 . \square

Step 2. We show that if the marked tiles are unmatched, with good probability we can get to a “balanced configuration” where both copies contain at least two unmatched tiles, and that the marked tiles at the end of the step are both unmatched.

Lemma 4.9 *Suppose s is not a refreshment time. Let \mathcal{B}_s denote the event that \bar{X}_s and \bar{Y}_s contain at least two unmatched entries each, and that the second marker is in one of these unmatched tiles for both \bar{X}_s, \bar{Y}_s at the end of the step (i.e., \mathcal{M}_s holds). Then*

$$\mathbb{P}(\mathcal{B}_s \cup \mathcal{K}_s | \mathcal{F}_{s-1}; T'_\delta \geq s; \mathcal{M}_{s-1}) \geq \delta^2. \tag{54}$$

Suppose now that s is a refreshment time. Then

$$\mathbb{P}(\mathcal{B}_s | \mathcal{F}_{s-1}; T'_\delta \geq s) \geq \delta^4. \tag{55}$$

Proof Let u, v be the two markers for step s . Suppose first that s is not a refreshment time, so we aim to prove (54). We treat several cases, according to whether \mathcal{B}_{s-1} holds or not. We start by assuming that \mathcal{B}_{s-1} does not hold. The idea is that in that case, at time $s - 1$, one copy (say \bar{Y}_{s-1}) has one unmatched entry, while the other one has at least three. It then suffices to fragment the unmatched entry in \bar{Y}_{s-1} and to coagulate the other two entries in \bar{X}_{s-1} . Since \mathcal{M}_{s-1} holds, and s is not a refreshment time, it suffices for v (the marker corresponding to the copy which has the smallest unmatched entry, which is necessarily \bar{X}_{s-1}) to fall in any of the other unmatched entries of \bar{X}_{s-1} : this necessarily results in a balanced configuration. Note also that this always results in both marked tiles to be unmatched at the end of the step, so \mathcal{B}_s indeed holds in that case. Moreover, this event has probability at least δ^2 since $T'_\delta \geq s$.

Suppose now that \mathcal{B}_{s-1} holds. Then let us show directly \mathcal{K}_s can occur with good probability. Indeed, if the second marker $v' = \Phi(v)$ (this is the marker associated with the copy, say \bar{Y}_s , that contains the larger of the two marked unmatched tiles) falls in another unmatched tile of \bar{Y}_s , then in this case a coagulation of two unmatched entries is guaranteed to occur in both copies. Hence \mathcal{K}_s occurs with probability at least δ^2 . Either way, (54) is proved.

Now suppose that s is a refreshment time. In that case it suffices to require that the first marker falls in an unmatched tile (which occurs with probability δ^2) and from then on we argue exactly as in the proof of (54) to obtain a proof of (55). All in all the lemma is proved. □

We point out that, combining Lemmas 4.8 and 4.9, regardless of whether s is a refreshment time, $\mathbb{P}(\mathcal{B}_s | \mathcal{F}_{s-1}) \geq \delta^4$.

Step 3. Having reached a balanced configuration with one marked unmatched entry in both copies, we show that a coagulation of two unmatched entries in both copies has a good chance of occurring. In that case, the number of unmatched entries has decreased by two or four.

Lemma 4.10 *We have*

$$\mathbb{P}(\mathcal{K}_s | \mathcal{F}_{s-1}; \mathcal{B}_{s-1}; T'_\delta \geq s) \geq \delta^4.$$

Proof We again need to distinguish between the cases where s is refreshment time or not. If not, then since \mathcal{B}_{s-1} holds, then the first marker is in an unmatched tile for both copies. If the second marker v' corresponding to the copy with the larger of these two unmatched tiles falls in a different unmatched tile (which has probability at least δ^2) then a coagulation is guaranteed to occur in both copies so \mathcal{K}_s holds.

If s is refreshment time, then the same argument applies, but the first marker must first fall in an unmatched component (which has probability at least δ^2 since $T'_\delta \geq s$). This gives a lower bound of δ^4 on the probability of \mathcal{K}_s , as desired. \square

Combining these three steps, it is now relatively easy to deduce the following:

Lemma 4.11 *We have that for all $\delta > 0$ small enough*

$$\lim_{\delta \rightarrow 0} \limsup_{n \rightarrow \infty} \mathbb{P}(\bar{X}_\Delta \neq \bar{Y}_\Delta | A(\delta)) = 0$$

Proof Initially there are at most 6 unmatched entries. Due to parity there can be either 6, 4 or 0 unmatched entries (note in particular that 2 is excluded, as a quick examination shows that no configuration can give rise to two unmatched entries). Furthermore, from the properties of the coupling, the number of unmatched entries either remains the same or decreases at each step. Once all the entries are matched they remain matched thereafter.

We have just shown that in any sequence of three transpositions, the probability that the number of unmatched decreases is at least δ^8 , unless T'_δ occurs during this sequence. Let Z be a binomial random variable with parameters $m = \lfloor (\Delta - 1)/3 \rfloor$ and $p = \delta^8$. Thus the event that $\{X_\Delta \neq Y_\Delta\}$ implies that there has been at most one success (i.e. $Z \leq 1$, so

$$\begin{aligned} \mathbb{P}(X_\Delta \neq Y_\Delta | A(\delta)) &\leq \mathbb{P}(Z \leq 1) + \mathbb{P}(T'_\delta \leq \Delta | A(\delta)) \\ &\leq (1 - p)^m + mp(1 - p)^{m-1} + O(\delta). \end{aligned}$$

Since $\Delta \gtrsim \delta^{-9}$ and $p = \delta^8$, the first two terms tend to 0 as $\delta \rightarrow 0$, which proves Lemma 4.11. \square

4.2.3 Coupling for $(s_2, s_3]$

The walks \tilde{X}^{id} and $\tilde{X}^{\tau_1 \circ \tau_2}$ are uniformly distributed on their conjugacy class. Thus one can couple \tilde{X}^{id} and $\tilde{X}^{\tau_1 \circ \tau_2}$ so that

- on the event $A(\delta)^c$ we have that $d(\tilde{X}_{s_2}^{\text{id}}, \tilde{X}_{s_2}^{\tau_1 \circ \tau_2}) = 2$,
- we have that using Lemma 4.11

$$\liminf_{\delta \downarrow 0} \liminf_{n \rightarrow \infty} \mathbb{P}\left(\tilde{X}_{s_2}^{\text{id}} = \tilde{X}_{s_2}^{\tau_1 \circ \tau_2} | A(\delta)\right) = 1,$$

- on the event $\{\tilde{X}_{s_2}^{\text{id}} \neq \tilde{X}_{s_2}^{\tau_1 \circ \tau_2}\}$, note that the walks \bar{X} and \bar{Y} have at most 6 unmatched entries. Hence the coupling is such that $d(\tilde{X}_{s_2}^{\text{id}}, \tilde{X}_{s_2}^{\tau_1 \circ \tau_2}) \leq 4$ no matter what.

Combining this with Lemma 4.2 we have just shown the following lemma.

Lemma 4.12 *There exists a coupling of \tilde{X}^{id} and $\tilde{X}^{\tau_1 \circ \tau_2}$ such that*

$$\limsup_{\delta \downarrow 0} \limsup_{n \rightarrow \infty} \mathbb{E} \left[d \left(\tilde{X}_{s_2}^{\text{id}}, \tilde{X}_{s_2}^{\tau_1 \circ \tau_2} \right) \right] \leq 2(1 - \theta(c)^4)$$

The theorem now follows immediately:

Proof of Theorem 1.2 It remains to see the coupling during the time interval $(s_2, s_3]$. During this time interval we apply the same transpositions to both \tilde{X}^{id} and $\tilde{X}^{\tau_1 \circ \tau_2}$ which keeps their distance constant throughout $(s_2, s_3]$. Thus we have that

$$d \left(X_t^{\text{id}}, X_t^{\tau_1 \circ \tau_2} \right) = d \left(\tilde{X}_{s_3}^{\text{id}}, \tilde{X}_{s_3}^{\tau_1 \circ \tau_2} \right) = d \left(\tilde{X}_{s_2}^{\text{id}}, \tilde{X}_{s_2}^{\tau_1 \circ \tau_2} \right).$$

Thus using Lemma 4.12 we see that (46) holds which finishes the proof. □

Acknowledgements We thank Yuval Peres and Spencer Hughes for useful discussions on discrete Ricci curvature.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

A Lower bound on mixing

In this section we give a proof of the lower bound on $t_{\text{mix}}(\delta)$ for some arbitrary $\delta \in (0, 1)$. This is for the most part a well-known argument, which shows that the number of fixed points at time $(1 - \epsilon) t_{\text{mix}}$ is large. In the case of random transpositions or more generally of a conjugacy class Γ such that $|\Gamma|$ is finite, this follows easily from the coupon collector problem. When $|\Gamma|$ is allowed to grow with n , we present here a self-contained argument for completeness.

Let $\Gamma \subset \mathcal{S}_n$ be a conjugacy class and set $k = k(n) = |\Gamma|$.

Lemma A.1 *We have that for any $\epsilon \in (0, 1)$,*

$$\lim_{n \rightarrow \infty} d_{TV}((1 - \epsilon) t_{\text{mix}}) = 1$$

Proof Let $K_m \subset \mathcal{S}_n$ be the set of permutations which have at least m fixed points. Recall that μ is the invariant measure, which is a uniform probability measure on \mathcal{S}_n or \mathcal{A}_n depending on the parity of Γ . Let U denote the uniform measure on \mathcal{S}_n . Either way,

$$\mu(K_m) \leq 2U(K_m).$$

Now, $U(K_m) \rightarrow \sum_{j=m}^{\infty} e^{-1} \frac{1}{j!}$ as $n \rightarrow \infty$, hence we deduce that

$$\limsup_{m \rightarrow \infty} \limsup_{n \rightarrow \infty} \mu(K_m) = 0. \tag{56}$$

Fix $\beta > 0$ and let

$$t_\beta = \frac{1}{k} n (\log n - \log \beta).$$

Assume that β is such that t_β is an integer. For each $i \geq 0$, γ_i write $N(\gamma_i) \subset \{1, \dots, n\}$ for the set of non-fixed points of γ_i . Then we have that for each $i \geq 0$, $|N(\gamma_i)| = k$ and further $\{N(\gamma_i)\}_{i=1}^{\infty}$ are i.i.d. subsets of $\{1, \dots, n\}$ chosen uniformly among the subsets of size $k = |\Gamma|$.

Consider for $1 \leq i \leq n$ the event A_i that the i -th card is not collected by time t_β , that is $i \notin \bigcup_{\ell=1}^{t_\beta} N(\gamma_\ell)$. Thus for $1 \leq i_1 < \dots < i_\ell \leq n$ and $\ell \leq n - k$,

$$\mathbb{P}(A_{i_1} \cap \dots \cap A_{i_\ell}) = \left(\frac{\binom{n-\ell}{k}}{\binom{n}{k}} \right)^{t_\beta}.$$

Let $N = N(n) \in \mathbb{N}$ be increasing to infinity such that $N^2 = o(n)$ and $N = o(n^2 k^{-2})$. Then by the inclusion-exclusion formula we have that

$$\mathbb{P}(A_1 \cup \dots \cup A_N) = \sum_{\ell=1}^N (-1)^{\ell+1} \binom{n}{\ell} \left(\frac{\binom{n-\ell}{k}}{\binom{n}{k}} \right)^{t_\beta}. \tag{57}$$

Writing out the fraction of binomials on the right hand side we have

$$\left(1 - \frac{k}{n-\ell} \right)^{\ell t_\beta} \leq \left(\frac{\binom{n-\ell}{k}}{\binom{n}{k}} \right)^{t_\beta} \leq \left(1 - \frac{k}{n} \right)^{\ell t_\beta}.$$

Now $-x/(1-x) \leq \log(1-x) \leq -x$ for $x \in (0, 1)$ thus we have that

$$\exp\left(-\frac{\ell k t_\beta}{n-k-\ell}\right) \leq \left(\frac{\binom{n-\ell}{k}}{\binom{n}{k}} \right)^{t_\beta} \leq \exp\left(-\frac{\ell k t_\beta}{n}\right). \tag{58}$$

On the other hand we have that

$$\frac{(n-\ell)^\ell}{\ell!} \leq \binom{n}{\ell} \leq \frac{n^\ell}{\ell!}. \tag{59}$$

Note that $ne^{-t_\beta k/n} = \beta$, then combining (58) and (59) we get

$$\left(1 - \frac{\ell}{n} \right)^\ell \exp\left(-\frac{k(k+\ell)\ell t_\beta}{n(n-k-\ell)}\right) \frac{\beta^\ell}{\ell!} \leq \binom{n}{\ell} \left(\frac{\binom{n-\ell}{k}}{\binom{n}{k}} \right)^{t_\beta} \leq \frac{\beta^\ell}{\ell!}. \tag{60}$$

Let us lower bound the error term on the left hand side of (60). First $(1 - \ell/n)^\ell \geq e^{-\ell^2/(n-\ell)}$, hence it follows that

$$\inf_{\ell \leq N} \left(1 - \frac{\ell}{n}\right)^\ell \exp\left(-\frac{k(k+\ell)\ell t_\beta}{n(n-k-\ell)}\right) \geq \inf_{\ell \leq N} \exp\left(-\frac{\ell^2}{n-\ell} - \frac{k(k+\ell)\ell t_\beta}{n(n-k-\ell)}\right).$$

It is easy to see that the right hand side above converges to 1 as $n \rightarrow \infty$. Using this and (60) it follows that

$$\lim_{n \rightarrow \infty} \sum_{\ell=1}^N (-1)^{\ell+1} \binom{n}{\ell} \left(\frac{\binom{n-\ell}{k}}{\binom{n}{k}}\right)^{t_\beta} = \lim_{n \rightarrow \infty} \sum_{\ell=1}^N (-1)^{\ell+1} \frac{\beta^\ell}{\ell!} = 1 - e^{-\beta}.$$

For integers $a < b$ let $K_{[a,b]} = A_{a+1} \cup A_{a+2} \cup \dots \cup A_b$. Then we have shown

$$\liminf_{n \rightarrow \infty} \mathbb{P}(X_{t_\beta} \in K_{[1,N]}) \geq 1 - e^{-\beta}.$$

Likewise, for any $j < \lfloor n/N \rfloor$,

$$\liminf_{n \rightarrow \infty} \mathbb{P}(X_{t_\beta} \in K_{[jN, (j+1)N]}) \geq 1 - e^{-\beta}.$$

Hence

$$\liminf_{n \rightarrow \infty} \mathbb{P}\left(X_{t_\beta} \in \bigcap_{j=1}^m K_{[jN, (j+1)N]}\right) \geq 1 - me^{-\beta}.$$

Let $\epsilon > 0$. Then for any $\beta > 0$, if $t = (1 - \epsilon)t_{\text{mix}}$ then $t < t_\beta$ for n sufficiently large, and hence

$$\liminf_{n \rightarrow \infty} \mathbb{P}\left(X_t \in \bigcap_{j=1}^m K_{[jN, (j+1)N]}\right) = 1.$$

But it is obvious that $\bigcap_{j=1}^m K_{[jN, (j+1)N]} \subset K_m$ and hence for $t = (1 - \epsilon)t_{\text{mix}}$,

$$\liminf_{n \rightarrow \infty} \mathbb{P}(X_t \in K_m) = 1. \tag{61}$$

Comparing with (56) the result follows. □

B Proof of Theorem 3.6

Let $\Gamma \subset \mathcal{S}_n$ be a conjugacy class with cycle structure (k_2, k_3, \dots) . Let $X = (X_t : t = 0, 1, \dots)$ be a random walk on \mathcal{S}_n which at each step applies an independent uniformly random element of Γ . Let $\rho = \sum_j (j - 1)k_j$ and let \tilde{X} be the transposition walk associated to the walk X using (49). In particular for $t \geq 0$, $\tilde{X}_{t\rho} = X_t$. Finally let $Z = (Z_1, Z_2, \dots)$ denote a Poisson–Dirichlet random variable.

For convenience we restate Theorem 3.6 here.

Theorem B.1 *Let $s \geq 0$ be such that $sk/(n\rho) \rightarrow c$ for some $c > c_\Gamma$. Then for each $m \in \mathbb{N}$ we have that as $n \rightarrow \infty$,*

$$\left(\frac{\tilde{\mathfrak{X}}_1(\tilde{X}_s)}{\theta(c)}, \dots, \frac{\tilde{\mathfrak{X}}_m(\tilde{X}_s)}{\theta(c)} \right) \rightarrow (Z_1, \dots, Z_m)$$

in distribution where $\theta(c)$ is given by (16).

The proof of this result is very similar to the proof of Theorem 1.1 in [24]. We give the details here.

Recall the hypergraph process $H = (H_t : t = 0, 1, \dots)$ associated with the walk X defined in Sect. 3.1. Analogously let $\tilde{G} = (\tilde{G}_t : t = 0, 1, \dots)$ be a process of graphs on $\{1, \dots, n\}$ such that the edge $\{x, y\}$ is present in \tilde{G}_t if and only if the transposition (x, y) has been applied to \tilde{X} prior to and including time t . Hence we have that for each $t = 0, 1, \dots, \tilde{G}_{t\rho} = H_t$.

Recall that \tilde{X} satisfies conditional uniformity as described in Proposition 4.1. Using the graph process \tilde{G} above and the conditional uniformity of \tilde{X} the following lemma, which is the analogue of Lemma 2.4 in [24], follows almost verbatim from Schramm’s arguments.

Lemma B.2 *Let $s \geq 0$ be such that $sk/(n\rho) \rightarrow c$ for some $c > c_\Gamma$ and let $\epsilon \in (0, 1/8)$. Let $M = M(\epsilon, n, s)$ be the minimum number of cycles of \tilde{X}_s which are needed to cover at least $(1 - \epsilon)$ proportion of the vertices in the giant component of \tilde{G}_s . Then for $\alpha \in (0, 1/8)$ we have that*

$$\limsup_{n \rightarrow \infty} \mathbb{P}(M > \alpha^{-1} |\log(\alpha\epsilon)|^2) \leq C\alpha$$

for some constant C which does not depend on α nor ϵ .

Henceforth fix some time $s \geq 0$ such that $sk/(n\rho) \rightarrow c$ for some $c > c_\Gamma$. Fix $\epsilon \in (0, 1/8)$ and define

$$\begin{aligned} \Delta &:= \lfloor \epsilon^{-1} \rfloor \\ s_0 &:= s - \Delta. \end{aligned}$$

For $t = 0, \dots, \Delta$ define $\tilde{X}_t = \mathfrak{X}(\tilde{X}_{s_0+t})$. We can assume that for $t \leq \Delta$, \tilde{X}_{s_0+t} satisfies the relaxed conditional uniformity assumption described in Definition 4.2. Indeed by making this assumption we are disregarding the constraint on the transpositions described in Proposition 4.1 applied to \tilde{X}_t for $t = s_0, \dots, s$. However the probability that we violate this constraint is at most $2\Delta k/n$.

Colour an element of $\tilde{X}_0 = \mathfrak{X}(\tilde{X}_{s_0})$ green if the cycle whose renormalised cycle length of this element lies in the giant component of \tilde{G}_{s_0} . We colour all the other elements of \tilde{X}_0 red. Thus asymptotically in n , the sum of the green elements is $\theta(c)$ and the sum of the red elements is $1 - \theta(c)$. In the evolution of $(\tilde{X}_t : t = 0, 1, \dots)$ we keep the colour scheme as follows. If an element fragments, both fragments retain the

same colour. If we coagulate two elements of the same colour then the new element retains the colour of the previous two elements. If we coagulate a green element and a red element, then the colour of the resulting element is green.

Define $\bar{X}' = (\bar{X}'_t : t = 0, \dots, \Delta)$ and $\bar{X}'' = (\bar{X}''_t : t = 0, \dots, \Delta)$ as follows. Initially $\bar{X}'_0 = \bar{X}''_0 = \bar{X}_0$. Apply the same colouring scheme to \bar{X}' and \bar{X}'' as we did to \bar{X} . Each step evolution is described as follows. Then the walks evolve as follows.

- \bar{X}'_t : Evolves the same as \bar{X} except we ignore any transition which involves a red entry.
- \bar{X}''_t : Evolves the same as \bar{X}' except that the markers u, v used in the transitions of \bar{X}'' are distributed uniformly on $[0, 1]$.

Lemma 3.1 states that the second largest component of \tilde{G}_{s_0} has size $o(n)$. Hence, initially each red element has size $o(1)$ as $n \rightarrow \infty$. Now Δ does not increase with n , hence for any $s = 0, 1, \dots, \Delta$, we are unlikely to make a coagulation (or fragmentation) in \bar{X}'_s without coagulating (or fragmenting) entries of \bar{X}'_s of similar size. Similar considerations for the processes \bar{X}' and \bar{X}'' leads to the following lemma.

Lemma B.3 *There exists a coupling between the walks \bar{X} and \bar{X}' , and between \bar{X}' and \bar{X}'' such that for each $\eta > 0$,*

$$\lim_{n \rightarrow \infty} \mathbb{P} \left(\sup_{i \in \mathbb{N}} |\bar{X}_i(\Delta) - \bar{X}'_i(\Delta)| > \eta \right) = \lim_{n \rightarrow \infty} \mathbb{P} \left(\sup_{i \in \mathbb{N}} |\bar{X}'_i(\Delta) - \bar{X}''_i(\Delta)| > \eta \right) = 0.$$

Using the preceding lemma, it suffices now to find an appropriate coupling between \bar{X}'' and Z . To do this we modify Schramm’s coupling in [24]. First we let $\{J_1, \dots, J_L\}$ be the set of times $s \in \{0, \dots, \Delta\}$ such that $\bar{X}''_{s-1} \neq \bar{X}''_s$. It is easy to see that $\lim_{n \rightarrow \infty} \mathbb{P}(L > \sqrt{\Delta}) = 1$ and henceforth we will condition on the event that $\{L > \sqrt{\Delta}\}$ and set $\Delta' = \lfloor \sqrt{\Delta} \rfloor$. Define a process $\bar{Y} = (\bar{Y}_t : t = 0, \dots, \Delta')$ as follows. Initially $\bar{Y}_0 = \bar{X}''_0$. For $t = 1, \dots, \Delta'$ we let \bar{Y}_t be X''_t renormalised so that $\sum_i \bar{Y}_i(t) = 1$ where $\bar{Y}_i(t)$ is the i -th element of \bar{Y}_t .

We define a process $\bar{Z} = (Z_t : t = 0, 1, \dots, \sqrt{\Delta})$ as follows. Initially Z_0 has the distribution of a Poisson–Dirichlet random variable, independent of \bar{Y} . Then for $t = 1, \dots, \Delta'$ define \bar{Z}_t by applying the coupling in Sect. 4.2.2 to \bar{Y} and \bar{Z} but with the following modifications:

- the markers $u, v \in [0, 1]$ are taken uniformly at random,
- we always take $v' = v$,
- we modify the definition of a refreshment time: s is a refreshment time if either $J_{s-1} + 2 \leq J_s$ or $J_s + s_0$ is a refreshment time in the sense of Definition 4.1,
- when a marked tile of size a fragments, it creates a tile of length v and a tile of length $a - v$. We mark the tile of length $a - v$.

It is not hard to check that the Poisson–Dirichlet distribution is invariant under this evolution and hence we have that for each $t = 0, 1, \dots, \Delta'$, Z_t has the law of a Poisson–Dirichlet.

Our coupling agrees with the coupling in [24, Section 3] when $\Gamma = T$ is the set of all transpositions. In this case each time s is a refreshment time and hence the

marked tile at time s is always chosen by the marker u . One can adapt the arguments in Chapter 3 of Schramm’s paper to our case by using the following idea. Note first that all the estimates of Schramm apply at s when s is a refreshment time. When s is not a refreshment time and Schramm considers the event that the marker u at time s falls inside an unmatched tile, instead we consider the event that the marker v at time $s - 1$ falls inside an unmatched tile. By the properties of the coupling, this guarantees that at time s the marked tile is unmatched.

Adapting Schramm’s arguments leads to the following lemma, which is the analogue of [24, Corollary 3.4].

Lemma B.4 *Define*

$$N^0 := \#\{i \in \mathbb{N} : \bar{Y}_i(0) > \epsilon\} + \#\{i \in \mathbb{N} : \bar{Z}_i(0) > \epsilon\}$$

and let

$$\bar{\epsilon} := \epsilon + \sum_{i=1}^{\infty} \bar{Y}_i(0) \mathbb{1}_{\{\bar{Y}_i(0) < \epsilon\}} + \sum_{i=1}^{\infty} \bar{Z}_i(0) \mathbb{1}_{\{\bar{Z}_i(0) < \epsilon\}}.$$

Define the event

$$\mathcal{B} = \left\{ \bar{\epsilon}^{4/5} \leq \frac{1}{\Delta'} \leq \frac{\bar{\epsilon}^{1/5}}{N^0 \vee 1} \right\}.$$

Let $q \in \{1, \dots, \Delta'\}$ be distributed uniformly, independent of the processes \bar{Y} and \bar{Z} . Then we have that for each $\rho > 0$,

$$\mathbb{P} \left(\sup_{i \in \mathbb{N}} |\bar{Y}_i(q) - \bar{Z}_i(q)| > \rho \right) \leq C \frac{\mathbb{P}(\mathcal{B})}{\rho \log \Delta'}$$

for some constant $C > 0$.

Using Lemma B.4 it suffices to show that $\mathbb{P}(\mathcal{B}) / \log \Delta' \rightarrow 0$ as $\epsilon \downarrow 0$. The following lemma shows a stronger result.

Lemma B.5 *Suppose that \mathcal{B} is defined as in Lemma B.4, then*

$$\lim_{\epsilon \downarrow 0} \mathbb{P}(\mathcal{B}) = 1.$$

Proof Let

$$\mathcal{B}_1 := \left\{ \bar{\epsilon}^{4/5} \leq \frac{1}{2} \epsilon^{1/2} \right\}$$

$$\mathcal{B}_2 := \left\{ 2\epsilon^{1/2} \leq \frac{\bar{\epsilon}^{1/5}}{N^0 \vee 1} \right\}$$

Now as $(1/2)\epsilon^{-1/2} \leq \Delta' \leq 2\epsilon^{-1/2}$ we have that $\mathcal{B} \supset \mathcal{B}_1 \cap \mathcal{B}_2$. First let us bound $\mathbb{P}(\mathcal{B}_1^c)$. Note that on the event \mathcal{B}_1^c we have that $\bar{\epsilon} > 2^{-5/4}\epsilon^{5/8}$. Note that a size biased sample from a Poisson–Dirichlet random variable has a uniform law on $[0, 1]$. Hence it follows that

$$\mathbb{E} \left[\sum_{i=1}^{\infty} \bar{Z}_i(0) \mathbb{1}_{\{\bar{Z}_i(0) < \epsilon\}} \right] = \epsilon$$

and thus

$$\mathbb{P} \left(\sum_{i=1}^{\infty} \bar{Z}_i(0) \mathbb{1}_{\{\bar{Z}_i(0) < \epsilon\}} > \epsilon^{5/6} \right) \leq \frac{\mathbb{E} \left[\sum_{i=1}^{\infty} \bar{Z}_i(0) \mathbb{1}_{\{\bar{Z}_i(0) < \epsilon\}} \right]}{\epsilon^{5/6}} \leq \epsilon^{1/6}.$$

Next consider the random variable M in Lemma B.2 at time $s_0 = s - \Delta$ where we recall that $\bar{Y}(0) = \bar{\mathfrak{X}}(\bar{X}_{s_0})$. We have that

$$\sum_{i=1}^{\infty} Y_i(0) \mathbb{1}_{\{\bar{Y}_i(0) : \bar{Y}_i(0) < \epsilon\}} \leq \epsilon(M + 1)$$

Then applying Lemma B.2 at time s_0 we have that

$$\mathbb{P} \left(\sum_{i=1}^{\infty} \bar{Y}_i(0) \mathbb{1}_{\{\bar{Y}_i(0) < \epsilon\}} > \epsilon^{5/6} \right) \leq \mathbb{P}(M > \epsilon^{-1/5}) \leq C\epsilon^{1/6}$$

for some constant $C > 0$. Hence it follows that for $\epsilon > 0$ small

$$\mathbb{P}(\mathcal{B}_1^c) = \mathbb{P}(\bar{\epsilon} > 2^{-5/4}\epsilon^{5/8}) \leq \mathbb{P}(\bar{\epsilon} > \epsilon^{5/6}) \leq \epsilon^{1/6} + C\epsilon^{1/6}.$$

which shows that $\mathbb{P}(\mathcal{B}_1) \rightarrow 1$ as $\epsilon \downarrow 0$.

Now we bound $\mathbb{P}(\mathcal{B}_2^c)$. Firstly we use the bound $\bar{\epsilon} \geq \epsilon$ and so we are left to bound N^0 from above. Using the stick breaking construction of Poisson–Dirichlet random variables (see for example [3, Definition 1.4]) one can show that

$$\mathbb{P} \left(\#\{i \in \mathbb{N} : Z_i(0) > \epsilon\} > \epsilon^{-1/4} \right) \leq C'\epsilon$$

for some constant $C' > 0$. On the other hand we have that

$$\#\{i \in \mathbb{N} : Y_i(0) > \epsilon\} \leq M$$

and hence using Lemma B.2 we obtain

$$\mathbb{P} \left(\#\{i \in \mathbb{N} : Y_i(0) > \epsilon\} > \epsilon^{-1/4} \right) \leq C''\epsilon^{1/5}$$

for some constant $C'' > 0$. Hence it follows that $\mathbb{P}(\mathcal{B}_2^c) \leq C'\epsilon + C''\epsilon^{1/5}$ and the result now follows. □

Theorem B.1 now follows from Lemmas B.3, B.4 and B.5.

References

1. Aldous, D.: Random walks on finite groups and rapidly mixing Markov chains. In: Seminar on Probability, XVII, Volume 986 of Lecture Notes in Mathematics, pp. 243–297. Springer, Berlin (1983)
2. Berestycki, N.: The hyperbolic geometry of random transpositions. *Ann. Probab.* **34**(2), 429–467 (2006)
3. Berestycki, N.: Recent Progress in Coalescent Theory, Volume 16 of *Ensaios Matemáticos [Mathematical Surveys]*. Sociedade Brasileira de Matemática, Rio de Janeiro (2009)
4. Berestycki, N.: Emergence of giant cycles and slowdown transition in random transpositions and k -cycles. *Electron. J. Probab.* **16**, 152–173 (2011)
5. Berestycki, N., Durrett, R.: A phase transition in the random transposition random walk. *Probab. Theor. Relat. Fields* **136**, 203–233 (2006)
6. Berestycki, N., Schramm, O., Zeitouni, O.: Mixing times for random k -cycles and coalescence-fragmentation chains. *Ann. Probab.* **39**(5), 1815–1843 (2011)
7. Bormashenko, O.: A coupling argument for the random transposition walk. arXiv preprint [arXiv:1109.3915](https://arxiv.org/abs/1109.3915)
8. Bubley, R., Dyer, M.: Path coupling: a technique for proving rapid mixing in Markov chains. In: 38th Annual Symposium on Foundations of Computer Science, 1997. Proceedings, pp. 223–231. IEEE (1997)
9. Diaconis, P.: Group Representations in Probability and Statistics. Institute of Mathematical Statistics Lecture Notes—Monograph Series, 11. Institute of Mathematical Statistics, Hayward (1988)
10. Diaconis, P., Shahshahani, M.: Generating a random permutation with random transpositions. *Z. Wahrscheinlichkeitstheorie Verwandte Geb.* **57**(2), 159–179 (1981)
11. Durrett, R.: Probability: Theory and Examples. Cambridge Series in Statistical and Probabilistic Mathematics, 4th edn. Cambridge University Press, Cambridge (2010)
12. Durrett, R.: Random Graph Dynamics. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, Cambridge (2010)
13. Flatto, L., Odlyzko, A.M., Wales, D.B.: Random shuffles and group representations. *Ann. Probab.* **13**(1), 154–178 (1985)
14. Jerrum, M.: A very simple algorithm for estimating the number of k -colorings of a low-degree graph. *Random Struct. Algorithms* **7**(2), 157–165 (1995)
15. Karoński, M., Łuczak, T.: The phase transition in a random hypergraph. *J. Comput. Appl. Math.* **142**(1), 125–135 (2002)
16. Levin, D.A., Peres, Y., Wilmer, E.L.: Markov Chains and Mixing Times. With a Chapter by James G. Propp and David B. Wilson. American Mathematical Society, Providence (2009)
17. Lulov, N., Pak, I.: Rapidly mixing random walks and bounds on characters of the symmetric group. *J. Algebr. Combin.* **16**(2), 151–163 (2002)
18. Ollivier, Y.: Ricci curvature of Markov chains on metric spaces. *J. Funct. Anal.* **256**(3), 810–864 (2009)
19. Paulin, D.: Mixing and concentration by Ricci curvature. *J. Funct. Anal.* **270**(5), 1623–1662 (2016)
20. Roichman, Y.: Upper bound on the characters of the symmetric groups. *Invent. Math.* **125**(3), 451–485 (1996)
21. Roichman, Y.: Characters of the symmetric groups: formulas, estimates and applications. In: Emerging Applications of Number Theory (Minneapolis, MN, 1996), Volume 109 of The IMA Volumes in Mathematics and its Applications, pp. 525–545. Springer, New York (1999)
22. Roussel, S.: Marches aléatoires sur le groupe symétrique. Ph.D. thesis, Toulouse (1999)
23. Roussel, S.: Phénomène de cutoff pour certaines marches aléatoires sur le groupe symétrique. *Colloq. Math.* **86**(1), 111–135 (2000)
24. Schramm, O.: Compositions of random transpositions. *Israel J. Math.* **147**, 221–243 (2005)
25. Vershik, A.M., Kerov, S.V.: Asymptotic theory of the characters of a symmetric group. *Funkt. Anal. Prilozh.* **15**(4), 15–27, 96 (1981)