

Mixing time of fractional random walk on finite fields

Jimmy He* Huy Tuan Pham† Max Wenqiang Xu‡

Abstract

We study a random walk on \mathbb{F}_p defined by $X_{n+1} = 1/X_n + \varepsilon_{n+1}$ if $X_n \neq 0$, and $X_{n+1} = \varepsilon_{n+1}$ if $X_n = 0$, where ε_{n+1} are independent and identically distributed. This can be seen as a non-linear analogue of the Chung–Diaconis–Graham process. We show that the mixing time is of order $\log p$, answering a question of Chatterjee and Diaconis.

Keywords: mixing times; finite field; spectral gap.

MSC2020 subject classifications: Primary 60J10, Secondary 11T23; 05C81.

Submitted to EJP on March 6, 2022, final version accepted on September 26, 2022.

1 Introduction

This paper studies a non-linear analogue of the Chung–Diaconis–Graham process, defined on \mathbb{F}_p by

$$X_{n+1} = aX_n + \varepsilon_{n+1},$$

where $a \in \mathbb{F}_p^\times$ is fixed and the ε_i are independent and identically distributed. The mixing time of this Markov chain has been extensively studied, and it is now known that for certain a (such as $a = 2$) and almost all integers p such that $(a, p) = 1$, cutoff occurs at time $c \log p$ for an explicit constant c [18].

Since simple random walk on \mathbb{F}_p requires order p^2 steps to mix (see [28] for example), the Chung–Diaconis–Graham process provides an explicit example of a random walk where applying a deterministic bijection (in this case, $x \mapsto ax$) between steps of the walk exponentially speeds up mixing. This was studied in [12], where it was asked whether other explicit examples could be provided.

In this paper, we consider the random walk on \mathbb{F}_p defined by

$$X_{n+1} = \iota(X_n) + \varepsilon_{n+1},$$

where $\iota(x) = 1/x$ if $x \neq 0$ and $\iota(0) = 0$, and the ε_i are independent and identically distributed. Chatterjee and Diaconis asked about the order of the mixing time for this

*Department of Mathematics, MIT, Cambridge, MA. E-mail: jimmyhe@mit.edu

†Department of Mathematics, Stanford University, Stanford, CA. E-mail: huypham@stanford.edu

‡Department of Mathematics, Stanford University, Stanford, CA. E-mail: maxxu@stanford.edu

walk [12], which was originally suggested by Soundararajan. We solve this problem by showing that this random walk mixes in order $\log p$ steps. Since simple random walk on \mathbb{F}_p mixes in order p^2 steps, this gives an exponential speedup for the mixing time, and provides an explicit example of when adding a deterministic function between steps of a Markov chain gives an exponential speedup.

This model is similar in spirit to random walk on \mathbb{F}_p given by moving from x to one of $x + 1$, $x - 1$ or x^{-1} , which is known to be an expander (see Theorem 8.8 of [36] for example). Our results are not directly implied by this, but similar ideas underlie the results for both models. We discuss the relationship with this model in more detail in Remark 1.3.

The proof uses comparison theory developed by Smith [34] to relate the random walk on \mathbb{F}_p to a random walk on the projective line $\mathbb{P}^1(\mathbb{F}_p)$. The walk on $\mathbb{P}^1(\mathbb{F}_p)$ is the quotient of a random walk on $SL_2(\mathbb{F}_p)$ which is known to have a constant order spectral gap by results of Bourgain and Gamburd [6]. While our methods identify the order of the mixing time, they are not strong enough to identify the constant or obtain cutoff, and we leave this as an open problem.

Our result also implies bounds on the number of solutions to the congruence $xy \equiv 1 \pmod{p}$ for $x \in I$ and $y \in J$ intervals of the same length, although these bounds are weaker than the ones established in [14].

1.1 Main results

We now formally state our main results. For two probability measures μ and ν , let

$$\|\mu - \nu\|_{TV} = \sup_A |\mu(A) - \nu(A)|$$

denote the *total variation distance*. Let X_n be an ergodic Markov chain with transition matrix P and stationary distribution π . Let $P^n(x, y)$ denote the probability that $X_n = y$ given that $X_0 = x$. Then the *mixing time* is defined to be

$$t_{mix}(\varepsilon) = \inf \left\{ n : \sup_x \|P^n(x, \cdot) - \pi\|_{TV} \leq \varepsilon \right\}.$$

Let $\iota : \mathbb{F}_p \rightarrow \mathbb{F}_p$ be the function defined by $\iota(x) = 1/x$ if $x \neq 0$ and $\iota(0) = 0$. Let ε_i be independent and identically distributed on \mathbb{Z} . Let K denote the transition matrix for the Markov chain on \mathbb{F}_p defined by

$$X_{n+1} = \iota(X_n) + \varepsilon_{n+1}. \tag{1.1}$$

We show that for this Markov chain, $t_{mix}(\varepsilon)$ is $\Theta(\log p)$, where the implied constants are allowed to depend on both μ and ε . Informally, this means that order $\log p$ steps are necessary and sufficient for the Markov chain to converge to its stationary distribution. The following theorems give the desired lower and upper bounds on t_{mix} respectively.

Theorem 1.1 (Lower bound). *Let μ denote a probability measure on \mathbb{Z} , and assume that*

$$H(\mu) = \sum -\mu(x) \log \mu(x) < \infty.$$

Let K be the transition matrix for the Markov chain on \mathbb{F}_p defined by

$$X_{n+1} = \iota(X_n) + \varepsilon_{n+1},$$

where the ε_n are independent and identically distributed according to μ . Let π denote the uniform distribution on \mathbb{F}_p . Then for all x in \mathbb{F}_p

$$\|K^n(x, \cdot) - \pi\|_{TV} \geq 1 - \frac{nH(\mu) + \log 2}{\log p}.$$

Theorem 1.2 (Upper bound). *Let μ denote a probability measure on \mathbb{Z} whose support contains at least two values. Let K be the transition matrix for the Markov chain on \mathbb{F}_p defined by*

$$X_{n+1} = \iota(X_n) + \varepsilon_{n+1},$$

where the ε_n are independent and identically distributed according to μ . Let π denote the uniform distribution on \mathbb{F}_p . Then there exists a constant $C > 0$ depending only on μ such that for all x in \mathbb{F}_p and sufficiently large p ,

$$\|K^n(x, \cdot) - \pi\|_{TV} \leq \frac{\sqrt{p}}{2} e^{-Cn}.$$

The lower bound follows easily from entropy considerations. We make an assumption that the entropy of μ is finite, which cannot be removed entirely (see Remark 2.2). However, this assumption is relatively weak, and indeed all finitely supported distributions have finite entropy.

We now describe our plan to prove the upper bound. We establish the upper bound by connecting the Markov chain X_n with a Markov chain on $\mathbb{P}^1(\mathbb{F}_p)$, which in turn is a projection of a Markov chain on a Cayley graph of $\text{SL}_2(\mathbb{F}_p)$. We next sketch this connection.

Recall that given a group G and a generating set S , we define the *Cayley graph* $\mathcal{G}(G, S)$ as the graph with vertex set G and two elements x and y are connected if and only if $x = \sigma y$ for some $\sigma \in S$. We use $A(\mathcal{G})$ to denote the normalised adjacency matrix of \mathcal{G} , whose largest eigenvalue is 1. Let S be a subset of $\text{SL}_2(\mathbb{Z})$ so that S_p generates $\text{SL}_2(\mathbb{F}_p)$, where S_p is the set $S \bmod p$. Results of Bourgain and Gamburd [6] give conditions for when $\mathcal{G}(\text{SL}_2(\mathbb{F}_p), S_p)$ has a constant spectral gap, i.e., $\lambda_2(A(\mathcal{G}(\text{SL}_2(\mathbb{F}_p), S_p))) \leq 1 - c$ for some constant $c > 0$ independent of p , and thus also a $O(\log p)$ mixing time.

The random walk on $\mathcal{G}(\text{SL}_2(\mathbb{F}_p), S_p)$ has a natural projection on $\mathbb{P}^1(\mathbb{F}_p)$, defined by the action

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix} \cdot x = \frac{ax + b}{cx + d}.$$

This can also be seen as a random walk on the Schreier graph corresponding to this action, which is defined as the graph on vertices $\mathbb{P}^1(\mathbb{F}_p)$ where x, y are connected if $y = gx$ for some $g \in S_p$. Since the spectral gaps of the walks on $\text{SL}_2(\mathbb{F}_p)$ and $\mathbb{P}^1(\mathbb{F}_p)$ are related, this gives an $O(\log p)$ mixing time for the walk on $\mathbb{P}^1(\mathbb{F}_p)$.

Finally, we connect the Markov chain on $\mathbb{P}^1(\mathbb{F}_p)$ with the fractional Markov chain X_n on \mathbb{F}_p . With an appropriate choice of the set S , the Markov chain on $\mathbb{P}^1(\mathbb{F}_p)$ induced by the random walk on $\mathcal{G}(\text{SL}_2(\mathbb{F}_p), S_p)$ and X_n have very similar transition probabilities. The mixing time bound of X_n can thus be obtained from that of the chain on $\mathbb{P}^1(\mathbb{F}_p)$ via comparison theory of Markov chains on different state spaces.

A key input to the above plan is the spectral gap of random walks on Cayley graphs of $\text{SL}_2(\mathbb{F}_p)$, which we discuss in further detail in Section 1.4.

Remark 1.3. The random walk on \mathbb{F}_p obtained by moving from x to $x + 1$, $x - 1$ or x^{-1} with equal probability has been previously studied. In Section 3.3 of [31], the walk on $\mathbb{P}^1(\mathbb{F}_p)$ is shown to have a constant order spectral gap, and Theorem 8.8 of [36] states that the same is true on \mathbb{F}_p . This model can be seen as an explicit version of the speedup proposed in [24], who study random walks on graphs with an additional random perfect matching.

The proofs in both models also follow similar ideas, although the random walk we study is more challenging due to the lack of reversibility. In particular, our results imply that the walks above have a constant spectral gap, while the reverse is not true. Theorem 8.8 of [36] is stated without proof, and we could not find an explanation for the reduction

between the study of the random walk on \mathbb{F}_p to $\mathbb{P}^1(\mathbb{F}_p)$ in the literature, and so we note that the obvious analogue of Proposition 4.2 in this setting would give a proof.

1.2 Bounds on solutions to a congruence equation

We also give an application of these ideas to bounding the number of solutions to the congruence

$$xy \equiv 1 \pmod{p}$$

with $x \in I$ and $y \in J$ for intervals I and J of the same length $m \leq p/2$. In particular, for large enough p and m , we establish that

$$|\{(x, y) \in I \times J \mid xy \equiv 1 \pmod{p}\}| \leq (1 - \delta)m$$

for some absolute constant $\delta > 0$ (see Theorem 5.1). The proof is simple given the spectral gap estimates in the proof of Theorem 1.2.

Our bound on the number of solutions to the congruence $xy \equiv 1 \pmod{p}$ is weaker than the bounds obtained in [11] and [14] for intervals that are not too large and weaker than the standard ones coming from estimates on incomplete Kloosterman sums for intervals that are not too small. The number of solutions has also been estimated in [19, 15, 33]. In particular, the method used in [33] also relies on SL_2 action. While this result is not new, the proof is straightforward and we hope that our ideas can lead to further applications.

1.3 Related work

The random walk we study may be viewed as a non-linear version of the Chung–Diaconis–Graham process, or the $ax + b$ process, which is the random walk on \mathbb{F}_p (or more generally $\mathbb{Z}/n\mathbb{Z}$ for composite n) defined by

$$X_{n+1} = aX_n + \varepsilon_{n+1},$$

where $a \in \mathbb{F}_p^\times$ is fixed and the ε_i are independent and identically distributed, drawn from some distribution. It was introduced in [13], where the case of $a = 2$ and ε_i distributed uniformly on $\{-1, 0, 1\}$ was studied in detail. They showed that the mixing time was at most order $\log p \log \log p$, and that for almost all odd p , order $\log p$ steps were sufficient, although for infinitely many odd p , order $\log p \log \log p$ steps were necessary. This process was subsequently studied in [25, 26, 29, 27, 10, 18]. In particular, Eberhard and Varjú recently showed in [18] that for almost all p , the walk exhibits cutoff at $c \log p$ for some explicit constant $c \approx 1.01136$.

The Chung–Diaconis–Graham process gives an example of a speedup phenomenon that occurs when applying bijections between steps of a Markov chain. This was studied in [12], where it was shown that for a Markov chain on a state space with n elements, almost all bijections, when applied between steps of a Markov chain, cause the chain to mix in $O(\log n)$ steps. Our work gives another example of this phenomenon, which fits into a broader theme of studying how convergence can be sped up, via either deterministic or random means [24, 5, 2].

In [22], the first author studied a more general non-linear version of the Chung–Diaconis–Graham process, defined by

$$X_{n+1} = f(X_n) + \varepsilon_{n+1},$$

where f is a bijection on \mathbb{F}_p . It was shown that for functions $f(x)$ which were extensions of rational functions, the mixing time was of order at most $p^{1+\varepsilon}$ for any $\varepsilon > 0$, where the implicit constant depends only on ε and the degree of the polynomials appearing in $f(x)$.

In particular, this applies for the function ι that we consider. However, the only known lower bound on the mixing time is of order $\log p$. Our work closes the large gap between the upper and lower bounds in this special case.

1.4 Random walk on Cayley graphs of $SL_2(\mathbb{F}_p)$

Selberg's Theorem [32] shows that if S is a subset of $SL_2(\mathbb{Z})$ such that S generates a subgroup of $SL_2(\mathbb{Z})$ of finite index, then

$$\limsup_{p \rightarrow \infty} \lambda_2(A(\mathcal{G}(SL_2(\mathbb{F}_p), S_p))) < 1. \tag{1.2}$$

Bourgain and Gamburd [6] strengthen this result. (See Theorem 4.7 for more details).

Theorem 1.4 ([6, Theorem 1]). *The property in (1.2) holds if and only if S generates a non-elementary subgroup of $SL_2(\mathbb{Z})$.*

Weigel [35] gives a convenient characterization of when S generates a non-elementary subgroup (see also [30]).

Theorem 1.5 ([35]). *Let $S \subseteq SL_2(\mathbb{Z})$. The subgroup $\langle S \rangle$ is non-elementary if and only if $\langle S_p \rangle = SL_2(\mathbb{F}_p)$ for some prime $p \geq 5$*

We remark that a simple reduction allows us to deduce Theorem 1.2 from the special case where the distribution of ε_n is supported on two points. In this case, we can construct the set S corresponding to the fractional Markov chain and easily verify that $\langle S \rangle$ is non-elementary.

Golumbev and Kamper [21] show that under more stringent assumptions on the set of generators S , the Cayley graphs $\mathcal{G}(SL_2(\mathbb{F}_p), S_p)$ have a stronger property that implies cut-off for the associated projected random walk on $\mathbb{P}^1(\mathbb{F}_p)$. This property also holds with high probability for random generating sets S_p . We also remark that the result of Bourgain-Gamburd has been further developed in various aspects, e.g., [8, 9, 20, 21]. These results also have many applications in number theory, including sum product problems on finite fields [7, 23].

1.5 Outline

In Section 2, we prove Theorem 1.1. In Section 3, we explain the comparison theory for Markov chains used to relate the spectral gaps for the walks on \mathbb{F}_p and $\mathbb{P}^1(\mathbb{F}_p)$. In Section 4, we give a proof of Theorem 1.2. Finally, in Section 5, we give a bound for the number of solutions to $xy \equiv 1 \pmod{p}$ in a square $I_1 \times I_2 \subseteq \mathbb{F}_p^2$.

1.6 Notations

Throughout, we let $\iota : \mathbb{F}_p \rightarrow \mathbb{F}_p$ be the function defined by $\iota(x) = 1/x$ if $x \neq 0$ and $\iota(x) = 0$ if $x = 0$. We view $\mathbb{P}^1(\mathbb{F}_p)$ as a superset of \mathbb{F}_p with one extra element ∞ . We define the function $\bar{\iota} : \mathbb{P}^1(\mathbb{F}_p) \rightarrow \mathbb{P}^1(\mathbb{F}_p)$ by $\bar{\iota}(x) = 1/x$ if $x \neq 0, \infty$, and $\bar{\iota}(0) = \infty$ and $\bar{\iota}(\infty) = 0$.

If P is a symmetric matrix, we let $\lambda_i(P)$ denote the i th largest eigenvalue of P .

2 Lower bound

In this section, we prove Theorem 1.1, which follows from entropy considerations. For a discrete probability measure μ , we let

$$H(\mu) = \sum -\mu(x) \log \mu(x)$$

denote the *entropy* of μ , with the convention that $0 \log 0 = 0$. If X is a random variable, we let $H(X)$ denote the entropy of the law of X . We need two basic properties of entropy.

The first is that it is subadditive, in the sense that if X and Y are independent, then $H(X + Y) \leq H(X) + H(Y)$. The second is that for any function f , $H(f(X)) \leq H(X)$, with equality if f is bijective.

The lower bound follows by noting that the entropy of the random walk at time n is at most $nH(\mu)$, which is too low to be close to uniform if n is too small. This idea is certainly not new, and similar arguments have appeared before, see [18, 1, 4] for example. Our argument follows [1], although we fill in some details and fix a minor error.

The following lemma shows that measures close to uniform in total variation must have large entropy. It was stated without the $\log 2$ term in [1], but this cannot be correct, as can be seen if π is uniform and ν is concentrated at a single point. It can be seen as a special case of Theorem 1 of [3] (in fact, it is established as part of the proof).

Lemma 2.1. *Let X be a finite set of size n , and let π denote the uniform measure on X . Let ν be any probability measure on X , and let $\delta = \|\nu - \pi\|_{TV}$. Then*

$$\begin{aligned} |H(\pi) - H(\nu)| &\leq \delta \log(n - 1) - \delta \log \delta - (1 - \delta) \log(1 - \delta) \\ &\leq \delta \log n + \log(2). \end{aligned}$$

Proof. The first inequality is Equation 11 of [3] and the second inequality is clear. □

Proof of Theorem 1.1. Because ι is a bijection, $H(X) = H(\iota(X))$. Furthermore, if μ_p denotes the distribution given by reducing $\mu \bmod p$, then $H(\mu_p) \leq H(\mu)$. Then by subadditivity of entropy, $H(K^n(x, \cdot)) \leq nH(\mu)$. Thus, Lemma 2.1 gives

$$\|K^n(x, \cdot) - \pi\|_{TV} \geq 1 - \frac{nH(\mu) + \log 2}{\log p}. \quad \square$$

Remark 2.2. We note that the finite entropy assumption cannot be removed entirely. Fix any constant $\delta \in (0, 1)$. Consider the measure on \mathbb{Z}^+ given by

$$\mu(i) \propto \frac{1}{i(\log i)^{2-\delta}}.$$

There is some constant $c = c(\delta) > 0$ and $C > 0$ such that for all p and $i \leq p$,

$$\sum_{k \in \mathbb{Z}^+} \mu(i + kp) \geq \sum_{k \geq p} \frac{C}{kp(\log(kp))^{2-\delta}} \geq \frac{c}{p(\log p)^{1-\delta}}.$$

Then on \mathbb{F}_p , the walk defined by (1.1) satisfies

$$\mathbb{P}(X_{n+1} = x | X_n = y) \geq \frac{c}{p(\log p)^{1-\delta}}$$

for all x and y , and so we can couple two copies of the walk, X_n and X'_n , so that independent of the initial states,

$$\mathbb{P}(X_n \neq X'_n) \leq \left(1 - \frac{c}{(\log p)^{1-\delta}}\right)^n.$$

Since the distance to stationarity can be controlled by this probability (see [28, Corollary 5.5]), this implies that

$$\|K^n(x, \cdot) - \pi\|_{TV} \leq \left(1 - \frac{c}{(\log p)^{1-\delta}}\right)^n,$$

and so $t_{mix}(\varepsilon) = O_\varepsilon((\log p)^{1-\delta})$, which is much smaller than $\log p$.

For positive integers k , define $\log^{(k)}$ iteratively by $\log^{(1)}(x) = \log x$ and $\log^{(k)}(x) = \log(\log^{(k-1)}(x))$ for $k \geq 2$. The measure

$$\mu(i) \propto \frac{1}{i(\log i)(\log^{(2)} i) \cdots (\log^{(k-1)} i)(\log^{(k)} i)^2}$$

gives a probability distribution on \mathbb{Z} with infinite entropy for which the mixing time $t_{mix}(\varepsilon) = O_\varepsilon(\log^{(k)} p)$.

3 Comparison theory

Comparison theory for Markov chains was introduced by Diaconis and Saloff-Coste in [17, 16]. This theory allows the spectral gaps (and also log-Sobolev constants) of different Markov chains on the same state space to be compared. Unfortunately, since we wish to compare a Markov chain on \mathbb{F}_p with one on $\mathbb{P}^1(\mathbb{F}_p)$, this theory does not immediately apply. Smith extended these ideas in [34] to the case of random walks on state spaces $X_0 \subseteq X$.

We first recall the relationship between the Dirichlet form and the spectral gap, and then explain the comparison theory developed by Smith for Markov chains on different state spaces. This is used to compare the spectral gaps of the random walks on $\mathbb{P}^1(\mathbb{F}_p)$ and \mathbb{F}_p . For a more thorough treatment and proofs, we refer the reader to [34].

3.1 The Dirichlet form

Recall the following standard notions. Let P be a reversible Markov chain on a finite space X with stationary distribution π . Define the function

$$V_\pi(f) = \frac{1}{2} \sum_{x,y \in X} |f(x) - f(y)|^2 \pi(x)\pi(y)$$

and the Dirichlet form

$$\mathcal{E}_P(f, f) = \frac{1}{2} \sum_{x,y \in X} |f(x) - f(y)|^2 P(x, y)\pi(x)$$

for $f : X \rightarrow \mathbb{R}$. Then the spectral gap can be computed by

$$1 - \lambda_2(P) = \inf_{\substack{f : X \rightarrow \mathbb{R} \\ f \text{ not constant}}} \frac{\mathcal{E}_P(f, f)}{V_\pi(f)}. \tag{3.1}$$

3.2 Comparison theory for Markov chains on different state spaces

Let $X_0 \subseteq X$, and let P_0 be a Markov chain on X_0 and P be a Markov chain on X , with stationary distributions π_0 and π respectively. Assume that both are ergodic and reversible. The following results of Smith compare V_{π_0} and \mathcal{E}_{P_0} with V_π and \mathcal{E}_P .

Lemma 3.1 ([34, Lemma 2]). *Let $f_0 : X_0 \rightarrow \mathbb{R}$ and let $f : X \rightarrow \mathbb{R}$ be any extension of f_0 . Then*

$$V_{\pi_0}(f_0) \leq CV_\pi(f),$$

where

$$C = \sup_{x \in X_0} \frac{\pi_0(x)}{\pi(x)}.$$

The analogous comparison result for the Dirichlet form requires further setup. For each $x \in X$, fix a probability measure Q_x on X_0 such that $Q_x = \delta_x$ if $x \in X_0$. For any $f_0 : X_0 \rightarrow \mathbb{R}$, this defines an extension $f : X \rightarrow \mathbb{R}$ by

$$f(x) = \sum_{y \in X_0} Q_x(y)f(y).$$

Next, choose a coupling of Q_x and Q_y for all $x, y \in X$ for which $P(x, y) > 0$ (note this is a symmetric relation). Denote these couplings by $Q_{x,y}$.

A path in X_0 from x to y is a sequence of $x_i \in X_0$ for $i = 0, \dots, k$ for which $x_0 = x$, $x_k = y$ and $P_0(x_i, x_{i+1}) > 0$ for all i . For a path γ from x to y , call x the *initial vertex* and y the *final vertex*, and denote them by $i(\gamma)$ and $o(\gamma)$ respectively. Let $|\gamma|$ denote the length of the path.

A flow on X_0 is a function F on the set of paths in X_0 whose restriction to paths from x to y gives a probability measure for all x, y . We require a choice of flow on X_0 . Actually it suffices to define the flow only for paths from a to b such that there exists $x, y \in X$ with $Q_{x,y}(a, b) > 0$. For our purposes, it will suffice to choose the flow such that when restricted to paths from a to b , it concentrates on a single path.

The following theorem compares \mathcal{E}_{P_0} with \mathcal{E}_P in terms of the chosen data.

Theorem 3.2 ([34, Theorem 4]). *Consider the setup defined above. Then for $f_0 : X_0 \rightarrow \mathbb{R}$ and f the extension to X defined by the measures Q_x , we have*

$$\mathcal{E}_P(f, f) \leq \mathcal{A} \mathcal{E}_{P_0}(f_0, f_0),$$

where

$$\mathcal{A} = \sup_{P_0(x,y)>0} \frac{1}{P_0(x, y)\pi_0(x)} (S_1 + S_2 + S_3)$$

with

$$\begin{aligned} S_1 &:= \sum_{\gamma \ni (x,y)} |\gamma| F(\gamma) P(i(\gamma), o(\gamma)) \pi(i(\gamma)) \\ S_2 &:= 2 \sum_{\gamma \ni (x,y)} |\gamma| F(\gamma) \sum_{b \notin X_0} Q_b(o(\gamma)) P(i(\gamma), b) \pi(i(\gamma)) \\ S_3 &:= \sum_{\gamma \ni (x,y)} |\gamma| F(\gamma) \sum_{\substack{a, b \notin X_0 \\ P(x,y)>0}} Q_{a,b}(i(\gamma), o(\gamma)) P(a, b) \pi(a). \end{aligned}$$

Note that by (3.1), an immediate consequence of these comparison results is that

$$1 - \lambda_2(P_0) \geq \frac{1}{\mathcal{CA}} (1 - \lambda_2(P)).$$

4 Upper bound

In this section, we prove Theorem 1.2. The proof of the upper bound consists of three main steps. First, we control the mixing time by the spectral gap of a symmetrized random walk.

Proposition 4.1. *Let π denote the uniform measure on \mathbb{F}_p . Then for all $x \in \mathbb{F}_p$,*

$$\|K^k(x, \cdot) - \pi\|_{TV} \leq \frac{\sqrt{p}}{2} \lambda_2^{(k-2)/4}.$$

Next, we relate the spectral gap of the symmetrized walk on \mathbb{F}_p with a walk on $\mathbb{P}^1(\mathbb{F}_p)$.

Proposition 4.2. *Let L_0 and L denote the transition matrices for the random walk on \mathbb{F}_p and $\mathbb{P}^1(\mathbb{F}_p)$ respectively. Then there exists an absolute constant $c > 0$ such that*

$$1 - \lambda_2(L_0) \geq c(1 - \lambda_2(L)).$$

Finally, we show that the spectral gap of the walk on $\mathbb{P}^1(\mathbb{F}_p)$ is of constant order.

Proposition 4.3. *Let L denote the transition matrix for the random walk on $\mathbb{P}^1(\mathbb{F}_p)$ which moves from $x \in \mathbb{F}_p$ to one of $x + b, x - b, \bar{\iota}(\bar{\iota}(x + a_1) + b) - a_1$ and $\bar{\iota}(\bar{\iota}(x + a_1) - b) - a_1$ uniformly at random. Then*

$$\lambda_2(L) \leq 1 - c$$

for some constant $c > 0$ independent of p .

Proof of Theorem 1.2. Let P denote the Markov matrix encoding the step $X \mapsto X + \varepsilon$ and let Π denote the Markov matrix encoding $X \mapsto \iota(X)$. Note that P^T encodes the step $X \mapsto X - \varepsilon$. By Proposition 4.1, it suffices to show that

$$\lambda_2(P^T \Pi P^T P \Pi P) \leq 1 - c$$

for some constant $c > 0$, independent of p .

Let $a_1, a_2 \in \mathbb{Z}$ be two distinct elements in the support of μ and write $b = a_1 - a_2$. Then in the Markov chain given by the transition matrix $P^T \Pi P^T P \Pi P$, there is $u > 0$ depending only on μ such that we transition from x to one of $x + b, x - b, \iota(\iota(x + a_1) + b) - a_1$ and $\iota(\iota(x + a_1) - b) - a_1$ with probability at least u . For example, note that $x + b = \iota(\iota(x + a_1) + a_1 - a_1) - a_2$ and so there is a positive probability of moving from x to $x + b$.

Then we can write $P^T \Pi P^T P \Pi P = uL_0 + (1 - u)L'$ where L_0 is the transition matrix for the random walk going from x to one of $x + b, x - b, \iota(\iota(x + a_1) + b) - a_1$ and $\iota(\iota(x + a_1) - b) - a_1$ with equal probability, and L' a symmetric stochastic matrix. We have

$$\lambda_2(P^T \Pi P^T P \Pi P) \leq u\lambda_2(L_0) + 1 - u,$$

and so it suffices to show that $\lambda_2(L_0) \leq 1 - c$ for some constant $c > 0$.

Let $\bar{\iota} : \mathbb{P}^1(\mathbb{F}_p) \rightarrow \mathbb{P}^1(\mathbb{F}_p)$ be the function defined by $\bar{\iota}(x) = 1/x$ if $x \neq 0, \infty$, and $\bar{\iota}(0) = \infty$, and $\bar{\iota}(\infty) = 0$. Let L denote the transition matrix for the random walk on $\mathbb{P}^1(\mathbb{F}_p)$ going from x to one of $x + b, x - b, \bar{\iota}(\bar{\iota}(x + a_1) + b) - a_1$ and $\bar{\iota}(\bar{\iota}(x + a_1) - b) - a_1$ with equal probability. By Proposition 4.2, we can instead show

$$\lambda_2(L) \leq 1 - c$$

for the walk defined by L instead of the walk defined by L_0 , and this is exactly given by Proposition 4.3. □

4.1 Reduction to the symmetric walk

Recall that we are interested in studying the walk $X_{n+1} = \iota(X_n) + \varepsilon_{n+1}$ on \mathbb{F}_p , where ε_n are independent and identically distributed according to μ . This random walk is non-reversible, so the first step is to relate it to a suitable symmetrization.

Let Π denote the transition matrix for the (deterministic) walk $X \mapsto \iota(X)$ on \mathbb{F}_p and let P denote the walk $X \mapsto X + \varepsilon$. Note that P^T is also a Markov matrix and encodes the walk $X \mapsto X - \varepsilon$. Let λ_2 be the second largest eigenvalue of $P^T \Pi P^T P \Pi P$. Note that all eigenvalues of $P^T \Pi P^T P \Pi P$ are real and non-negative.

The following result from [12] relates the original walk $K = P\Pi$ to the symmetrized walk $P^T \Pi P^T P \Pi P$.

Lemma 4.4 ([12, Corollary 5.4]). *Let $k \geq 2$ and let $x \in \mathbb{R}^p$ such that $\sum x_i = 0$. Then*

$$\|(K^T)^k x\|_2 \leq \lambda_2^{(k-2)/4} \|x\|_2.$$

This result means that if $P^T \Pi P^T P \Pi P$ can be shown to have a spectral gap of constant order, then the walk K mixes in order $\log p$ steps. Proposition 4.1 is an easy consequence of Lemma 4.4 (see Equation 5.5 in [12]).

4.2 Comparison of the random walks on $\mathbb{P}^1(\mathbb{F}_p)$ and \mathbb{F}_p

We now explain how to apply the comparison results to bound the spectral gap of the walk on \mathbb{F}_p in terms of the spectral gap of the walk on $\mathbb{P}^1(\mathbb{F}_p)$.

Recall from the proof of Theorem 1.2 that we denote by L_0 the transition matrix of the random walk on \mathbb{F}_p which moves from $x \in \mathbb{F}_p$ to one of $x+b, x-b, \iota(\iota(x+a_1)+b)-a_1$ and $\iota(\iota(x+a_1)-b)-a_1$ with equal probability. We denote by L the transition matrix of the walk on $\mathbb{P}^1(\mathbb{F}_p)$ which moves from $x \in \mathbb{F}_p$ to one of $x+b, x-b, \bar{\iota}(\bar{\iota}(x+a_1)+b)-a_1$ and $\bar{\iota}(\bar{\iota}(x+a_1)-b)-a_1$. Here, recall that $\bar{\iota}(x) = 1/x$ if $x \neq 0, \infty$, and $\bar{\iota}(0) = \infty$, and $\bar{\iota}(\infty) = 0$. The following lemma is useful in constructing the data required in Theorem 3.2.

Lemma 4.5. *For all $x, y \in \mathbb{F}_p$ except $x = -a_1$ and $y = -a_1$, if $L(x, y) > 0$, then $L_0(x, y) > 0$.*

Proof. First, note that if $x \neq -a_1, b^{-1} - a_1, -b^{-1} - a_1$, then the statement is clear since in this case the functions ι and $\bar{\iota}$ are identical for the transitions involved. It can be checked that if $x = b^{-1} - a_1$ or $x = -b^{-1} - a_1$, and $y \neq \infty$ and $L(x, y) > 0$, then $L_0(x, y) > 0$. Similarly, we can check that if $x = -a_1$ and $L(x, y) > 0$ then $L_0(x, y) > 0$ unless $y = -a_1$. □

We are now ready to prove Proposition 4.2.

Proof of Proposition 4.2. We begin by defining the data needed to apply Theorem 3.2. Take $X_0 = \mathbb{F}_p$ and $X = \mathbb{P}^1(\mathbb{F}_p)$. Define for $x \in \mathbb{F}_p$, $Q_x = \delta_x$ and Q_∞ is uniform on the set of $y \in \mathbb{F}_p$ such that $L(\infty, y) > 0$. Take all couplings to be independent, so $Q_{x,y}(x_0, y_0) = Q_x(x_0)Q_y(y_0)$. Then note that the condition that $Q_{x,y}(x_0, y_0) > 0$ for some $x, y \in X$ is exactly that either $L(x_0, y_0) > 0$, or $L(x_0, \infty) > 0$ and $L(y_0, \infty) > 0$.

By Lemma 4.5, if $L(x_0, y_0) > 0$, then $L_0(x_0, y_0) > 0$ except when $x_0 = -a_1$ and $y_0 = -a_1$. Thus, if $L(x_0, y_0) > 0$, we pick our flow F to be concentrated on the path consisting of the single edge (x_0, y_0) , except when $x_0 = y_0 = -a_1$ where we take the path of length 2 given by $-a_1 \mapsto b^{-1} - a_1 \mapsto -a_1$.

If $L(x_0, \infty) > 0$ and $L(y_0, \infty) > 0$, then there must be a path of length 2 connecting x_0 to y_0 . This is because $x_0, y_0 \in \{b^{-1} - a_1, -b^{-1} - a_1\}$ and so we can take our flow to be concentrated on the path $x_0 \mapsto -a_1 \mapsto y_0$. Thus, our flow is concentrated entirely on paths of lengths 1 and 2.

We can now compute the constants C and \mathcal{A} . It's easy to see that $C \leq 2$. To bound \mathcal{A} , note that for each (x, y) for which $L_0(x, y) > 0$, there is at most one path γ of length 1 on which $F(\gamma) \neq 0$ containing (x, y) , namely $\gamma = (x, y)$, and there is a bounded number of paths γ of length 2 containing (x, y) for which $F(\gamma) > 0$. Thus, for any (x, y) for which $L_0(x, y) > 0$, there are a bounded number of terms in the summation in the definition of \mathcal{A} , each of bounded size. Hence, \mathcal{A} is bounded above by a constant. This implies that

$$1 - \lambda_2(L_0) \geq c(1 - \lambda_2(L))$$

for some absolute constant $c > 0$. □

Remark 4.6. For a reversible Markov chain P with stationary distribution π , let

$$\phi(P) = \min_{\pi(S) \leq 1/2} \frac{\sum_{x \in S, y \in S^c} \pi(x)P(x, y)}{\pi(S)}$$

denote the *bottleneck ratio* (sometimes also called the *Cheeger constant* or *conductance*). We note that a constant order spectral gap for L_0 could also be derived using Cheeger's inequality (see [28, Theorem 13.10] for example), which states that

$$\frac{\phi(P)^2}{2} \leq 1 - \lambda_2(P) \leq 2\phi(P).$$

The bottleneck ratios of L_0 and L can be easily compared, which would give a uniform lower bound for $1 - \lambda_2(L_0)$ using the one for $1 - \lambda_2(L)$. We prefer to use comparison theory because the argument can be more easily generalized. In particular, it is more robust and would give a sharper result if L_0 did not have a constant order spectral gap.

4.3 Spectral gap for random walk on $\mathbb{P}^1(\mathbb{F}_p)$

In this section, we prove that the random walk on $\mathbb{P}^1(\mathbb{F}_p)$ has a constant order spectral gap.

Recall that $SL_2(\mathbb{F}_p)$ has a transitive action on $\mathbb{P}^1(\mathbb{F}_p)$, viewed as lines in \mathbb{F}_p^2 . This is through linear fractional transformations, and may be formally defined by

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix} \cdot x = \begin{cases} \frac{ax+b}{cx+d} & \text{if } cx + d \neq 0 \text{ and } x \in \mathbb{F}_p \\ \infty & \text{if } cx + d = 0 \text{ and } x \in \mathbb{F}_p \\ \frac{a}{c} & \text{if } x = \infty \text{ and } c \neq 0 \\ \infty & \text{if } x = \infty \text{ and } c = 0. \end{cases} \tag{4.1}$$

The random walk on $\mathbb{P}^1(\mathbb{F}_p)$ given by moving from x to one of $x + b$, $x - b$, $\bar{i}(\bar{i}(x + a_1) + b) - a_1$ and $\bar{i}(\bar{i}(x + a_1) - b) - a_1$ uniformly at random can be viewed as the quotient of a random walk on $SL_2(\mathbb{F}_p)$. The following result of Bourgain and Gamburd gives a constant order spectral gap for random walks on $SL_2(\mathbb{F}_p)$.

Theorem 4.7 ([6, Theorem 1]). *Let $S \subseteq SL_2(\mathbb{Z})$ be a symmetric set which generates a non-elementary subgroup of $SL_2(\mathbb{Z})$. Let S_p denote the set of generators mod p for a prime p . Let P denote the transition matrix for the random walk on $SL_2(\mathbb{F}_p)$ defined by*

$$X_{n+1} = X_n \varepsilon_{n+1},$$

where the ε_i are independent and uniformly distributed on S_p . Then for all primes p large enough,

$$\lambda_2(P) \leq 1 - c$$

for some constant $c > 0$ independent of p .

To show that this applies to the random walk we wish to study, we need the following lemma.

Lemma 4.8. *Let*

$$S = \left\{ \begin{pmatrix} 1 & b \\ 0 & 1 \end{pmatrix}, \begin{pmatrix} 1 - a_1 b & -a_1^2 b \\ b & a_1 b + 1 \end{pmatrix}, \begin{pmatrix} 1 & -b \\ 0 & 1 \end{pmatrix}, \begin{pmatrix} 1 + a_1 b & a_1^2 b \\ -b & 1 - a_1 b \end{pmatrix} \right\}.$$

Then S generates a non-elementary subgroup of $SL_2(\mathbb{Z})$.

Proof. By Theorem 2.5 of [30], non-elementary subgroups are the same as Zariski-dense subgroups in $SL_2(\mathbb{Z})$. By Theorem 1.5, it suffices to show that after reducing mod p for some $p \geq 5$, S_p generates $SL_2(\mathbb{F}_p)$.

If $b \neq 0 \pmod{p}$, then the matrices $\begin{pmatrix} 1 & b \\ 0 & 1 \end{pmatrix}$ and $\begin{pmatrix} 1 & -b \\ 0 & 1 \end{pmatrix}$ generate all matrices of the form $\begin{pmatrix} 1 & t \\ 0 & 1 \end{pmatrix}$ for $t \in \mathbb{F}_p$.

Using $\begin{pmatrix} 1 & t \\ 0 & 1 \end{pmatrix}$ and $\begin{pmatrix} 1 - a_1 b & -a_1^2 b \\ b & a_1 b + 1 \end{pmatrix}$, we can generate all matrices of the form

$$\begin{pmatrix} 1 & t' \\ 0 & 1 \end{pmatrix} \cdot \begin{pmatrix} 1 - a_1 b & -a_1^2 b \\ b & a_1 b + 1 \end{pmatrix} \cdot \begin{pmatrix} 1 & t \\ 0 & 1 \end{pmatrix}.$$

In particular, we can generate all matrices in the subset X_b of $SL_2(\mathbb{F}_p)$ of matrices whose lower left corner equal to b .

Similarly, S_p generates matrices in the subset X_{-b} of $SL_2(\mathbb{F}_p)$ of matrices with lower left corner equal to $-b$, using $\begin{pmatrix} 1 + a_1b & a_1^2b \\ -b & 1 - a_1b \end{pmatrix}$ instead of $\begin{pmatrix} 1 - a_1b & -a_1^2b \\ b & a_1b + 1 \end{pmatrix}$.

Finally, we can easily check that X_b and X_{-b} generate $SL_2(\mathbb{F}_p)$. Thus, S_p generates $SL_2(\mathbb{F}_p)$ for all $p > |b|$, and hence, S generates a non-elementary subgroup of $SL_2(\mathbb{Z})$. \square

Since the random walk on $\mathbb{P}^1(\mathbb{F}_p)$ is a quotient of a random walk on $SL_2(\mathbb{F}_p)$, we can obtain the desired bound on the spectral gap, and we prove Proposition 4.3.

Proof of Proposition 4.3. Note that the formulas given in 4.1 actually define a $GL_2(\mathbb{F}_p)$ action. It's easy to see that ι and addition by $a \in \mathbb{F}_p$ are both linear fractional transformations, represented by the matrices

$$\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \begin{pmatrix} 1 & a \\ 0 & 1 \end{pmatrix},$$

respectively, in $GL_2(\mathbb{F}_p)$. Then we may write the function $x \mapsto x + b$ as a product of these matrices, and similarly for $x - b$, $\bar{\iota}(\bar{\iota}(x + a_1) + b) - a_1$ and $\bar{\iota}(\bar{\iota}(x + a_1) - b) - a_1$. This allows us to view the walk defined by L as the quotient of the walk on $SL_2(\mathbb{F}_p)$ generated by the set

$$S = \left\{ \begin{pmatrix} 1 & b \\ 0 & 1 \end{pmatrix}, \begin{pmatrix} 1 - a_1b & -a_1^2b \\ b & a_1b + 1 \end{pmatrix}, \begin{pmatrix} 1 & -b \\ 0 & 1 \end{pmatrix}, \begin{pmatrix} 1 + a_1b & a_1^2b \\ -b & 1 - a_1b \end{pmatrix} \right\},$$

with transition matrix \tilde{L} . This means that the spectrum of L is contained in the spectrum of \tilde{L} , and in particular $\lambda_2(L) \leq \lambda_2(\tilde{L})$. But since S generates a non-elementary subgroup of $SL_2(\mathbb{Z})$ by Lemma 4.8, Theorem 4.7 implies that

$$\lambda_2(\tilde{L}) \leq 1 - c$$

for some constant $c > 0$ independent of p , giving us the desired bound. \square

5 Points on modular hyperbolas

In this section, we prove the following theorem stating that for all subsets $A \subseteq \mathbb{F}_p$ of size at most $p/2$, the sets A and $\iota(A)$ cannot both be close to intervals. As a corollary, we obtain bounds on the number of solutions to the congruence $xy = 1 \pmod{p}$ lying in $I \times J \subseteq \mathbb{F}_p^2$ for I and J intervals of the same length.

Theorem 5.1. *Let I and J be two intervals in \mathbb{F}_p , each of length $m \leq p/2$, and let $A \subseteq \mathbb{F}_p$ with $A \subseteq I$ and $\iota(A) \subseteq J$. There exists an absolute constant $\delta > 0$ such that for all p and m sufficiently large, $|A| \leq (1 - \delta)m$.*

Proof. Throughout the proof, we use the notation $[-k, k]$ to denote the set of integers $\{-k, -k + 1, \dots, k - 1, k\}$. Let $Q = P^T \Pi P^T P \Pi P$, where P denotes the transition matrix for the random walk on \mathbb{F}_p generated by the uniform measure on $[-1, 1]$ and Π encodes the bijection ι . Then Q is the transition matrix of a reversible Markov chain on \mathbb{F}_p , and from the proof of Theorem 1.2 it has a constant order spectral gap, i.e. for large enough p , $1 - \lambda_2(Q) \geq \gamma$ for some absolute constant $\gamma > 0$.

Assume that there exists intervals I, J such that $A \subseteq I$, $\iota(A) \subseteq J$, and $|A| \geq (1 - \delta)m$. We will show using Cheeger's inequality that this implies the spectral gap cannot be of constant order, which gives a contradiction.

Observe that $Q(x, y) \geq c > 0$ for some constant c independent of p . Since the stationary distribution is uniform, the bottleneck ratio of the chain X_n is given by

$$\Phi = \min_{S \subseteq \mathbb{F}_p, |S| \leq p/2} \frac{\sum_{x \in S, y \notin S} Q(x, y)}{|S|}.$$

By Cheeger's inequality, we have Φ is bounded below by a positive constant γ' depending only on γ . Thus,

$$\frac{\sum_{x \in A, y \notin A} Q(x, y)}{|A|} \geq \gamma'.$$

For $x \in A$, we consider $y \in \mathbb{F}_p$ for which $Q(x, y) > 0$. Then

$$y \in \iota(\iota(A + [-1, 1]) + [-2, 2]) + [-1, 1].$$

Note that $\iota(A + [-1, 1]) \subseteq \iota(I + [-1, 1])$, and

$$\begin{aligned} \iota(I + [-1, 1]) &= \iota(A \cup ((I + \{-1, +1\}) \setminus A)) \\ &= \iota(A) \cup \iota((I + \{-1, +1\}) \setminus A). \end{aligned}$$

Let $S_1 = \iota((I + \{-1, +1\}) \setminus A)$. Then $\iota(A + [-1, 1]) \subseteq \iota(A) \cup S_1$, and $|S_1| \leq 2 + \delta m$.

Similarly, since $\iota(A) \cup S_1 \subseteq J \cup S_1$, we have

$$\begin{aligned} \iota((\iota(A) \cup S_1) + [-2, 2]) &\subseteq \iota(J + [-2, 2]) \cup \iota(S_1 + [-2, 2]) \\ &\subseteq A \cup \iota((J + [-2, 2]) \setminus \iota(A)) \cup \iota(S_1 + [-2, 2]). \end{aligned}$$

Here, $|\iota((J + [-2, 2]) \setminus \iota(A))| \leq 4 + \delta m$, and $|\iota(S_1 + [-2, 2])| \leq 5|S_1|$. Letting

$$S_2 = \iota((J + [-2, 2]) \setminus \iota(A)) \cup \iota(S_1 + [-2, 2]),$$

we have

$$\iota(\iota(A + [-1, 1]) + [-2, 2]) \subseteq A \cup S_2,$$

where $|S_2| \leq 14 + 6\delta m$.

Finally, we have $y \in (A \cup S_2) + [-1, 1]$, and

$$\begin{aligned} (A \cup S_2) + [-1, 1] &\subseteq (I \cup S_2) + [-1, 1] \\ &\subseteq (I + [-1, 1]) \cup (S_2 + [-1, 1]) \\ &\subseteq A \cup ((I + [-1, 1]) \setminus A) \cup (S_2 + [-1, 1]). \end{aligned}$$

Let $S_3 = ((I + [-1, 1]) \setminus A) \cup (S_2 + [-1, 1])$, and note that $|S_3| \leq 44 + 19\delta m$.

Thus, if $x \in A$, $y \notin A$, and $Q(x, y) > 0$, then we must have $y \in S_3$. Hence,

$$\frac{\sum_{x \in A, y \notin A} Q(x, y)}{|A|} \leq \frac{\sum_{x \in A, y \in S_3} Q(x, y)}{|A|} \leq \frac{44 + 19\delta m}{(1 - \delta)m} \leq \frac{20\delta}{1 - \delta},$$

where we assumed that m is sufficiently large (in δ), and used that

$$\sum_{x \in A, y \in S_3} Q(x, y) \leq \sum_{x \in \mathbb{F}_p, y \in S_3} Q(y, x) = |S_3|.$$

For δ sufficiently small, $20\delta/(1 - \delta) < \gamma'$, and we have a contradiction. □

The following corollary follows easily from Theorem 5.1 by taking $A = \{x \in I \mid \iota(x) \in J\}$, which away from 0 counts the solutions of interest.

Corollary 5.2. *Let I and J be two intervals in \mathbb{F}_p of the same length $m \leq p/2$. There exists an absolute constant $\delta > 0$ such that for sufficiently large p and m ,*

$$|\{(x, y) \in I \times J \mid xy \equiv 1 \pmod{p}\}| \leq (1 - \delta)m.$$

References

- [1] David Aldous, *Random walks on finite groups and rapidly mixing Markov chains*, Seminar on probability, XVII, Lecture Notes in Math., vol. 986, Springer, Berlin, 1983, pp. 243–297. MR770418
- [2] Noga Alon, Itai Benjamini, Eyal Lubetzky, and Sasha Sodin, *Non-backtracking random walks mix faster*, Commun. Contemp. Math. **9** (2007), no. 4, 585–603. MR2348845
- [3] Koenraad M. R. Audenaert, *A sharp continuity estimate for the von Neumann entropy*, J. Phys. A **40** (2007), no. 28, 8127–8136. MR2344161
- [4] Charles Bordenave, Pietro Caputo, and Justin Salez, *Cutoff at the “entropic time” for sparse Markov chains*, Probab. Theory Related Fields **173** (2019), no. 1-2, 261–292. MR3916108
- [5] Charles Bordenave, Yanqi Qiu, and Yiwei Zhang, *Spectral gap of sparse bistochastic matrices with exchangeable rows*, Ann. Inst. Henri Poincaré Probab. Stat. **56** (2020), no. 4, 2971–2995. MR4164862
- [6] Jean Bourgain and Alex Gamburd, *Uniform expansion bounds for Cayley graphs of $SL_2(\mathbb{F}_p)$* , Ann. of Math. (2) **167** (2008), no. 2, 625–642. MR2415383
- [7] Jean Bourgain, Alex Gamburd, and Peter Sarnak, *Affine linear sieve, expanders, and sum-product*, Invent. Math. **179** (2010), no. 3, 559–644. MR2587341
- [8] Jean Bourgain and Péter P. Varjú, *Expansion in $SL_d(\mathbb{Z}/q\mathbb{Z})$, q arbitrary*, Invent. Math. **188** (2012), no. 1, 151–173. MR2897695
- [9] Emmanuel Breuillard and Alex Gamburd, *Strong uniform expansion in $SL(2, p)$* , Geom. Funct. Anal. **20** (2010), no. 5, 1201–1209. MR2746951
- [10] Emmanuel Breuillard and Péter P. Varjú, *Cut-off phenomenon for the $ax+b$ Markov chain over a finite field*, 2019.
- [11] Tsz Ho Chan and Igor E. Shparlinski, *On the concentration of points on modular hyperbolas and exponential curves*, Acta Arith. **142** (2010), no. 1, 59–66. MR2601049
- [12] Sourav Chatterjee and Persi Diaconis, *Speeding up Markov chains with deterministic jumps*, Probab. Theory Related Fields **178** (2020), no. 3-4, 1193–1214. MR4168397
- [13] F. R. K. Chung, Persi Diaconis, and R. L. Graham, *Random walks arising in random number generation*, Ann. Probab. **15** (1987), no. 3, 1148–1165. MR893921
- [14] Javier Cilleruelo and Moubariz Z. Garaev, *Concentration of points on two and three dimensional modular hyperbolas and applications*, Geom. Funct. Anal. **21** (2011), no. 4, 892–904. MR2827013
- [15] Javier Cilleruelo and Ana Zumalacárregui, *Saving the logarithmic factor in the error term estimates of some congruence problems*, Math. Z. **286** (2017), no. 1-2, 545–558. MR3648508
- [16] Persi Diaconis and Laurent Saloff-Coste, *Comparison techniques for random walk on finite groups*, Ann. Probab. **21** (1993), no. 4, 2131–2156. MR1245303
- [17] Persi Diaconis and Laurent Saloff-Coste, *Comparison theorems for reversible Markov chains*, Ann. Appl. Probab. **3** (1993), no. 3, 696–730. MR1233621
- [18] Sean Eberhard and Péter P. Varjú, *Mixing time of the Chung-Diaconis-Graham random process*, Probab. Theory Related Fields **179** (2021), no. 1-2, 317–344. MR4221660
- [19] M. Z. Garaev, *On the logarithmic factor in error term estimates in certain additive congruence problems*, Acta Arith. **124** (2006), no. 1, 27–39. MR2262138
- [20] A. Salehi Golsefidy and Péter P. Varjú, *Expansion in perfect groups*, Geom. Funct. Anal. **22** (2012), no. 6, 1832–1891. MR3000503
- [21] Konstantin Golubev and Amitay Kamber, *Cutoff on graphs and the Sarnak-Xue density of eigenvalues*, 2019.
- [22] Jimmy He, *Markov chains on finite fields with deterministic jumps*, Electron. J. Probab. **27** (2022), Paper No. 28, 17. MR4385831
- [23] H. A. Helfgott, *Growth and generation in $SL_2(\mathbb{Z}/p\mathbb{Z})$* , Ann. of Math. (2) **167** (2008), no. 2, 601–623. MR2415382
- [24] Jonathan Hermon, Allan Sly, and Perla Sousi, *Universality of cutoff for graphs with an added random matching*, Ann. Probab. **50** (2022), no. 1, 203–240. MR4385126

- [25] Martin Hildebrand, *On the Chung-Diaconis-Graham random process*, Electron. Comm. Probab. **11** (2006), 347–356. MR2274529
- [26] Martin Hildebrand, *A lower bound for the Chung-Diaconis-Graham random process*, Proc. Amer. Math. Soc. **137** (2009), no. 4, 1479–1487. MR2465674
- [27] Martin Hildebrand, *On a lower bound for the Chung-Diaconis-Graham random process*, Statist. Probab. Lett. **152** (2019), 121–125. MR3953053
- [28] David A. Levin and Yuval Peres, *Markov chains and mixing times*, American Mathematical Society, Providence, RI, 2017, Second edition of [MR2466937], With contributions by Elizabeth L. Wilmer, With a chapter on “Coupling from the past” by James G. Propp and David B. Wilson. MR3726904
- [29] Richard Neville, III, *On lower bounds of the Chung-Diaconis-Graham random process*, ProQuest LLC, Ann Arbor, MI, 2011, Thesis (Ph.D.)–State University of New York at Albany. MR2873460
- [30] Igor Rivin, *Zariski density and genericity*, Int. Math. Res. Not. IMRN (2010), no. 19, 3649–3657. MR2725508
- [31] Peter Sarnak, *Some applications of modular forms*, Cambridge Tracts in Mathematics, vol. 99, Cambridge University Press, Cambridge, 1990. MR1102679
- [32] Atle Selberg, *On the estimation of Fourier coefficients of modular forms*, Proc. Sympos. Pure Math., Vol. VIII, Amer. Math. Soc., Providence, R.I., 1965, pp. 1–15. MR0182610
- [33] I. D. Shkredov, *Modular hyperbolas and bilinear forms of Kloosterman sums*, J. Number Theory **220** (2021), 182–211. MR4177541
- [34] Aaron Smith, *Comparison theory for Markov chains on different state spaces and application to random walk on derangements*, J. Theoret. Probab. **28** (2015), no. 4, 1406–1430. MR3422936
- [35] Thomas Weigel, *On the profinite completion of arithmetic groups of split type*, Lois d’algèbres et variétés algébriques (Colmar, 1991), Travaux en Cours, vol. 50, Hermann, Paris, 1996, pp. 79–101. MR1600994
- [36] Avi Wigderson, *Mathematics and computation*, Princeton University Press, Princeton, NJ, 2019, A theory revolutionizing technology and science. MR4205964

Acknowledgments. The authors thank Sourav Chatterjee, Persi Diaconis, Jacob Fox, Sean Eberhard, Ilya Shkredov, Kannan Soundararajan, Péter Varjú, Thuy Duong Vuong, and Yuval Wigderson for their help and comments on earlier drafts. We are greatly indebted to Péter Varjú for pointing out the many useful references to the study of spectral gap of Cayley graphs of $SL_2(\mathbb{F}_p)$, in particular, [6]. We thank the anonymous referees for helpful comments and corrections.

Electronic Journal of Probability

Electronic Communications in Probability

Advantages of publishing in EJP-ECP

- Very high standards
- Free for authors, free for readers
- Quick publication (no backlog)
- Secure publication (LOCKSS¹)
- Easy interface (EJMS²)

Economical model of EJP-ECP

- Non profit, sponsored by IMS³, BS⁴, ProjectEuclid⁵
- Purely electronic

Help keep the journal free and vigorous

- Donate to the IMS open access fund⁶ (click here to donate!)
- Submit your best articles to EJP-ECP
- Choose EJP-ECP over for-profit journals

¹LOCKSS: Lots of Copies Keep Stuff Safe <http://www.lockss.org/>

²EJMS: Electronic Journal Management System: <https://vtex.lt/services/ejms-peer-review/>

³IMS: Institute of Mathematical Statistics <http://www.imstat.org/>

⁴BS: Bernoulli Society <http://www.bernoulli-society.org/>

⁵Project Euclid: <https://projecteuclid.org/>

⁶IMS Open Access Fund: <https://imstat.org/shop/donation/>