Cutoff for product replacement on finite groups

Yuval Peres² · Ryokichi Tanaka¹ · Alex Zhai³

Received: 6 February 2019 / Revised: 31 January 2020 / Published online: 18 February 2020
© The Author(s) 2020

Abstract

We analyze a Markov chain, known as the *product replacement chain*, on the set of generating n -tuples of a fixed finite group G . We show that as $n \rightarrow \infty$, the total-variation mixing time of the chain has a cutoff at time $\frac{3}{2}n \log n$ with window of order n . This generalizes a result of Ben-Hamou and Peres (who established the result for $G = \mathbb{Z}/2$) and confirms a conjecture of Diaconis and Saloff-Coste that for an arbitrary but fixed finite group, the mixing time of the product replacement chain is $O(n \log n)$.

Keywords Markov chain · Mixing time · Cutoff phenomenon · Product replacement algorithm

Mathematics Subject Classification Primary 60K35 · 60K35; Secondary 60K35

1 Introduction

Let G be a finite group, and let $[n] := \{1, 2, \dots, n\}$. We consider the set G^n of all functions $\sigma : [n] \rightarrow G$ (or “configurations”). We may define a Markov chain $(\sigma_t)_{t \geq 0}$ on G^n as follows: if we have a current state σ , then uniformly at random, choose an ordered pair (i, j) of distinct integers in $[n]$, and change the value of $\sigma(i)$ to $\sigma(i)\sigma(j)^{\pm 1}$, where the signs are chosen with equal probability.

Alex Zhai was affiliated with Stanford University, Stanford, CA, USA, when work was conducted.

✉ Ryokichi Tanaka
ryokichi.tanaka@tohoku.ac.jp

Yuval Peres
yuval@yuvalperes.com

Alex Zhai
alexlinzhai@gmail.com

¹ Mathematical Institute, Tohoku University, Sendai 980-8578, Japan

² Kent State University, Kent, OH, USA

³ Stanford University, Stanford, CA, USA

We will restrict the chain $(\sigma_t)_{t \geq 0}$ to the space of *generating n -tuples*, i.e. the set of σ whose values generate G as a group:

$$\mathcal{S} := \{ \sigma \in G^n : \langle \sigma(1), \dots, \sigma(n) \rangle = G \}.$$

It is not hard to see that for fixed G and large enough n , the chain on $\mathcal{S}$ is irreducible (see [8, Lemma 3.2]). We will always assume n is large enough so that this irreducibility holds. Note that the chain is also symmetric, and it is aperiodic because it has holding on some states. Thus, the chain has a uniform stationary distribution π with $\pi(\sigma) = 1/|\mathcal{S}|$.

This Markov chain was first considered in the context of computational group theory—it models the *product replacement algorithm* for generating random elements of a finite group introduced in [6]. By running the chain for a long enough time t and choosing a uniformly random index $k \in [n]$, the element $\sigma_t(k)$ is a (nearly) uniformly random element of G . The product replacement algorithm has been found to perform well in practice [6, 10], but the question arises: how large does t need to be in order to ensure near uniformity?

One way of answering the question is to estimate the mixing time of the Markov chain. It was shown by Diaconis and Saloff-Coste that for any fixed finite group G , there exists a constant C_G such that the ℓ^2 -mixing time is at most $C_G n^2 \log n$ [8, 9] (see also Chung and Graham [3] for a simpler proof of this fact with a different value for C_G).

In another line of work, Lubotzky and Pak [12] analyzed the mixing of the product replacement chain in terms of Kazhdan constants (see also subsequent quantitative estimates for Kazhdan constants by Kassabov [11]). We also mention a result of Pak [14] which shows mixing in $\text{polylog}(|G|)$ steps when $n = \Theta(\log |G| \log \log |G|)$. The reader may consult the survey [15] for further background on the product replacement algorithm.

Diaconis and Saloff-Coste conjectured that the mixing time bound can be improved to $C_G n \log n$ [9, Remark 2, Section 7, p. 290], based on the observation that at least $n \log n$ steps are needed by the classical coupon-collector's problem. This was confirmed in the case $G = \mathbb{Z}/2$ by Chung and Graham [4] and recently refined by Ben-Hamou and Peres, who show that when $G = \mathbb{Z}/2$, the chain in fact exhibits a cutoff at time $\frac{3}{2}n \log n$ in total-variation with window of order n [2].

In this paper, we extend the result of Ben-Hamou and Peres to all finite groups. Note that this also verifies the conjecture of Diaconis and Saloff-Coste for a fixed finite group. To state the result, let us denote the total variation distance between $\mathbb{P}_\sigma(\sigma_t \in \cdot)$ and π by

$$d_\sigma(t) := \max_{A \subseteq \mathcal{S}} |\mathbb{P}_\sigma(\sigma_t \in A) - \pi(A)|.$$

Theorem 1.1 *Let G be a finite group. Then, the Markov chain $(\sigma_t)_{t \geq 0}$ on the set of generating n -tuples of G has a total-variation cutoff at time $\frac{3}{2}n \log n$ with window of order n . More precisely, we have*

$$\lim_{\beta \rightarrow \infty} \limsup_{n \rightarrow \infty} \max_{\sigma \in \mathcal{S}} d_\sigma \left(\frac{3}{2}n \log n + \beta n \right) = 0 \quad (1)$$

and

$$\lim_{\beta \rightarrow \infty} \liminf_{n \rightarrow \infty} \max_{\sigma \in \mathcal{S}} d_{\sigma} \left(\frac{3}{2} n \log n - \beta n \right) = 1. \quad (2)$$

1.1 A connection to cryptography

We mention another motivation for studying the product replacement chain in the case $G = (\mathbb{Z}/q)^m$ for a prime $q \geq 2$ and integers $m \geq 1$. It comes from a public-key authentication protocol proposed by Sotiraki [16], which we now briefly describe. In the protocol, a verifier wants to check the identity of a prover based on the time needed to answer a challenge.

First, the prover runs the Markov chain with $G = (\mathbb{Z}/q)^m$ and $n = m$, which can be interpreted as performing a random walk on $SL_n(\mathbb{Z}/q)$, where $\sigma(k)$ is viewed as the k -th row of a $n \times n$ matrix. (In each step, a random row is either added to or subtracted from another random row.)

After t steps, the prover records the resulting matrix $A \in SL_n(\mathbb{Z}/q)$ and makes it public. To authenticate, the verifier gives the prover a vector $x \in (\mathbb{Z}/q)^n$ and challenges her to compute $y := Ax$. The prover can perform this calculation in $O(t)$ operations by retracing the trajectory of the random walk.

Without knowing the trajectory, if t is large enough, an adversary will not be able to distinguish A from a random matrix and will be forced to perform the usual matrix-vector multiplication (using n^2 operations) to complete the challenge. Thus, the question is whether $t \ll n^2$ is large enough for the matrix A to become sufficiently random, so that the prover can answer the challenge much faster than an adversary.

Note that when $n > m$, the product replacement chain on $G = (\mathbb{Z}/q)^m$ amounts to the projection of the random walk on $SL_n(\mathbb{Z}/q)$ onto the first m columns. Thus, Theorem 1.1 shows that when m is fixed and $n \rightarrow \infty$, the mixing time for the first m columns is around $\frac{3}{2}n \log n$. One then hopes that the mixing of several columns is enough to make it computationally intractable to distinguish A from a random matrix; this would justify the authentication protocol, as $n \log n \ll n^2$.

We remark that when t is much larger than the mixing time of the random walk on $SL_n(\mathbb{Z}/q)$ generated by row and additions and subtractions, it is information theoretically impossible for an adversary to distinguish A from a random matrix. However, the diameter of the corresponding Cayley graph on $SL_n(\mathbb{Z}/q)$ is known to be of order $\Theta\left(\frac{n^2}{\log_q n}\right)$ [1, 5], so a lower bound of the same order necessarily holds for the mixing time. Diaconis and Saloff-Coste [8, Section 4, p. 420] give an upper bound of $O(n^4)$, which was subsequently improved to $O(n^3)$ by Kassabov [11]. Closing the gap between n^3 and $\frac{n^2}{\log n}$ remains an open problem.

1.2 Outline of proof

The proof of Theorem 1.1 analyzes the mixing behavior in several stages:

- an initial “burn-in” period lasting around $n \log n$ steps, after which the group elements appearing in the configuration are not mostly confined to any proper subgroup of G ;

- an averaging period lasting around $\frac{1}{2}n \log n$ steps, after which the counts of group elements become close to their average value under the stationary distribution; and
- a coupling period lasting $O(n)$ steps, after which our chain becomes exactly coupled to the stationary distribution with high probability.

The argument is in the spirit of [2], but a more elaborate analysis is required in the second and third stages. To analyze the first stage, for a fixed proper subgroup H , the number of group elements in H appearing in the configuration is a birth-and-death process whose transition probabilities are easy to estimate. The analysis of the resulting chain is the same as in [2], and we can then union bound over all proper subgroups H .

In the second stage, for a given starting configuration $\sigma_0 \in \mathcal{S}$, we consider quantities $n_{a,b}(\sigma)$ counting the number of sites k where $\sigma_0(k) = a$ and $\sigma(k) = b$. A key observation (which also appears in [2]) is that by symmetry, projecting the Markov chain onto the values $(n_{a,b}(\sigma_t))_{a,b \in G}$ does not affect the mixing behavior. Thus, it is enough to understand the mixing behavior of the counts $n_{a,b}$.

One expects these counts to evolve towards their expected value $\mathbb{E}_{\sigma \sim \pi} n_{a,b}(\sigma)$ as the chain mixes. To carry out the analysis rigorously, we write down a stochastic difference equation for the $n_{a,b}$ and analyze it via the Fourier transform. Intuitively, as $n \rightarrow \infty$, the process approaches a “hydrodynamic limit” so that it becomes approximately deterministic. It turns out that after about $\frac{1}{2}n \log n$ steps, the $n_{a,b}$ are likely to be within $O(\sqrt{n})$ of their expected value. Our analysis requires a sufficiently “generic” initial configuration, which is why the first stage is necessary.

Finally, in the last stage, we show that if the $(n_{a,b}(\sigma))_{a,b \in G}$ and $(n_{a,b}(\sigma'))_{a,b \in G}$ for two configurations are within $O(\sqrt{n})$ in ℓ^1 distance, they can be coupled to be exactly the same with high probability after $O(n)$ steps of the Markov chain. A standard argument involving coupling to the stationary distribution then implies a bound on the mixing time.

The main idea to prove the coupling bound is that even if the ℓ^1 distance evolves like an unbiased random walk, there is a good chance that it will hit 0 due to random fluctuations. A similar argument is used to prove cutoff for lazy random walk on the hypercube [13, Chapter 18]. However, some careful accounting is necessary in our setting to ensure that in fact the ℓ^1 distance does not increase in expectation and to ensure sufficient fluctuations.

1.3 Organization of the paper

The rest of the paper is organized as follows. In Sect. 2, we state (without proof) the key lemmas describing the behavior in each of the three stages and use these to prove the upper bound (1) in Theorem 1.1. Sections 3 and 4 contain the proofs of these lemmas. Finally, in Sect. 5, we prove the lower bound (2) in Theorem 1.1; this is mostly a matter of verifying that the estimates used in the upper bound were tight.

1.4 Notation

Throughout this paper, we use $c, C, C', \dots$ to denote absolute constants whose exact values may change from line to line, and also use them with subscripts, for instance,

C_G to specify its dependency only on G . We also use subscripts with big- O notation, e.g. we write $O_G(\cdot)$ when the implied constant depends only on G .

2 Proof of Theorem 1.1 (1)

Let us fix a finite group G and denote its cardinality by $\mathcal{Q} := |G|$. For a configuration $\sigma \in \mathcal{S}$, let $n_a(\sigma)$ denote the number of sites having group element a , i.e.,

$$n_a(\sigma) := |\{i \in [n] : \sigma(i) = a\}|.$$

2.1 The burn-in period

For a proper subgroup $H \subseteq G$, let

$$n_{non}^H(\sigma) := \sum_{a \in G \setminus H} n_a(\sigma)$$

denote the number of sites not in H , and define for $c \in (0, 1)$ the set

$$\mathcal{S}_{non}(c) := \{\sigma \in \mathcal{S} : n_{non}^H(\sigma) \geq cn \text{ for all proper subgroups } H \subseteq G\}.$$

Thus, $\mathcal{S}_{non}(c)$ is the set of states σ where the group elements appearing in σ are not mostly confined to any particular proper subgroup of G . The next lemma shows that we reach $\mathcal{S}_{non}(1/3)$ in about $n \log n$ steps, and once we reach $\mathcal{S}_{non}(1/3)$, we remain in $\mathcal{S}_{non}(1/6)$ for n^2 steps with high probability. Note that n^2 is much larger than the overall mixing time, so we may essentially assume that we are in $\mathcal{S}_{non}(1/6)$ for all of the later stages.

Lemma 2.1 *Let $\tau_{1/3} := \min\{t \geq 0 : \sigma_t \in \mathcal{S}_{non}(1/3)\}$ be the first time to hit $\mathcal{S}_{non}(1/3)$. Then for all large enough n and for all large enough $\beta > 0$,*

$$\max_{\sigma \in \mathcal{S}} \mathbb{P}_{\sigma}(\tau_{1/3} > n \log n + \beta n) \leq \frac{120\mathcal{Q}}{\beta^2}.$$

Moreover, there exists a constant C_G depending only on G such that

$$\max_{\sigma \in \mathcal{S}_{non}(1/3)} \mathbb{P}_{\sigma}(\sigma_t \notin \mathcal{S}_{non}(1/6) \text{ for some } t \leq n^2) \leq C_G n^2 e^{-n/10}.$$

Proof Fix a proper subgroup $H \subset G$, and consider what happens to $n_{non}^H(\sigma_t)$ at time t . Suppose our next step is to replace $\sigma(i)$ with $\sigma(i)\sigma(j)$.

If $\sigma(j) \in H$, then $n_{non}^H(\sigma_{t+1}) = n_{non}^H(\sigma_t)$. If $\sigma(j) \notin H$ and $\sigma(i) \in H$, then $n_{non}^H(\sigma_{t+1}) = n_{non}^H(\sigma_t) + 1$. Finally, if $\sigma(j), \sigma(i) \notin H$, then $\sigma(i)\sigma(j)$ may or may not be in H , so $n_{non}^H(\sigma_{t+1}) \geq n_{non}^H(\sigma_t) - 1$.

Let $(N_t)_{t \geq 0}$ be the birth-and-death chain with the following transition probabilities for $1 \leq k \leq n$:

$$\begin{aligned}\mathbb{P}(N_{t+1} = k + 1 \mid N_t = k) &= \frac{k(n-k)}{n(n-1)} \\ \mathbb{P}(N_{t+1} = k - 1 \mid N_t = k) &= \frac{k(k-1)}{n(n-1)} \\ \mathbb{P}(N_{t+1} = k \mid N_t = k) &= \frac{n-k}{n}.\end{aligned}$$

We start this chain at $N_0 = n_{non}^H(\sigma_0)$; note that because the elements appearing in σ_0 generate G , we are guaranteed to have $n_{non}^H(\sigma_0) > 0$.

The above birth-and-death chain corresponds to the behavior of $(n_{non}^H(\sigma_t))$ if whenever $\sigma(j), \sigma(i) \notin H$, it always happened that $\sigma(i)\sigma(j) \in H$. Thus, $(n_{non}^H(\sigma_t))$ stochastically dominates (N_t) .

The chain (N_t) is precisely what is analyzed in [2] for the case $G = \mathbb{Z}/2$. Let

$$T_k := \min\{t \geq 0 : N_t = k\}.$$

Then, we have $\mathbb{E}_{k-1} T_k \leq \frac{n^2}{k(n-2k)}$ [2, (2) in the proof of Lemma 1] and thus $\mathbb{E}_1(T_{n/3}) = \sum_{k=2}^{n/3} \mathbb{E}_{k-1} T_k \leq n \log n + n$. On the other hand, setting $v_k = \text{Var}_{k-1}(T_k)$, we have $v_2 \leq n^2$,

$$v_{k+1} \leq \frac{k}{n-k} v_k + \frac{54n^2}{k^2},$$

and $\text{Var}_1(T_{n/3}) = \sum_{k=2}^{n/3} v_k \leq 110n^2$ [2, The proof of Lemma 1]. Hence by Chebyshev's inequality for all large enough $\beta > 0$,

$$\mathbb{P}_1(T_{n/3} > n \log n + \beta n) \leq \frac{120}{\beta^2}.$$

Moreover, we have $\mathbb{P}_{n/3}(T_{n/6} \leq n^2) \leq n^2 e^{-n/10}$. Indeed, this follows from the fact that for $m < k$, we have

$$\mathbb{P}_k(T_m \leq n^2) \leq n^2 \frac{\pi_{\text{BD}}(m)}{\pi_{\text{BD}}(k)},$$

where $\pi_{\text{BD}}(k) = \binom{n}{k} / (2^n - 1)$ [2, (5) and the following in the proof of Proposition 2].

We now take a union bound over all the proper subgroups H . $\square$

2.2 The averaging period

In the next stage, the counts $n_a(\sigma_t)$ go toward their average value. We actually analyze this stage in two substages, looking at a “proportion vector” and “proportion matrix”, as described below.

2.2.1 Proportion vector chain

For a configuration $\sigma \in \mathcal{S}$, we consider the Q -dimensional vector $(n_a(\sigma)/n)_{a \in G}$, which we call the *proportion vector* of σ . One may check that for a typical $\sigma \in \mathcal{S}$, each $n_a(\sigma)/n$ is about $1/Q$. For each $\delta > 0$, we define the δ -typical set

$$\mathcal{S}_*(\delta) := \left\{ \sigma \in \mathcal{S} : \left\| \left(\frac{n_a(\sigma)}{n} \right)_{a \in G} - \left(\frac{1}{Q} \right)_{a \in G} \right\| \leq \delta \right\},$$

where $\|\cdot\|$ denotes the ℓ^2 -norm in $\mathbb{R}^G$.

The following lemma implies that starting from $\sigma \in \mathcal{S}_{non}(1/3)$, we reach $\mathcal{S}_*(\delta)$ in $O_\delta(n)$ steps with high probability. The proof is given in Sect. 3.4.

Lemma 2.2 *Consider any $\sigma \in \mathcal{S}_{non}(1/3)$ and any constant $\delta > 0$. There exists a constant $C_{G,\delta}$ depending only on G and δ such that for any $T \geq C_{G,\delta}n$, we have*

$$\mathbb{P}_\sigma(\sigma_T \notin \mathcal{S}_*(\delta)) \leq \frac{1}{n}$$

for all large enough n .

2.2.2 Proportion matrix chain

We actually need a more precise averaging than what is provided by Lemma 2.2. Fix a configuration $\sigma_0 \in \mathcal{S}$. For any $\sigma \in \mathcal{S}$ and for any $a, b \in G$, define

$$n_{a,b}^{\sigma_0}(\sigma) := |\{i \in [n] : \sigma_0(i) = a, \sigma(i) = b\}|.$$

If we run the Markov chain $(\sigma_t)_{t \geq 0}$ with initial state σ_0 , then $n_{a,b}^{\sigma_0}(\sigma_t)$ is the number of sites that originally contained the element a (at time 0) but now contain b (at time t). Note that

$$\sum_{b \in G} n_{a,b}^{\sigma_0}(\sigma) = n_a(\sigma_0) \quad \text{and} \quad \sum_{a \in G} n_{a,b}^{\sigma_0}(\sigma) = n_b(\sigma).$$

We can then associate with $(\sigma_t)_{t \geq 0}$ another Markov chain $\left(n_{a,b}^{\sigma_0}(\sigma_t) \right)_{a,b \in G}$ for $t \geq 0$, which we call the *proportion matrix chain (with respect to σ_0)*. The state space for the proportion matrix chain is $\{0, 1, \dots, n\}^{G \times G}$, and the transition probabilities depend on σ_0 .

The proportion matrix acts like a “sufficient statistic” for analyzing our Markov chain started at σ_* , because of the permutation invariance of our dynamics. In fact, as the following lemma shows, the distance to stationarity of the proportion matrix chain is equal to the distance to stationarity of the original chain.

Lemma 2.3 *Let $\sigma_* \in \mathcal{S}$ be a configuration. For the Markov chain $(\sigma_t)_{t \geq 0}$ with initial state σ_* , we consider $\left(n_{a,b}^{\sigma_*}(\sigma_t)\right)_{a,b \in G}$. Let $\bar{\pi}^{\sigma_*}$ be the stationary measure for the Markov chain $\{(n_{a,b}^{\sigma_*}(\sigma_t))_{a,b \in G}\}_{t \geq 0}$ on $\{0, 1, \dots, n\}^{G \times G}$. Then, for every $t \geq 0$, we have*

$$\|\mathbb{P}_{\sigma_*}(\sigma_t \in \cdot) - \pi\|_{\text{TV}} = \left\| \mathbb{P}_{\sigma_*} \left((n_{a,b}^{\sigma_*}(\sigma_t))_{a,b \in G} \in \cdot \right) - \bar{\pi}^{\sigma_*} \right\|_{\text{TV}}.$$

Proof For any matrix $N = (N_{a,b})_{a,b \in G} \in \{0, 1, \dots, n\}^{G \times G}$, write

$$\mathcal{X}_{(N)} := \left\{ \sigma \in \mathcal{S} : n_{a,b}^{\sigma_*}(\sigma) = N_{a,b} \text{ for all } a, b \in G \right\}$$

for the set of configurations with N as their proportion matrix.

Since the distribution of σ_t is invariant under permutations on sites $i \in [n]$ preserving the set $\{i : \sigma_*(i) = a\}$ for every $a \in G$, the conditional probability measures $\mathbb{P}_{\sigma_*}(\sigma_t \in \cdot \mid \sigma_t \in \mathcal{X}_{(N)})$ and $\pi(\cdot \mid \mathcal{X}_{(N)})$ are both uniform on $\mathcal{X}_{(N)}$. This implies that for each $\sigma \in \mathcal{X}_{(N)}$,

$$|\mathbb{P}_{\sigma_*}(\sigma_t = \sigma) - \pi(\sigma)| = \frac{1}{|\mathcal{X}_{(N)}|} \left| \mathbb{P}_{\sigma_*} \left((n_{a,b}^{\sigma_*}(\sigma_t))_{a,b \in G} = N \right) - \bar{\pi}^{\sigma_*}(N) \right|,$$

and summing over all $\sigma \in \mathcal{X}_{(N)}$ and all N , we obtain the claim. $\square$

For $\sigma_0 \in \mathcal{S}$ and $r > 0$, define the set of configurations

$$\mathcal{S}_*(\sigma_0, r) := \left\{ \sigma \in \mathcal{S} : \left\| \left(\frac{n_{a,b}^{\sigma_0}(\sigma)}{n_a(\sigma_0)} \right)_{b \in G} - \left(\frac{1}{Q} \right)_{b \in G} \right\| \leq r \text{ for all } a \in G \right\}.$$

Roughly speaking, the following lemma shows that starting from a typical configuration $\sigma_* \in \mathcal{S}_*\left(\frac{1}{4Q}\right)$, we need about $\frac{1}{2}n \log n$ steps to reach $\mathcal{S}_*\left(\sigma_*, \frac{R}{\sqrt{n}}\right)$, where R is a constant. We show this fact in a slightly more general form where the initial state need not be σ_* ; the proof is given in Sect. 3.5.

Lemma 2.4 *Consider any $\sigma_*, \sigma'_* \in \mathcal{S}_*\left(\frac{1}{4Q}\right)$, and let $T := \lceil \frac{1}{2}n \log n \rceil$. There exists a constant $C_G > 0$ depending only on G such that for any given $R > 0$, we have*

$$\mathbb{P}_{\sigma'_*} \left(\sigma_T \notin \mathcal{S}_* \left(\sigma_*, \frac{R}{\sqrt{n}} \right) \right) \leq C_G e^{-R} + \frac{1}{n}$$

for all large enough n .

2.3 The coupling period

After reaching $\mathcal{S}_*\left(\sigma_*, \frac{R}{\sqrt{n}}\right)$, we show that only $O(n)$ additional steps are needed to mix in total variation distance. The main ingredient in the proof is a coupling of proportion matrix chains so that they coalesce in $O(n)$ steps when they both start from configurations $\sigma, \tilde{\sigma} \in \mathcal{S}_*\left(\sigma_*, \frac{R}{\sqrt{n}}\right)$. We construct such a coupling and prove the following lemma in Sect. 4.

Lemma 2.5 *Consider any $\sigma_* \in \mathcal{S}_*\left(\frac{1}{5Q^3}\right)$, and let $R > 0$. Suppose $\sigma, \tilde{\sigma} \in \mathcal{S}_*\left(\sigma_*, \frac{R}{\sqrt{n}}\right)$. Then, there exists a coupling $(\sigma_t, \tilde{\sigma}_t)$ of the Markov chains with initial states $(\sigma, \tilde{\sigma})$ such that for a given $\beta > 0$ and all large enough n ,*

$$\mathbb{P}_{\sigma, \tilde{\sigma}}(\tau > \beta n) \leq \frac{32Q^2R}{\sqrt{\beta}},$$

where $\tau := \min\{t \geq 0 : n_{a,b}^{\sigma_t}(\sigma_t) = n_{a,b}^{\tilde{\sigma}_t}(\tilde{\sigma}_t) \text{ for all } a, b \in G\}$.

To translate this coupling time into a bound on total variation distance, we need also the simple observation that the stationary measure π concentrates on $\mathcal{S}_*\left(\sigma_*, \frac{R}{\sqrt{n}}\right)$ except for probability $O(1/R^2)$, as given in the next lemma.

Lemma 2.6 *For the stationary distribution π of the chain $(\sigma_t)_{t \geq 0}$, for every $R > 0$ and for all $n > m$,*

$$\pi\left(\sigma \notin \mathcal{S}_*\left(\frac{R}{\sqrt{n}}\right)\right) \leq \frac{C_G}{R^2}.$$

Moreover for every $\delta < 1/(2Q)$, for every $R > 0$ and for all $n > m$,

$$\max_{\sigma_* \in \mathcal{S}_*(\delta)} \pi\left(\sigma \notin \mathcal{S}_*\left(\sigma_*, \frac{R}{\sqrt{n}}\right)\right) \leq \frac{2C_G Q}{R^2},$$

where C_G and m are constants depending only on G .

Proof Observe that since the stationary distribution π is uniform on $\mathcal{S}$, it is given by the uniform distribution Unif on G^n conditioned on $\mathcal{S}$. Note that we can always generate G using each of its $|G|$ elements, so we have an easy lower bound of $|\mathcal{S}| \geq |G|^{n-|G|}$. Consequently, we have

$$\begin{aligned} \pi\left(\sigma \notin \mathcal{S}_*\left(\frac{R}{\sqrt{n}}\right)\right) &\leq |G|^{|G|} \text{Unif}\left(\sigma \notin \mathcal{S}_*\left(\frac{R}{\sqrt{n}}\right)\right) \\ &\leq |G|^{|G|} \sum_{a \in G} \text{Unif}\left(\left|\frac{n_a(\sigma)}{n} - \frac{1}{Q}\right| \geq \frac{R}{\sqrt{n}}\right) \\ &\leq \frac{|G|^{|G|}}{R^2} \left(1 - \frac{1}{Q}\right). \end{aligned}$$

Concerning the second assertion, we note that $n_a(\sigma_*) \geq (1/\mathcal{Q} - \delta)n$ for each $a \in G$; the rest follows similarly, so we omit the details. $\square$

Remark 2.7 In Lemma 2.6 above, we have given a very loose bound on C_G for sake of simplicity. Actually, it is not hard to see that holding G fixed, we have $\lim_{n \rightarrow \infty} |S|/|G|^n = 1$. See also [9, Section 6.B.] for more explicit bounds for various families of groups.

Together, Lemmas 2.4, 2.5, and 2.6 imply the following bound for total variation distance.

Lemma 2.8 *Let $\beta > 0$ be given, and let $T := \lceil \frac{1}{2}n \log n \rceil + \lceil \beta n \rceil$. Then, for any $\sigma_* \in \mathcal{S}_*\left(\frac{1}{5\mathcal{Q}^3}\right)$, we have*

$$\|\mathbb{P}_{\sigma_*}(\sigma_T \in \cdot) - \pi\|_{\text{TV}} \leq \frac{C_G}{\beta^{1/4}},$$

where C_G is a constant depending only on G .

Proof Let $\tilde{\sigma}$ be drawn from the stationary distribution π . Define

$$\tau = \min \left\{ t \geq 0 : n_{a,b}^{\sigma_*}(\sigma_t) = n_{a,b}^{\sigma_*}(\tilde{\sigma}_t) \text{ for all } a, b \in G \right\},$$

where $(\tilde{\sigma}_t)$ is a Markov chain started at $\tilde{\sigma}$. Let $\bar{\pi}^{\sigma_*}$ denote the stationary distribution for the proportion matrix with respect to σ_* . Since $\tilde{\sigma}$ was drawn from π , the proportion matrix of $\tilde{\sigma}_t$ remains distributed as $\bar{\pi}^{\sigma_*}$ for all t .

We first run σ and $\tilde{\sigma}$ independently up until time $T_1 := \lceil \frac{1}{2}n \log n \rceil$. For a parameter R to be specified later, consider the events

$$\mathcal{G} := \left\{ \sigma_{T_1} \in \mathcal{S}_*\left(\sigma_*, \frac{R}{\sqrt{n}}\right) \right\}, \quad \tilde{\mathcal{G}} := \left\{ \tilde{\sigma}_{T_1} \in \mathcal{S}_*\left(\sigma_*, \frac{R}{\sqrt{n}}\right) \right\}.$$

Lemma 2.4 implies that $\mathbb{P}(\mathcal{G}^c) \leq C_G e^{-R} + \frac{1}{n}$, and Lemma 2.6 implies that $\mathbb{P}(\tilde{\mathcal{G}}^c) \leq \frac{2C_G \mathcal{Q}}{R^2}$.

Let $T_2 := \lceil \beta n \rceil$. Starting from time T_1 , as long as both $\mathcal{G}$ and $\tilde{\mathcal{G}}$ hold, we may use Lemma 2.5 to form a coupling $(\sigma_t, \tilde{\sigma}_t)$ so that

$$\mathbb{P}_{\sigma_*, \sigma_*} \left(n_{a,b}^{\sigma_*}(\sigma_{T_1+T_2}) \neq n_{a,b}^{\sigma_*}(\tilde{\sigma}_{T_1+T_2}) \text{ for some } a, b \in G \mid \mathcal{G} \cap \tilde{\mathcal{G}} \right) \leq \frac{C \mathcal{Q}^2 R}{\sqrt{\beta}}.$$

Setting $R = \beta^{1/4}$, we conclude that

$$\begin{aligned} & \mathbb{P}_{\sigma_*, \sigma_*} \left(n_{a,b}^{\sigma_*}(\sigma_{T_1+T_2}) \neq n_{a,b}^{\sigma_*}(\tilde{\sigma}_{T_1+T_2}) \text{ for some } a, b \in G \right) \\ & \leq \frac{CQ^2R}{\sqrt{\beta}} + \mathbb{P}(\mathcal{G}^c) + \mathbb{P}(\tilde{\mathcal{G}}^c) \\ & \leq \frac{CQ^2R}{\sqrt{\beta}} + \left(C_G e^{-R} + \frac{1}{n} \right) + \frac{2C_G Q}{R^2} = O_G \left(\frac{1}{\beta^{1/4}} \right). \end{aligned}$$

We have $T = T_1 + T_2$, and recall that the proportion matrix for $\tilde{\sigma}$ is stationary for all time. This yields

$$\left\| \mathbb{P}_{\sigma_*} \left((n_{a,b}^{\sigma_*}(\sigma_T))_{a,b \in G} \in \cdot \right) - \bar{\pi}^{\sigma_*} \right\|_{\text{TV}} = O_G \left(\frac{1}{\beta^{1/4}} \right).$$

The result then follows by Lemma 2.3. $\square$

2.4 Proof of the main theorem

We now combine the lemmas from the burn-in, averaging, and coupling periods to complete the proof of the upper bound in Theorem 1.1.

Proof of Theorem 1.1 (1) Define $T_1 := \lceil n \log n + \beta n \rceil$, $T_2 := \lceil \beta n \rceil$, and $T_3 := \lceil \frac{1}{2} n \log n \rceil + \lceil \beta n \rceil$.

Let $\tau_{1/3}$ be the first time to hit $\mathcal{S}_{\text{non}}(1/3)$ as in Lemma 2.1. Then, Lemma 2.1 implies that for any $\sigma_1 \in \mathcal{S}$ and any $t \geq 0$, we have

$$\begin{aligned} d_{\sigma_1}(T_1 + t) & \leq \mathbb{P}_{\sigma_1}(\tau_{1/3} > T_1) + \max_{\sigma \in \mathcal{S}_{\text{non}}(1/3)} d_{\sigma}(t) \\ & \leq \frac{120Q}{\beta^2} + \max_{\sigma \in \mathcal{S}_{\text{non}}(1/3)} d_{\sigma}(t). \end{aligned} \quad (3)$$

Next, by Lemma 2.2, for any $\sigma_2 \in \mathcal{S}_{\text{non}}(1/3)$ and when β and n are sufficiently large, we have that $\mathbb{P}_{\sigma_2} \left(\sigma_{T_2} \notin \mathcal{S}_* \left(\frac{1}{5Q^3} \right) \right) \leq \frac{1}{n}$. Consequently, for $\sigma_2 \in \mathcal{S}_{\text{non}}(1/3)$, we have

$$d_{\sigma_2}(T_2 + t) \leq \frac{1}{n} + \max_{\sigma_* \in \mathcal{S}_* \left(\frac{1}{5Q^3} \right)} d_{\sigma_*}(t). \quad (4)$$

Finally, Lemma 2.8 states that

$$\max_{\sigma_* \in \mathcal{S}_* \left(\frac{1}{5Q^3} \right)} d_{\sigma_*}(T_3) \leq \frac{C_G}{\beta^{1/4}}. \quad (5)$$

Thus, combining (3), (4), and (5), we obtain for any $\sigma \in S$ that

$$\begin{aligned} d_\sigma \left(\frac{3}{2} n \log n + 4\beta n \right) &\leq d_\sigma (T_1 + T_2 + T_3) \\ &\leq \frac{120Q}{\beta^2} + \frac{1}{n} + \frac{C_G}{\beta^{1/4}} \end{aligned}$$

sending $n \rightarrow \infty$ and then $\beta \rightarrow \infty$ yields (1). $\square$

3 Proofs for the averaging period

In this section, we prove Lemmas 2.2 and 2.4. The proofs are based on analyzing stochastic difference equations satisfied by the Fourier transform of the proportion vector or matrix.

3.1 The Fourier transform for G

We first establish some notation and preliminaries for the Fourier transform. Let G^* be a complete set of non-trivial irreducible representations of G . In other words, for each $\rho \in G^*$, we have a finite dimensional complex vector space V_ρ such that $\rho : G \rightarrow GL(V_\rho)$ is a non-trivial irreducible representation, and any non-trivial irreducible representation of G is isomorphic to some unique $\rho \in G^*$. Moreover, we may equip each V_ρ with an inner product for which $\rho \in G^*$ is unitary.

For a configuration $\sigma \in S$ and for each $\rho \in G^*$, we consider the matrix acting on V_ρ given by

$$x_\rho(\sigma) := \sum_{a \in G} \frac{n_a(\sigma)}{n} \rho(a),$$

so that $x_\rho(\sigma)$ is the Fourier transform of the proportion vector at the representation ρ . We write $x(\sigma) := (x_\rho(\sigma))_{\rho \in G^*}$.

Let $\tilde{V} := \bigoplus_{\rho \in G^*} \text{End}_{\mathbb{C}}(V_\rho)$, and write $d_\rho := \dim_{\mathbb{C}} V_\rho$. For an element $x = (x_\rho)_{\rho \in G^*} \in \tilde{V}$, we define a norm $\|\cdot\|_{\tilde{V}}$ given by

$$\|x\|_{\tilde{V}}^2 := \frac{1}{Q} \sum_{\rho \in G^*} d_\rho \|x_\rho\|_{\text{HS}}^2,$$

where $\langle A, B \rangle_{\text{HS}} = \text{Tr}(A^* B)$ denotes the Hilbert–Schmidt inner product in $\text{End}_{\mathbb{C}}(V_\rho)$ and $\|\cdot\|_{\text{HS}}$ denotes the corresponding norm. (Note that $\langle \cdot, \cdot \rangle_{\text{HS}}$ and $\|\cdot\|_{\text{HS}}$ depend on ρ , but for sake of brevity, we omit the ρ when there is no danger of confusion.)

The Peter–Weyl theorem [7, Chapter 2] says that

$$L^2(G) \cong \mathbb{C} \oplus \tilde{V},$$

where the isomorphism is given by the Fourier transform. The Plancherel formula then reads

$$\|x(\sigma)\|_{\tilde{V}}^2 = \left\| \left(\frac{n_a(\sigma)}{n} \right)_{a \in G} - \left(\frac{1}{Q} \right)_{a \in G} \right\|^2. \quad (6)$$

Thus, in order to show that $\sigma \in \mathcal{S}_*(\delta)$, it suffices to show that $\|x(\sigma)\|_{\tilde{V}}$ is small. A similar argument may be applied to the proportion matrix instead of the proportion vector.

Finally, for an element $A \in \text{End}_{\mathbb{C}}(V_{\rho})$, we will at times also consider the *operator norm* $\|A\|_{op} := \sup_{v \in V_{\rho}, v \neq 0} \|Av\|/\|v\|$. We will also sometimes use the following (equivalent) variational characterization of the operator norm:

$$\begin{aligned} \sup_{\substack{X \in \text{End}_{\mathbb{C}}(V_{\rho}) \\ \|X\|_{HS}=1}} \|XA\|_{HS}^2 &= \sup_{\substack{X \in \text{End}_{\mathbb{C}}(V_{\rho}) \\ \|X\|_{HS}=1}} \text{Tr}(XAA^*X^*) = \sup_{\substack{X \in \text{End}_{\mathbb{C}}(V_{\rho}) \\ \|X\|_{HS}=1}} \text{Tr}(X^*XAA^*) \\ &= \sup_{\substack{Y \in \text{End}_{\mathbb{C}}(V_{\rho}) \\ Y=Y^*, \text{Tr } Y=1}} \langle Y, AA^* \rangle_{HS} = \|AA^*\|_{op} = \|A\|_{op}^2. \end{aligned}$$

3.1.1 The special case of $G = \mathbb{Z}/q$

On a first reading of this section, the reader may wish to consider everything for the special case of $G = \mathbb{Z}/q$ for some integer $q \geq 2$. In that case, each representation is one-dimensional, and the representations can be indexed by $\ell = 0, 1, 2, \dots, q-1$. The Fourier transform is then particularly simple: the coefficients are scalar values

$$x_{\ell}(\sigma) = \sum_{a=0}^{q-1} \frac{n_a(\sigma)}{n} \omega^{a\ell},$$

where $\omega := e^{\frac{2\pi i}{q}}$ is a primitive q -th root of unity.

This special case already illustrates most of the main ideas while simplifying the estimates in some places (e.g. matrix inequalities we use will often be immediately obvious for scalars).

3.2 A stochastic difference equation for the n_a

For $a \in G$, we next analyze the behavior of $n_a(\sigma_t)$ over time. For convenience, we write $n_a(t) = n_a(\sigma_t)$. Let $\mathcal{F}_t$ denote the σ -field generated by the Markov chain $(\sigma_t)_{t \geq 0}$ up to time t . Then, our dynamics satisfy the equation

$$\mathbb{E}[n_a(t+1) - n_a(t) \mid \mathcal{F}_t] = \sum_{b \in G} \frac{n_{ab^{-1}}(t)n_b(t)}{2n(n-1)} + \sum_{b \in G} \frac{n_{ab}(t)n_b(t)}{2n(n-1)} - \frac{n_a(t)}{n}. \quad (7)$$

Note that $|n_a(t+1) - n_a(t)| \leq 1$ almost surely. Thus, for each $a \in G$, we can write the above as a stochastic difference equation

$$n_a(t+1) - n_a(t) = \sum_{b \in G} \frac{n_{ab^{-1}}(t)n_b(t)}{2n(n-1)} + \sum_{b \in G} \frac{n_{ab}(t)n_b(t)}{2n(n-1)} - \frac{n_a(t)}{n} + M_a(t+1), \quad (8)$$

where $\mathbb{E}[M_a(t+1) \mid \mathcal{F}_t] = 0$ and $|M_a(t)| \leq 2$ almost surely.

It is easiest to analyze this equation through the Fourier transform. Writing $x_\rho(t) = x_\rho(\sigma_t)$, we calculate from (8) that

$$x_\rho(t+1) - x_\rho(t) = \frac{1}{n-1} x_\rho(t) \left(\frac{x_\rho(t) + x_\rho(t)^*}{2} - \frac{n-1}{n} \right) + \widehat{M}_\rho(t+1),$$

where $\widehat{M}_\rho(t) := \frac{1}{n} \sum_{a \in G} M_a(t) \rho(a)$. For convenience, write

$$X_\rho(t) = \frac{1}{n-1} \left(\frac{x_\rho(t) + x_\rho(t)^*}{2} - \frac{n-1}{n} \right),$$

so that our equation becomes

$$x_\rho(t+1) - x_\rho(t) = x_\rho(t) X_\rho(t) + \widehat{M}_\rho(t+1). \quad (9)$$

Note that we have

$$\|x_\rho(t)\|_{\text{HS}} \leq \sqrt{d_\rho}, \quad \mathbb{E}[\widehat{M}_\rho(t+1) \mid \mathcal{F}_t] = 0, \quad \text{and} \quad \|\widehat{M}_\rho(t)\|_{\text{HS}} \leq \frac{2Q\sqrt{d_\rho}}{n},$$

and thus,

$$\|x(t)\|_{\widetilde{V}} \leq 1 \quad \text{and} \quad \|\widehat{M}(t)\|_{\widetilde{V}} \leq \frac{2Q}{n},$$

where $\widehat{M} = (\widehat{M}_\rho)_{\rho \in G^*}$.

3.3 A general estimate for stochastic difference equations

Before proving Lemma 2.2, we also need a technical lemma for controlling the behavior of stochastic difference equations, which will be used to analyze (9) as well as other similar equations.

Lemma 3.1 *Let $(z(t))_{t \geq 0}$ be a sequence of $[0, 1]$ -valued random variables adapted to a filtration $(\mathcal{F}_t)_{t \geq 0}$. Let $\varepsilon \in (0, 1)$ be a small constant, and let $\varphi : \mathbb{R}^+ \rightarrow (0, 1]$ be a non-decreasing function.*

Suppose that there are $\mathcal{F}_t$ -measurable random variables $M(t)$ for which

$$z(t+1) - z(t) \leq -\varepsilon\varphi(t+1)z(t) + M(t+1) \quad (10)$$

and which, for some constant D , satisfy the bounds

$$\mathbb{E}[M(t+1) \mid \mathcal{F}_t] \leq D\varepsilon\sqrt{\varepsilon}, \quad |M(t)| \leq D\varepsilon.$$

Then, for each t and each $\lambda > 0$, we have

$$\mathbb{P}\left(z(t) \geq \lambda\sqrt{\varepsilon} + e^{-\varepsilon \int_0^t \varphi(s) ds} \cdot z(0)\right) \leq C_{D,\varphi} e^{-c_{D,\varphi}\lambda^2}$$

for constants $c_{D,\varphi}, C_{D,\varphi}$ depending only on D and φ .

Proof Let us define for integers $t \geq 1$,

$$\Phi(t) := \varepsilon^{-1} \sum_{k=1}^t \log \frac{1}{1 - \varepsilon\varphi(k)}, \quad \text{and} \quad \Phi(0) := 0.$$

Taking conditional expectations in the inequality relating $z(t+1)$ to $z(t)$, we have

$$\mathbb{E}[z(t+1) \mid \mathcal{F}_t] \leq (1 - \varepsilon\varphi(t+1))z(t) + D\varepsilon\sqrt{\varepsilon}.$$

Rearranging and using the fact that $\varphi(t)$ is non-decreasing, we have

$$\begin{aligned} \mathbb{E}[z(t+1) \mid \mathcal{F}_t] - \frac{D\sqrt{\varepsilon}}{\varphi(0)} &\leq (1 - \varepsilon\varphi(t+1))z(t) - \frac{D\sqrt{\varepsilon}(1 - \varepsilon\varphi(t+1))}{\varphi(0)} \\ &\leq (1 - \varepsilon\varphi(t+1)) \left(z(t) - \frac{D\sqrt{\varepsilon}}{\varphi(0)} \right). \end{aligned}$$

Consequently,

$$Z_t := e^{\varepsilon\Phi(t)} \left(z(t) - \frac{D\sqrt{\varepsilon}}{\varphi(0)} \right)$$

is a supermartingale, and its increments are bounded by

$$|Z_{t+1} - Z_t| \leq e^{\varepsilon\Phi(t+1)} (|M(t+1)| + D\varepsilon) \leq 2D\varepsilon e^{\varepsilon\Phi(t+1)}. \quad (11)$$

Recall that φ is non-decreasing, so that for all $t \geq s \geq 0$, we have

$$\Phi(t) = \Phi(s) + \varepsilon^{-1} \sum_{k=s+1}^t \log \frac{1}{1 - \varepsilon\varphi(k)} \geq \Phi(s) + (t-s)\varphi(0).$$

Using this with (11), we see that the sum of the squares of the first t increments is at most

$$\begin{aligned}
\sum_{s=1}^t 4D^2 \varepsilon^2 e^{2\varepsilon \Phi(s)} &\leq 4D^2 \varepsilon^2 \sum_{s=1}^t e^{2\varepsilon \Phi(t) - 2\varepsilon \varphi(0)(t-s)} \leq 4D^2 \varepsilon^2 e^{2\varepsilon \Phi(t)} \cdot \frac{1}{1 - e^{-2\varepsilon \varphi(0)}} \\
&\leq 4D^2 \varepsilon^2 e^{2\varepsilon \Phi(t)} \cdot \frac{1}{1 - (1 - \frac{1}{2}\varepsilon \varphi(0))} = \frac{8D^2 \varepsilon}{\varphi(0)} \cdot e^{2\varepsilon \Phi(t)}.
\end{aligned}$$

By the Azuma–Hoeffding inequality, this yields

$$\mathbb{P}\left(Z_t \geq \lambda \sqrt{\varepsilon} e^{\varepsilon \Phi(t)} + Z_0\right) \leq \exp\left(-\frac{\varphi(0)\lambda^2 \varepsilon \cdot e^{2\varepsilon \Phi(t)}}{16D^2 \varepsilon \cdot e^{2\varepsilon \Phi(t)}}\right) = \exp\left(-\frac{\varphi(0)\lambda^2}{16D^2}\right),$$

which in turn implies

$$\mathbb{P}\left(z(t) \geq \frac{D\sqrt{\varepsilon}}{\varphi(0)} + e^{-\varepsilon \Phi(t)} z(0) + \lambda \sqrt{\varepsilon}\right) \leq \exp\left(-\frac{\varphi(0)\lambda^2}{16D^2}\right).$$

Finally, observe that $\Phi(t) \geq \sum_{k=1}^t \varphi(k) \geq \int_0^t \varphi(s) ds$. The result then follows upon shifting and rescaling of λ . $\square$

3.4 Proportion vector chain: Proof of Lemma 2.2

We first prove a bound for the Fourier coefficients $x_\rho(t)$.

Lemma 3.2 Consider any $\sigma \in \mathcal{S}_{\text{non}}(1/3)$ and any $\rho \in G^*$. We have a constant c_G depending only on G for which

$$\mathbb{P}_\sigma \left(\bigcup_{t=1}^{n^2} \left\{ \|x_\rho(t)\|_{\text{HS}} \geq \frac{1}{n^{1/8}} + e^{-c_G t/n} \cdot \|x_\rho(0)\|_{\text{HS}} \right\} \right) \leq \frac{1}{n^3}.$$

for all large enough n .

This immediately implies Lemma 2.2.

Proof of Lemma 2.2 With c_G defined as in Lemma 3.2, take $C_{G,\delta}$ large enough so that for any $T \geq C_{G,\delta} n$,

$$\frac{1}{n^{1/8}} + e^{-c_G T/n} \sqrt{d_\rho} \leq \delta.$$

Then, Lemma 3.2 and Plancherel's formula yield

$$\begin{aligned}
\mathbb{P}_\sigma(\sigma_T \notin \mathcal{S}_*(\delta)) &\leq \mathbb{P}_\sigma(\|x_\rho(T)\|_{\text{HS}} \geq \delta \text{ for some } \rho \in G^*) \\
&\leq \frac{Q}{n^3} \leq \frac{1}{n},
\end{aligned}$$

for large enough n , as desired. $\square$

We are now left with proving Lemma 3.2, which relies on the following bound on the operator norm.

Lemma 3.3 *There exists a positive constant γ_G depending on G such that for any $\rho \in G^*$ and any $\sigma \in \mathcal{S}_{non}(1/6)$,*

$$\|I_{d_\rho} + X_\rho(\sigma)\|_{op} \leq 1 - \frac{\gamma_G}{n}.$$

Proof Let Δ_G denote the set of all probability distributions on G , and for $c \in (0, 1)$, let $\Delta_G(c) \subset \Delta_G$ denote the set of all probability distributions μ such that $\mu(H) \leq 1 - c$ for all proper subgroups $H \subset G$.

Consider a representation $\rho \in G^*$, and consider the function $h : \Delta_G(1/6) \rightarrow \text{End}_{\mathbb{C}}(V_\rho)$ given by

$$h(\mu) = \sum_{a \in G} \mu(a) \frac{\rho(a) + \rho(a)^*}{2}.$$

Then, $h(\mu)$ is hermitian, and since ρ is unitary, we clearly have

$$\lambda(\mu) := \max_{v \in V_\rho, \|v\|=1} \langle h(\mu)v, v \rangle \leq 1.$$

We claim that $\lambda(\mu) < 1$ for each $\mu \in \Delta_G(c)$. Indeed, suppose the contrary. Then, there exists a non-zero vector $v \in V_\rho$ such that $\text{Re}\langle \rho(a)v, v \rangle = 1$ for all $a \in G$ with $\mu(a) > 0$. This implies that the support of μ is included in the subgroup

$$H = \{a \in G : \rho(a)v = v\}.$$

Since ρ is a (non-trivial) irreducible representation, H is a proper subgroup of G , and $\mu(H) = 1$, contradicting the assumption that $\mu \in \Delta_G(c)$.

Note that $\mu \mapsto \lambda(\mu)$ is continuous. We may define

$$\gamma_\rho := \max_{\mu \in \Delta_G(1/6)} \lambda(\mu) < 1 \quad \text{and} \quad \tilde{\gamma}_G := \max_{\rho \in G^*} \gamma_\rho < 1.$$

Then, we have for any $\sigma \in \mathcal{S}_{non}(1/6)$,

$$\frac{x_\rho(\sigma) + x_\rho(\sigma)^*}{2} = \sum_{a \in G} \frac{n_a(\sigma)}{n} \frac{\rho(a) + \rho(a)^*}{2} \preceq \tilde{\gamma}_G I_{d_\rho}.$$

Taking $0 < \gamma_G < 1 - \tilde{\gamma}_G$, and plugging this into the definition of X_ρ gives $X_\rho(\sigma) \preceq -\frac{\gamma_G}{n} I_{d_\rho}$. Note that $X_\rho(\sigma) \succeq -\frac{2}{n-1} I_{d_\rho}$. Combining these together gives the result. $\square$

Remark 3.4 A much more direct approach is possible in the case $G = \mathbb{Z}/q$. The condition $\sigma \in \mathcal{S}_{non}(1/6)$ implies that $n_0(\sigma) \leq \frac{5}{6}$. Then, we have

$$\operatorname{Re} x_\ell(\sigma) := \operatorname{Re} \sum_{a=0}^{q-1} \frac{n_a(\sigma)}{n} \omega^{a\ell} \leq \frac{5}{6} + \frac{1}{6} \max_{1 \leq a \leq q-1} \operatorname{Re} \omega^{a\ell} = \frac{5}{6} + \frac{1}{6} \cos \frac{2\pi}{q} < 1 - \gamma_G$$

for some positive γ_G . Some rearranging of equations then yields the desired result.

Proof of Lemma 3.2 Fix $\rho \in G^*$. Let $\mathcal{G}_t$ denote the event where for all $0 \leq s \leq t$, we have $\|I_{d_\rho} + X_\rho(s)\|_{op} \leq 1 - \frac{\gamma_G}{n}$, where γ_G is taken as in Lemma 3.3. Since our chain starts at $\sigma \in \mathcal{S}_{non}(1/3)$, Lemmas 2.1 and 3.3 together imply that

$$\mathbb{P}_\sigma(\mathcal{G}_{n^2}^c) \leq C_G n^2 e^{-n/10}.$$

Next, we turn to (9). Rearranging (9) and squaring, we have

$$\begin{aligned} \|x_\rho(t+1)\|_{\text{HS}}^2 &= \|x_\rho(t)(I_{d_\rho} + X_\rho(t))\|_{\text{HS}}^2 + \|\widehat{M}_\rho(t+1)\|_{\text{HS}}^2 \\ &\quad + 2\operatorname{Re}\langle x_\rho(t)(I_{d_\rho} + X_\rho(t)), \widehat{M}_\rho(t+1) \rangle_{\text{HS}} \end{aligned} \quad (12)$$

Let $z_t := \mathbf{1}_{\mathcal{G}_t} \|x_\rho(t)\|_{\text{HS}}^2$ and

$$M'(t+1) := \|\widehat{M}_\rho(t+1)\|_{\text{HS}}^2 + 2\operatorname{Re}\langle x_\rho(t)(I_{d_\rho} + X_\rho(t)), \widehat{M}_\rho(t+1) \rangle_{\text{HS}}.$$

Substituting into (12), we obtain

$$z_{t+1} \leq \|I_{d_\rho} + X_\rho(t)\|_{op}^2 \cdot z_t + \mathbf{1}_{\mathcal{G}_t} M'(t+1) \leq \left(1 - \frac{\gamma_G}{n}\right)^2 z_t + \mathbf{1}_{\mathcal{G}_t} M'(t+1).$$

Note that we have the bounds

$$\begin{aligned} \mathbb{E}[M'(t+1) \mid \mathcal{F}_t] &= \mathbb{E}[\|\widehat{M}_\rho(t+1)\|_{\text{HS}}^2 \mid \mathcal{F}_t] \leq \frac{4Q^2 d_\rho}{n^2} \\ |M'(t+1)| &\leq \|\widehat{M}_\rho(t+1)\|_{\text{HS}}^2 + 2\sqrt{d_\rho} \left(1 + \frac{1}{n(n-1)}\right) \|\widehat{M}_\rho(t+1)\|_{\text{HS}} \\ &\leq \frac{6Q^2 d_\rho}{n}. \end{aligned}$$

We now apply Lemma 3.1 with $\varepsilon = \frac{1}{n}$, $\varphi(t) = \gamma_G$, $D = 6Q^2 d_\rho$, and $\lambda = n^{1/4}$. This yields

$$\mathbb{P}\left(z_t \geq n^{-1/4} + e^{-\gamma_G t/n} \cdot z_0\right) \leq C'_G e^{-c'_G \sqrt{n}}.$$

Consequently,

$$\mathbb{P}_\sigma\left(\|x_\rho(t)\|_{\text{HS}} \geq n^{-1/8} + e^{-\gamma_G t/2n} \cdot \|x_\rho(0)\|_{\text{HS}}\right) \leq C'_G e^{-c'_G \sqrt{n}} + C_G n^2 e^{-n/10}.$$

The lemma with $c_G = \gamma_G/2$ then follows from union bounding over all $1 \leq t \leq n^2$ and taking n sufficiently large. $\square$

3.5 Proportion matrix chain: Proof of Lemma 2.4

We carry out a similar albeit more refined strategy to analyze the proportion matrix. Throughout this section, we assume our Markov chain $(\sigma_t)_{t \geq 0}$ starts at an initial state $\sigma_* \in \mathcal{S}_*\left(\frac{1}{4Q}\right)$. We again write $n_a(t) = n_a(\sigma_t)$ and $n_{a,b}(t) = n_{a,b}^{\sigma_*}(\sigma_t)$, and similar to before, the $n_{a,b}(t)$ satisfy the difference equation

$$\begin{aligned} n_{a,b}(t+1) - n_{a,b}(t) &= \sum_{c \in G} \frac{n_{a,bc^{-1}}(t)n_c(t)}{2n(n-1)} + \sum_{c \in G} \frac{n_{a,bc}(t)n_c(t)}{2n(n-1)} \\ &\quad - \frac{n_{a,b}(t)}{n} + M_{a,b}(t+1), \end{aligned} \quad (13)$$

where $\mathbb{E}[M_{a,b}(t+1) \mid \mathcal{F}_t] = 0$ and $|M_{a,b}(t)| \leq 2$ for all $t \geq 0$.

We can again analyze this equation via the Fourier transform. In this case, for each $a \in G$, we take the Fourier transform of $(n_{a,b}(t)/n_a(\sigma_*))_{b \in G}$. For $\rho \in G^*$, let

$$y_{a,\rho}(t) = y_{a,\rho}^{\sigma_*}(t) := \sum_{b \in G} \frac{n_{a,b}(t)}{n_a(\sigma_*)} \rho(b)$$

denote the Fourier coefficient at ρ . Let $\widehat{M}_{a,\rho}(t) := \frac{1}{n_a(\sigma_*)} \sum_{b \in G} M_{a,b}(t) \rho(b)$. Then, (13) becomes

$$y_{a,\rho}(t+1) - y_{a,\rho}(t) = y_{a,\rho}(t)X_\rho(t) + \widehat{M}_{a,\rho}(t+1). \quad (14)$$

Note that $\mathbb{E}_\sigma[\widehat{M}_{a,\rho}(t+1) \mid \mathcal{F}_t] = 0$. Also, since we assumed $\sigma_* \in \mathcal{S}_*\left(\frac{1}{4Q}\right)$, it follows that $\frac{n_a(\sigma_*)}{n} \geq \frac{1}{2Q}$. Thus, we also know $\|\widehat{M}_{a,\rho}(t+1)\|_{\text{HS}} \leq \frac{4Q^2\sqrt{d_\rho}}{n}$.

Again, our main step is a bound on the Fourier coefficients $y_{a,\rho}(t)$, which will also be useful later in proving Lemma 2.5.

Lemma 3.5 *Consider any $\sigma_*, \sigma'_* \in \mathcal{S}_*\left(\frac{1}{4Q}\right)$. There exist constants $c_G, C_G > 0$ depending only on G such that for all large enough n , we have*

$$\mathbb{P}_{\sigma'_*} \left(\|y_{a,\rho}^{\sigma_*}(t)\|_{\text{HS}} \geq R \left(\frac{1}{\sqrt{n}} + e^{-t/n} \|y_{a,\rho}^{\sigma_*}(0)\|_{\text{HS}} \right) \right) \leq e^{-\Omega_G(R^2) + O_G(1)} + \frac{2}{n^3}$$

for all t and $R > 0$.

The above lemma directly implies Lemma 2.4.

Proof of Lemma 2.4 We apply Lemma 3.5 to each $a \in G$ and $\rho \in G^*$. Recall that $T = \lceil \frac{1}{2}n \log n \rceil$, so that

$$\frac{1}{\sqrt{n}} + e^{-T/n} \|y_{a,\rho}^{\sigma_*}(0)\|_{\text{HS}} \leq \frac{2\sqrt{d_\rho}}{\sqrt{n}}.$$

Then, Lemma 3.5 implies

$$\mathbb{P}_{\sigma'_*} \left(\|y_{a,\rho}^{\sigma_*}(T)\|_{\text{HS}} \geq \frac{R}{\sqrt{n}} \right) \leq e^{-\Omega_G(R^2) + O_G(1)} + \frac{2}{n^3}.$$

Union bounding over all $a \in G$ and $\rho \in G^*$ and using the Plancherel formula, this yields

$$\begin{aligned} \mathbb{P}_{\sigma'_*} \left(\sigma_* \notin \mathcal{S}_* \left(\sigma_*, \frac{R}{\sqrt{n}} \right) \right) &\leq \mathbb{P}_{\sigma'_*} \left(\max_{a,\rho} \|y_{a,\rho}^{\sigma_*}(T)\|_{\text{HS}} \geq \frac{R}{\sqrt{n}} \right) \\ &\leq e^{-\Omega_G(R^2) + O_G(1)} + \frac{2Q^2}{n^3} \leq C_G e^{-R} + \frac{1}{n} \end{aligned}$$

for sufficiently large C_G and n . $\square$

We now prove Lemma 3.5. Before proceeding with the main proof, we need the following routine estimate as a preliminary lemma.

Lemma 3.6 Let $\theta_n : \mathbb{R}^d \rightarrow \mathbb{R}^+$ be the function given by $\theta_n(x) = \|x\| + \frac{1}{\sqrt{n}} e^{-\sqrt{n}\|x\|} - \frac{1}{\sqrt{n}}$. Then, we have the inequalities

$$\|\nabla \theta_n(x)\| \leq 1, \quad \theta_n(x+h) \leq \theta_n(x) + \langle h, \nabla \theta_n(x) \rangle + \frac{\sqrt{n}}{2} \|h\|^2.$$

Proof We can write $\theta_n(x) = f(\|x\|)$, where $f(r) = r + \frac{1}{\sqrt{n}} e^{-\sqrt{n}r} - \frac{1}{\sqrt{n}}$. By spherical symmetry, we have

$$\|\nabla \theta_n(x)\| = f'(\|x\|) = 1 - e^{-\sqrt{n}\|x\|} \leq 1,$$

which is the first inequality. Again by spherical symmetry, the eigenvalues of the Hessian $\nabla^2 \theta_n(x)$ can be directly computed to be $f''(\|x\|)$ and $f'(\|x\|)/\|x\|$. But these are bounded by

$$f''(r) \leq \sqrt{n} e^{-\sqrt{n}r} \leq \sqrt{n}, \quad f'(r)/r \leq \frac{1 - e^{-\sqrt{n}r}}{r} \leq \sqrt{n}.$$

Thus, $\nabla^2 \theta_n(x) \leq \sqrt{n}I$, and the second inequality follows from Taylor expansion. $\square$

Proof of Lemma 3.5 Let γ_G and c_G be the constants from Lemmas 3.3 and 3.2, respectively. Define the events

$$\mathcal{G}_t := \bigcap_{s=0}^t \left\{ X_\rho(\sigma_s) \leq -\frac{\gamma_G}{n} \right\},$$

$$\mathcal{G}'_t := \bigcap_{s=0}^t \left\{ X_\rho(\sigma_s) \leq -\frac{1 - \sqrt{d_\rho} e^{-c_G s/n} - 2n^{-1/8}}{n} \right\}.$$

Note that $\sigma'_* \in \mathcal{S}_*\left(\frac{1}{4Q}\right) \subseteq \mathcal{S}_{non}(1/3)$. Hence, by Lemmas 2.1 and 3.3, we have $\mathbb{P}(\mathcal{G}_{n^2}^c) \leq C_G n^2 e^{-n/10}$. We also have

$$\begin{aligned} X_\rho(s) &= \frac{1}{n-1} \left(\frac{x_\rho(s) + x_\rho(s)^*}{2} - \frac{n-1}{n} I_{d_\rho} \right) \leq -\frac{1}{n} \left(1 - \frac{n \|x_\rho(s)\|_{\text{HS}}}{n-1} \right) I_{d_\rho} \\ &\leq -\frac{1}{n} \left(1 - \|x_\rho(s)\|_{\text{HS}} - \frac{\sqrt{d_\rho}}{n-1} \right) I_{d_\rho}, \end{aligned}$$

where we have used the fact that $\left\| \frac{x_\rho(s) + x_\rho(s)^*}{2} \right\|_{\text{op}} \leq \|x_\rho(s)\|_{\text{op}} \leq \|x_\rho(s)\|_{\text{HS}}$.

Lemma 3.2 then implies that $\mathbb{P}(\mathcal{G}'_{n^2}) \leq \frac{1}{n^3}$. Thus, setting

$$\begin{aligned} \varphi(t) &:= \max(\gamma_G, 1 - \sqrt{d_\rho} e^{-c_G t/n} - 2n^{-1/8}), \\ \mathcal{H}_t &:= \mathcal{G}_t \cap \mathcal{G}'_t = \bigcap_{s=0}^t \left\{ X_\rho(\sigma_s) \leq -\frac{\varphi(t)}{n} \right\}, \end{aligned}$$

we conclude that

$$\mathbb{P}(\mathcal{H}_{n^2}^c) \leq \mathbb{P}(\mathcal{G}_{n^2}^c) + \mathbb{P}(\mathcal{G}'_{n^2}) \leq \frac{2}{n^3}$$

for all large enough n .

Next, we turn to (14) and apply θ_n to both sides, where we identify $\mathbb{C}^{d_\rho^2}$ with $\mathbb{R}^{2d_\rho^2}$. Using Lemma 3.6 and taking the conditional expectation, we obtain

$$\begin{aligned} \mathbb{E}[\theta_n(y_{a,\rho}(t+1)) \mid \mathcal{F}_t] &\leq \theta_n(y_{a,\rho}(t)(I_{d_\rho} + X_\rho(t))) + \frac{8Q^4 d_\rho}{n\sqrt{n}} \\ &\leq \theta_n(\|I_{d_\rho} + X_\rho(t)\|_{\text{op}} \cdot y_{a,\rho}(t)) + \frac{8Q^4 d_\rho}{n\sqrt{n}} \\ &\leq \|I_{d_\rho} + X_\rho(t)\|_{\text{op}} \cdot \theta_n(y_{a,\rho}(t)) + \frac{8Q^4 d_\rho}{n\sqrt{n}}, \end{aligned}$$

where the second inequality follows from the variational formula for operator norm (i.e. that $\|BA\|_{\text{HS}} \leq \|A\|_{\text{op}} \|B\|_{\text{HS}}$), and the third inequality follows from the fact that θ_n is convex with $\theta_n(0) = 0$. Thus, we may write

$$\theta_n(y_{a,\rho}(t+1)) \leq \|I_{d_\rho} + X_\rho(t)\|_{\text{op}} \cdot \theta_n(y_{a,\rho}(t)) + M'(t+1)$$

where

$$\mathbb{E}[M'(t+1) \mid \mathcal{F}_t] \leq \frac{8Q^4 d_\rho}{n\sqrt{n}}, \quad |M'(t+1)| \leq \frac{8Q^2 \sqrt{d_\rho}}{n}.$$

Now, let $z_t := \mathbf{1}_{\mathcal{H}_t} \theta_n(y_{a,\rho}(t))$, and note that since $X_\rho(\sigma) \geq -\frac{2}{n-1} I_{d_\rho}$, we have $\|I_{d_\rho} + X_\rho(t)\|_{\text{op}} \leq 1 - \frac{\varphi(t)}{n}$ whenever $\mathcal{H}_t$ holds. Thus,

$$z_{t+1} \leq \|I_{d_\rho} + X_\rho(t)\|_{\text{op}} \cdot z_t + \mathbf{1}_{\mathcal{H}_t} M'(t+1) \leq \left(1 - \frac{1}{n} \varphi(t)\right) z_t + \mathbf{1}_{\mathcal{H}_t} M'(t+1).$$

We may then apply Lemma 3.1 with $\varepsilon = \frac{1}{n}$ and $D = 8Q^4 d_\rho$. Note that

$$\int_0^t \varphi(s) ds \geq \left(1 - 2n^{-\frac{1}{8}}\right)t - \sqrt{d_\rho} \int_0^\infty e^{-\frac{c_G s}{n}} ds \geq t - O_G(n)$$

for all large enough n . Thus, Lemma 3.1 implies that

$$\mathbb{P}\left(z_t \geq \frac{\lambda}{\sqrt{n}} + C_G e^{-t/n} \cdot z_0\right) \leq C'_G e^{-c'_G \lambda^2}. \quad (15)$$

Consequently,

$$\begin{aligned} \mathbb{P}\left(\|y_{a,\rho}(t)\|_{\text{HS}} \geq R\left(\frac{1}{\sqrt{n}} + e^{-\frac{t}{n}} \|y_{a,\rho}(0)\|_{\text{HS}}\right)\right) \\ &\leq \mathbb{P}\left(\theta_n(y_{a,\rho}(t)) \geq \frac{R-1}{\sqrt{n}} + R e^{-\frac{t}{n}} \|y_{a,\rho}(0)\|_{\text{HS}}\right) \\ &\leq \mathbb{P}\left(\theta_n(y_{a,\rho}(t)) \geq \frac{R-1}{\sqrt{n}} + R e^{-\frac{t}{n}} \theta_n(y_{a,\rho}(0))\right) \\ &\leq \mathbb{P}\left(z_t \geq \frac{R-1}{\sqrt{n}} + R e^{-\frac{t}{n}} z_0\right) + \mathbb{P}(\mathcal{H}_{n^2}^c) \\ &\leq e^{-\Omega_G(R^2) + O_G(1)} + \frac{2}{n^3}, \end{aligned}$$

as desired. $\square$

4 Construction of the coupling: Proof of Lemma 2.5

For each $\delta > 0$, we define a subset of $\{0, 1, \dots, n\}^{G \times G}$ by

$$\mathcal{M}_\delta := \left\{ (n_{a,b})_{a,b \in G} : n_{a,b} \geq \frac{(1-\delta)n}{Q^2} \text{ for every } a, b \in G \text{ and } \sum_{a,b \in G} n_{a,b} = n \right\}.$$

Lemma 4.1 Consider a configuration $\sigma_* \in \mathcal{S}$ and a constant $0 < \delta \leq \frac{1}{2Q^2}$, and assume that $(1-\delta)n/Q^2$ is an integer. Let $(\sigma_t)_{t \geq 0}$ and $(\tilde{\sigma}_t)_{t \geq 0}$ be two product replacement chains started at σ and $\tilde{\sigma}$, respectively. Then, there exists a coupling $(\sigma_t, \tilde{\sigma}_t)$ of the Markov chains satisfying the following:

Let

$$D_t := \frac{1}{2} \sum_{a,b \in G} |n_{a,b}^{\sigma_*}(\sigma_t) - n_{a,b}^{\sigma_*}(\tilde{\sigma}_t)|.$$

Then, on the event $\{(n_{a,b}^{\sigma_*}(\sigma_t))_{a,b \in G}, (n_{a,b}^{\sigma_*}(\tilde{\sigma}_t))_{a,b \in G} \in \mathcal{M}_\delta\}$ and $\{D_t > 0\}$, one has

$$\mathbb{E}_{\sigma, \tilde{\sigma}}[D_{t+1} - D_t \mid \sigma_t, \tilde{\sigma}_t] \leq 0, \quad (16)$$

$$\mathbb{P}_{\sigma, \tilde{\sigma}}(D_{t+1} - D_t \neq 0 \mid \sigma_t, \tilde{\sigma}_t) \geq \frac{(1-\delta)^2}{4Q^3}. \quad (17)$$

Proof Let us abbreviate $n_{a,b}(t) = n_{a,b}^{\sigma_*}(\sigma_t)$ and $\tilde{n}_{a,b}(t) = n_{a,b}^{\sigma_*}(\tilde{\sigma}_t)$. Let $m_{a,b}(t) := \min(n_{a,b}(t), \tilde{n}_{a,b}(t))$. For each $a \in G$, we define the quantity

$$d_a(t) := \frac{1}{2} \sum_{b \in G} |n_{a,b}(t) - \tilde{n}_{a,b}(t)| = \sum_{b \in G} n_{a,b}(t) - \sum_{b \in G} m_{a,b}(t),$$

so that $D_t = \sum_{a \in G} d_a(t)$.

For accounting purposes, it is helpful to introduce two sequences

$$(x_1, x_2, \dots, x_n) \text{ and } (\tilde{x}_1, \tilde{x}_2, \dots, \tilde{x}_n)$$

of elements of $G \times G$. These sequences are chosen so that the number of x_k equal to (a, b) is exactly $n_{a,b}$, and similarly the number of $\tilde{x}_k$ equal to (a, b) is $\tilde{n}_{a,b}$. Moreover, we arrange their indices in a coordinated fashion, as described below.

We define three families of disjoint sets: $P_{a,b}$, Q_a , and $R_a \subset [n]$.

- For each $a, b \in G$, let $P_{a,b}$ be a set of size $(1-\delta)n/Q^2$ such that for any $k \in P_{a,b}$, we have $x_k = \tilde{x}_k = (a, b)$. (This is possible provided that $(n_{a,b}(t)), (\tilde{n}_{a,b}(t)) \in \mathcal{M}_\delta$ holds.)
- For each $a \in G$, let Q_a be a set of size $\sum_{b \in G} (m_{a,b} - |P_{a,b}|)$ such that for any $k \in Q_a$, $x_k = \tilde{x}_k = (a, b)$ for some b . (Note that Q_a may be empty.)

Fig. 1 Illustration of cases (i)–(iv)

- For each $a \in G$, let R_a be a set of size d_a such that for any $k \in R_a$, x_k and $\tilde{x}_k$ both have a as their first coordinate. (This R_a is well-defined since $\sum_b n_{a,b} = \sum_b \tilde{n}_{a,b}$ for each a ; it may also be empty.)

Define

$$P := \bigsqcup_{a,b \in G} P_{a,b}, \quad Q := \bigsqcup_{a \in G} Q_a, \quad R := \bigsqcup_{a \in G} R_a.$$

Suppose that $D_t > 0$, so that for some $a_*, b_*, b'_* \in G$ we have $n_{a_*, b_*} > \tilde{n}_{a_*, b_*}$ and $n_{a_*, b'_*} < \tilde{n}_{a_*, b'_*}$. Let us consider all possible ways to sample a pair of indices and a sign $(k, l, s) \in \{1, 2, \dots, n\}^2 \times \{\pm 1\}$ with $k \neq l$.

Suppose $x_k = (a_k, b_k)$ and $x_l = (a_l, b_l)$. We think of $(k, l, +1)$ as corresponding to a move on $(n_{a,b}(t))$ where n_{a_k, b_k} is decremented and $n_{a_k, (b_k \cdot b_l)}$ is incremented. Similarly, $(k, l, -1)$ corresponds to a move where n_{a_k, b_k} is decremented and $n_{a_k, (b_k \cdot b_l^{-1})}$ is incremented. We may also think of $(k, l, \pm 1)$ as corresponding to moves on $(\tilde{n}_{a,b}(t))$ in an analogous way.

We now analyze four cases, as illustrated in Fig. 1.

(i) *Case* $(k, l) \in (P \sqcup Q) \times (P \sqcup Q)$. For all but an exceptional situation described below, we apply the move corresponding to (k, l, s) to both states $(n_{a,b}(t))$ and $(\tilde{n}_{a,b}(t))$. In these cases, $D_{t+1} = D_t$.

We now describe the exceptional situation. Define

$$S = P_{a_*, b_*} \times \left(\bigsqcup_{c \in G} P_{c, (b_*^{-1} \cdot b'_*)} \right) \quad \text{and} \quad S' = P_{a_*, b'_*} \times \left(\bigsqcup_{c \in G} P_{c, \text{id}} \right).$$

Then, the exceptional situation occurs when $s = +1$ and $(k, l) \in S \sqcup S'$.

Take any bijection τ from S to S' . If $(k, l) \in S$, then we apply $(k, l, +1)$ to $(n_{a,b}(t))$ while applying $(\tau(k, l), +1)$ to $(\tilde{n}_{a,b}(t))$. This increments n_{a_*, b'_*} , decrements n_{a_*, b_*} , and has no effect on the $(\tilde{n}_{a,b}(t))$. The overall effect is that $D_{t+1} = D_t - 1$.

If instead $(k, l) \in S'$, then we apply $(k, l, +1)$ to $(n_{a,b}(t))$ and $(\tau^{-1}(k, l), +1)$ to $(\tilde{n}_{a,b}(t))$. A similar analysis shows that in this case $D_{t+1} = D_t + 1$.

The exceptional event occurs with probability $\frac{(1-\delta)^2}{2Q^3}$, and when it occurs, D_t increases or decreases by 1 with equal probability. Thus, the exceptional situation plays the role of introducing some unbiased fluctuation in D_t and gives us (17).

(ii) *Case* $(k, l) \in (Q \sqcup R) \times (Q \sqcup R)$ but $(k, l) \notin Q \times Q$. This occurs with probability

$$\frac{1}{n(n-1)} ((|Q| + |R|)(|Q| + |R| - 1) - |Q|(|Q| - 1))$$

which is at most

$$\frac{2}{n(n-1)} (|Q| + |R|)|R| = \frac{2\delta}{n-1} D_t.$$

Apply the move corresponding to (k, l, s) to both states. This increases D_t by at most 1. We will see later that the effect of this case is small compared to the other cases.

(iii) *Case* $(k, l) \in P \times R$. This occurs with probability

$$\frac{1}{n(n-1)} |P||R| = \frac{1-\delta}{n-1} D_t.$$

Apply the move corresponding to (k, l, s) to both states. Again, this increases D_t by at most 1, but there is also a chance not to increase.

Suppose that $x_l = (a_1, b_1)$ and $\tilde{x}_l = (a_1, \tilde{b}_1)$, and suppose that $k \in P_{a_2, b_2}$. Then the move has the effect of decreasing n_{a_2, b_2} and $\tilde{n}_{a_2, b_2}$ while increasing $n_{a_2, (b_2 \cdot b_1^s)}$ and $\tilde{n}_{a_2, (b_2 \cdot \tilde{b}_1^s)}$. Note that conditioned on this case happening, (a_2, b_2) is distributed uniformly over $G \times G$. When $(a_2, (b_2 \cdot \tilde{b}_1^s)) = (a_*, b_*)$ or $(a_2, (b_2 \cdot b_1^s)) = (a_*, b_*')$, the move does not increase D_t . Therefore there is at least a $2/Q^2$ chance that D_t is actually not increased. Hence, the probability that D_t is increased by 1 is at most

$$\left(1 - \frac{2}{Q^2}\right) \frac{1-\delta}{n-1} D_t.$$

(iv) *Case* $(k, l) \in R \times P$. This occurs with probability

$$\frac{1}{n(n-1)} |R||P| = \frac{1-\delta}{n-1} D_t.$$

Suppose that $x_k = (a, b)$ and $\tilde{x}_k = (a, \tilde{b})$. Let τ be a permutation of P such that for $l \in P_{a,c}$, one has $\tau(l) \in P_{a, \tilde{b}^{-1} \cdot b \cdot c^s}$. Then apply (k, l, s) to $(n_{a,b}(t))$ and apply $(k, \tau(l), s)$ to $(\tilde{n}_{a,b}(t))$. This always decreases D_t by 1.

Let us now summarize what we know when $(n_{a,b}(t)), (\tilde{n}_{a,b}(t)) \in \mathcal{M}_\delta$ and $D_t > 0$. From Cases (i), (ii), and (iii), we have

$$\mathbb{P}_{\sigma, \tilde{\sigma}}(D_{t+1} = D_t + 1 \mid \sigma_t, \tilde{\sigma}_t) \leq \left(1 - \frac{2(1-\delta)}{Q^2} + \delta\right) \frac{D_t}{n-1} + \frac{(1-\delta)^2}{4Q^3}.$$

From Cases (i) and (iv), we have

$$\mathbb{P}_{\sigma, \tilde{\sigma}}(D_{t+1} = D_t - 1 \mid \sigma_t, \tilde{\sigma}_t) \geq (1 - \delta) \frac{D_t}{n - 1} + \frac{(1 - \delta)^2}{4Q^3}.$$

Therefore, if $0 < \delta \leq \frac{1}{2Q^2}$, then

$$\mathbb{E}_{\sigma, \tilde{\sigma}}[D_{t+1} - D_t \mid \sigma_t, \tilde{\sigma}_t] \leq 0,$$

verifying (16).

To fully define the coupling, when $D_t = 0$, we can couple σ_t and σ_t to be identical, and if either $(n_{a,b}(t)) \notin \mathcal{M}_\delta$ or $(\tilde{n}_{a,b}(t)) \notin \mathcal{M}_\delta$, we may run the two chains independently. $\square$

Proof of Lemma 2.5 Since $\sigma \in \mathcal{S}_*\left(\sigma_*, \frac{R}{\sqrt{n}}\right)$, we must have for each $a \in G$ and $\rho \in G^*$ that $\|y_{a,\rho}^{\sigma_*}(\sigma)\|_{\text{HS}} \leq \frac{R}{\sqrt{n}}$. Note that for large enough n , we have $\mathcal{S}_*\left(\sigma_*, \frac{R}{\sqrt{n}}\right) \subseteq \mathcal{S}_*\left(\frac{1}{5Q^3}\right)$. Thus, we may apply Lemma 3.5 to obtain

$$\mathbb{P}\left(\bigcup_{t=0}^{n^2} \left\{ \|y_{a,\rho}^{\sigma_*}(\sigma_t)\|_{\text{HS}} \geq \frac{1}{5Q^3} \right\}\right) \leq n^2 \left(e^{-\Omega_G(n) + O_G(1)} + \frac{2}{n^3} \right) \leq \frac{3}{n} \quad (18)$$

for large enough n . Define the event

$$\mathcal{G}_t := \left\{ \sigma_s \in \mathcal{S}_*\left(\sigma_*, \frac{1}{5Q^3}\right) \text{ for all } 1 \leq s \leq t \right\}.$$

The Plancherel formula applied to (18) implies that $\mathbb{P}(\mathcal{G}_{n^2}^c) \leq \frac{3Q^2}{n}$. We may analogously define an event $\tilde{\mathcal{G}}_t$ for $\tilde{\sigma}$ and let $\mathcal{A}_t := \mathcal{G}_t \cap \tilde{\mathcal{G}}_t$. Thus, $\mathbb{P}(\mathcal{A}_{n^2}^c) \leq \frac{6Q^2}{n}$.

Pick $\delta' \in \left(\frac{2}{5Q^2}, \frac{3}{7Q^2}\right)$ so that $(1 - \delta')n/Q^2$ is an integer. Note that when $\mathcal{A}_t$ holds, we have

$$\sigma_t \in \mathcal{S}_*\left(\sigma_*, \frac{1}{5Q^3}\right) \quad \text{and} \quad \sigma_* \in \mathcal{S}_*\left(\frac{1}{5Q^3}\right) \implies (n_{a,b}(t)) \in \mathcal{M}_{\frac{2}{5Q^2}} \subseteq \mathcal{M}_{\delta'},$$

and similarly $\tilde{\sigma}_t \in \mathcal{M}_{\delta'}$.

Thus, we may invoke Lemma 4.1 to give a coupling between σ and $\tilde{\sigma}$ where on the event $\mathcal{A}_t$, the quantity D_t is more likely to decrease than increase. Letting $\mathbf{D}_t := \mathbf{1}_{\mathcal{A}_t} D_t$, we see that $(\mathbf{D}_t)$ is a supermartingale with respect to $(\mathcal{F}_t)$.

Define

$$\tau := \min\{t \geq 0 : D_t = 0\}, \quad \tilde{\tau} := \min\{t \geq 0 : \mathbf{D}_t = 0\}.$$

Then, Lemma 4.1 ensures that on the event $\{\tilde{\tau} > t\}$, we have $\text{Var}(\mathbf{D}_{t+1} \mid \mathcal{F}_t) \geq \alpha^2$, where $\alpha^2 := \left(1 - \frac{1}{Q^2}\right) \frac{(1-\delta')^2}{4Q^3}$. By [13, Proposition 17.20], for every $u > 12/\alpha^2$,

$$\mathbb{P}(\tilde{\tau} > u) \leq \frac{4\mathbf{D}_0}{\alpha\sqrt{u}}. \quad (19)$$

Recall that $T = \lceil \beta n \rceil$ and $D_0 \leq \sqrt{Q}R\sqrt{n}$. As long as β is large enough, we may apply (19) with $u = T$ to get

$$\mathbb{P}_{\sigma, \tilde{\sigma}}(\tau > T) \leq \frac{16Q^2R}{(1-\delta')\sqrt{\beta}} + \mathbb{P}(\mathcal{A}_T^c) \leq \frac{32Q^2R}{\sqrt{\beta}}$$

for all large enough n , as desired. $\square$

5 Proof of Theorem 1.1 (2)

The lower bound is proved essentially by showing that the estimates of Lemmas 2.1 and 2.4 cannot be improved. Let $a_1, a_2, \dots, a_k$ be a set of generators for G . Let $\sigma_\star \in \mathcal{S}$ be the configuration given by

$$\sigma_\star(i) = \begin{cases} a_i & \text{if } i \leq k, \\ 0 & \text{otherwise.} \end{cases}$$

We will analyze the Markov chain started at $\sigma_\star$ and show that it does not mix too fast. Recall from Sect. 2 the notation

$$n_{\text{non}}^{\{\text{id}\}}(\sigma) = |\{i \in [n] : \sigma(i) \neq \text{id}\}|$$

for the number of sites in σ that do not contain the identity. We first show that if we run the chain for slightly less than $n \log n$ steps, most of the sites will still contain the identity.

Lemma 5.1 *Let $T := \lfloor n \log n - Rn \rfloor$. Then,*

$$\mathbb{P}_{\sigma_\star} \left(n_{\text{non}}^{\{\text{id}\}}(\sigma_T) \geq \frac{n}{3} \right) \leq \frac{4Q^2}{R^2}.$$

Proof Recall that in one step of our Markov chain, we pick two indices $i, j \in [n]$ and replace $\sigma(i)$ with $\sigma(i) \cdot \sigma(j)$ or $\sigma(i) \cdot \sigma(j)^{-1}$. The only way for $n_{\text{non}}^{\{\text{id}\}}(\sigma_t)$ to increase after this step is if $\sigma(j) \neq \text{id}$. Thus,

$$\mathbb{P}(n_{\text{non}}^{\{\text{id}\}}(\sigma_{t+1}) = n_{\text{non}}^{\{\text{id}\}}(\sigma_t) + 1 \mid n_{\text{non}}^{\{\text{id}\}}(\sigma_t)) \leq \frac{n_{\text{non}}^{\{\text{id}\}}(\sigma_t)}{n}. \quad (20)$$

Let $\tau := \min\{t \geq 0 : n_{non}^{\{id\}}(\sigma_t) \geq \frac{n}{3}\}$ be the first time that $n_{non}^{\{id\}}(\sigma_t)$ is at least $\frac{n}{3}$. We have that $n_{non}^{\{id\}}(\sigma_\star) = k$, so it follows from (20) that τ stochastically dominates the sum

$$G := \sum_{s=k}^{\lfloor n/3 \rfloor} G_s,$$

where the G_s are independent geometric variables with success probability $\frac{s}{n}$. Note that we have the bounds

$$\mathbb{E}G = \sum_{s=k}^{\lfloor n/3 \rfloor} \frac{n}{s} \geq n \left(\log \left\lfloor \frac{n}{3} \right\rfloor - \log k \right), \quad \text{Var}(G) = \sum_{s=k}^{\lfloor n/3 \rfloor} \frac{n(n-s)}{s^2} \leq n^2.$$

Hence,

$$\begin{aligned} \mathbb{P}(\tau < T) &\leq \mathbb{P}(G < T) \leq \mathbb{P}(G < \mathbb{E}G + n \log(3k) - Rn) \\ &\leq \frac{n^2}{n^2(R - \log(3k))^2} \leq \frac{4}{R^2} \end{aligned}$$

for $R \geq 2Q \geq 2 \log(3k)$. On the other hand, the bound claimed in the lemma is trivial for $R \leq 2Q$, so we have completed the proof. $\square$

Next, we show that it really takes about $\frac{1}{2}n \log n$ steps for the Fourier coefficients x_ρ to decay to $O\left(\frac{1}{\sqrt{n}}\right)$, as suggested by Lemma 2.4. Note that it suffices here to analyze the x_ρ instead of the $y_{a,\rho}$, which simplifies our analysis. Actually, it suffices to consider (the real part of) the trace of x_ρ . Here the orthogonality of characters reads $\frac{1}{Q} \sum_{a \in G} \text{Tr } \rho(a) = 0$, and it takes about $\frac{1}{2}n \log n$ steps for $\text{ReTr } x_\rho(t)$ to decay to $O\left(\frac{1}{\sqrt{n}}\right)$.

Lemma 5.2 Consider any $\rho \in G^*$ and any $R > 5$. Let $T := \lfloor \frac{1}{2}n \log n - Rn \rfloor$, and suppose that $\sigma \in \mathcal{S}$ satisfies $n_{non}^{\{id\}}(\sigma) \leq \frac{n}{3}$. Then,

$$\mathbb{P}_\sigma \left(\|x_\rho(\sigma_T)\|_{\text{HS}} \leq \frac{R}{\sqrt{n}} \right) \leq \frac{4Q^2}{R^2}.$$

Proof Let $z(t) := (1/d_\rho) \text{Tr}(x_\rho(t) + x_\rho(t)^*)/2$. Then, noting that (9) also holds for $x_\rho(t)^*$ since $x_{\rho^*}(t) = x_\rho(t)^*$, we have

$$z(t+1) - z(t) = \frac{1}{n-1} \frac{1}{d_\rho} \text{Tr} \left(\frac{x_\rho(t) + x_\rho(t)^*}{2} \right)^2 - \frac{1}{n} z(t) + M(t+1),$$

where

$$\mathbb{E}[M(t+1) \mid \mathcal{F}_t] = 0 \quad \text{and} \quad |M(t)| \leq \frac{2Q}{n}.$$

Here we have

$$\frac{1}{d_\rho} \text{Tr} \left(\frac{x_\rho(t) + x_\rho(t)^*}{2} \right)^2 \geq z(t)^2.$$

We compare $z(t)$ to another process $(w(t))_{t \geq 0}$ defined by $w(0) := \frac{1}{3}$ and

$$w(t+1) := \left(1 - \frac{1}{n}\right) w(t) + M(t+1). \quad (21)$$

We will show by induction that $z(t) \geq w(t)$ for all t . For the base case, note that since $n_{\text{non}}^{\{\text{id}\}}(\sigma) \leq \frac{n}{3}$, we have

$$z(0) = \frac{1}{d_\rho} \text{ReTr} \sum_{a \in G} \frac{n_a(t)}{n} \cdot \rho(a) \geq \frac{2}{3} - \frac{1}{3} = \frac{1}{3}.$$

Suppose now that $z(t) \geq w(t)$. Then,

$$\begin{aligned} z(t+1) &\geq z(t) + \frac{1}{n-1} z(t)^2 - \frac{1}{n} z(t) + M(t+1) \\ &\geq \left(1 - \frac{1}{n}\right) w(t) + M(t+1) = w(t+1), \end{aligned}$$

completing the induction.

It now suffices to lower bound $w(T)$. To this end, we first note that applying (21) repeatedly and taking expectations, we obtain

$$\mathbb{E}w(T) = \left(1 - \frac{1}{n}\right)^T \cdot \frac{1}{3} \geq \frac{e^{-R}}{6\sqrt{n}} \geq \frac{2R}{\sqrt{n}}.$$

In order to calculate the variance of $w(T)$, we can also square (21) and take the expectation, which gives us

$$\begin{aligned} \text{Var}(w(T)) &= \mathbb{E}w(T)^2 - (\mathbb{E}w(T))^2 \\ &\leq \left(1 - \frac{1}{n}\right)^{2T} \cdot \frac{1}{9} + n \cdot \left(\frac{2Q}{n}\right)^2 - \left(\left(1 - \frac{1}{n}\right)^T \cdot \frac{1}{3}\right)^2 \\ &= \frac{4Q^2}{n}. \end{aligned}$$

Then, by Chebyshev's inequality, we have

$$\begin{aligned}\mathbb{P}_{\sigma}\left(\|x_{\rho}(\sigma_T)\|_{\text{HS}} \leq \frac{R}{\sqrt{n}}\right) &\leq \mathbb{P}\left(z(T) \leq \frac{R}{\sqrt{n}}\right) \leq \mathbb{P}\left(w(T) \leq \frac{R}{\sqrt{n}}\right) \\ &\leq \frac{4Q^2/n}{(R/\sqrt{n})^2} = \frac{4Q^2}{R^2},\end{aligned}$$

as desired. $\square$

Proof of Theorem 1.1 (2) Let $T = T_1 + T_2$, where $T_1 := \lfloor n \log n - \beta n \rfloor$ and $T_2 := \lfloor \frac{1}{2} n \log n - \beta n \rfloor$. Fix any $\rho \in G^*$. By Lemma 5.1 followed by Lemma 5.2, we have for large enough β that

$$\mathbb{P}_{\sigma_*}\left(\sigma_T \in \mathcal{S}_*\left(\frac{\beta}{\sqrt{n}}\right)\right) \leq \mathbb{P}_{\sigma_*}\left(\|x_{\rho}(T)\|_{\text{HS}} \leq \sqrt{\frac{Q}{d_{\rho}}} \frac{\beta}{\sqrt{n}}\right) \leq \frac{8Q^2}{\beta^2}.$$

On the other hand, Lemma 2.6 tells us that

$$\pi\left(\mathcal{S}_*\left(\frac{\beta}{\sqrt{n}}\right)\right) \geq 1 - \frac{c_G}{\beta^2}.$$

Consequently,

$$d_{\sigma_*}(T) \geq 1 - \frac{c_G}{\beta^2} - \frac{8Q^2}{\beta^2},$$

which tends to 1 as $\beta \rightarrow \infty$, establishing (2). $\square$

Acknowledgements This work was initiated while R.T. and A.Z. were visiting Microsoft Research in Redmond. They thank Microsoft Research for the hospitality. R.T. was also visiting the University of Washington in Seattle and thanks Professor Christopher Hoffman for making his visit possible. R.T. is supported by JSPS Grant-in-Aid for Young Scientists (B) 17K14178 and partially by JST, ACT-X Grant Number JPMJAX190J, Japan. A.Z. is supported by a Stanford Graduate Fellowship. The authors thank an anonymous referee for helpful comments.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Andrén, D., Hellström, L., Markström, K.: On the complexity of matrix reduction over finite fields. *Adv. Appl. Math.* **39**(4), 428–452 (2007)

2. Ben-Hamou, A., Peres, Y.: Cutoff for a stratified random walk on the hypercube. *Electron. Commun. Probab.* **23**(32), 1–10 (2018)
3. Chung, F.R.K., Graham, R.L.: Random walks on generating sets for finite groups. *Electron. J. Comb.* **4**(2) (1997)
4. Chung, F.R.K., Graham, R.L.: Stratified random walks on the n -cube. *Random Struct. Algorithms* **11**(3), 199–222 (1997)
5. Christofides, D.: The asymptotic complexity of matrix reduction over finite fields. *arXiv preprint arXiv:1406.5826* (2014)
6. Celler, F., Leedham-Green, C.R., Murray, S.H., Niemeyer, A.C., O'Brien, E.A.: Generating random elements of a finite group. *Commun. Algebra* **23**(13), 4931–4948 (1995)
7. Diaconis, P.: *Group Representations in Probability and Statistics*. Institute of Mathematical Statistics, Lecture Notes-Monograph Series, vol. 11. Institute of Mathematical Statistics, Hayward (1988)
8. Diaconis, P., Saloff-Coste, L.: Walks on generating sets of abelian groups. *Probab. Theory Relat. Fields* **105**(3), 393–421 (1996)
9. Diaconis, P., Saloff-Coste, L.: Walks on generating sets of groups. *Inventiones Mathematicae* **134**(2), 251–299 (1998)
10. Holt, D.F., Rees, S.: An implementation of the Neumann–Praeger algorithm for the recognition of special linear groups. *Exp. Math.* **1**(3), 237–242 (1992)
11. Kassabov, M.: Kazhdan constants for $SL_n(\mathbb{Z})$. *Int. J. Algebra Comput.* **15**(5–6), 971–995 (2005)
12. Lubotzky, A., Pak, I.: The product replacement algorithm and Kazhdan's property (T). *J. Am. Math. Soc.* **14**(2), 347–363 (2001)
13. Levin, D., Peres, Y., Wilmer, E.: *Markov Chains and Mixing Times*, 2nd edn. American Mathematical Society, Providence (2017)
14. Pak, I.: The product replacement algorithm is polynomial. In: *Proceedings of the 41st Annual Symposium on the Foundations of Computer Science (FOCS)*, pp. 476–485. IEEE (2000)
15. Pak, I.: What do we know about the product replacement algorithm? In: *Groups and Computation, III* (Columbus, OH, 1999), Volume 8 of Ohio State University Mathematical Research Institute Publications, pp. 301–347. de Gruyter, Berlin (2001)
16. Sotiraki, A.: Authentication protocol using trapdoored matrices. Master Thesis, Massachusetts Institute of Technology (2016)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.