

Cutoff for a stratified random walk on the hypercube

Anna Ben-Hamou*

Yuval Peres†

Abstract

We consider the random walk on the hypercube which moves by picking an ordered pair (i, j) of distinct coordinates uniformly at random and adding the bit at location i to the bit at location j , modulo 2. We show that this Markov chain has cutoff at time $\frac{3}{2}n \log n$ with window of size n , solving a question posed by Chung and Graham (1997).

Keywords: Markov chains; mixing times; cutoff; hypercube.

AMS MSC 2010: 60J10.

Submitted to ECP on December 12, 2017, final version accepted on April 10, 2018.

1 Introduction

Setting and main result

Let $\mathcal{X} = \{0, 1\}^n \setminus \{0\}$ and consider the Markov chain $\{Z_t\}_{t \geq 0}$ on \mathcal{X} defined as follows: if the current state is x and if $x(i)$ denotes the bit at the i^{th} coordinate of x , then the walk proceeds by choosing uniformly at random an ordered pair (i, j) of distinct coordinates, and replacing $x(j)$ by $x(j) + x(i) \pmod{2}$.

The transition matrix P of this chain is symmetric, irreducible and aperiodic. Its stationary distribution π is the uniform distribution over \mathcal{X} , i.e. for all $x \in \mathcal{X}$, $\pi(x) = \frac{1}{2^n - 1}$. We are interested in the total-variation mixing time, defined as

$$t_{\text{mix}}(\varepsilon) = \min \{t \geq 0, d(t) \leq \varepsilon\},$$

where $d(t) = \max_{x \in \mathcal{X}} d_x(t)$ and $d_x(t)$ is the total-variation distance between $P^t(x, \cdot)$ and π :

$$d_x(t) = \sup_{A \subset \mathcal{X}} (\pi(A) - P^t(x, A)) = \sum_{y \in \mathcal{X}} (P^t(x, y) - \pi(y))_+.$$

Diaconis and Saloff-Coste [4] showed that the log-Sobolev constant of $\{Z_t\}_{t \geq 0}$ is $O(n^2)$, which yields an upper-bound of order $n^2 \log n$ on the ℓ_2 -mixing time. They however conjectured that the right order for the total-variation mixing was $n \log n$. Chung and Graham [3] confirmed this conjecture. They showed that the relaxation time of $\{Z_t\}$ was of order n (which yields a tight upper-bound of order n^2 for ℓ_2 -mixing) and that the total-variation mixing time $t_{\text{mix}}(\varepsilon)$ was smaller than $c_\varepsilon n \log n$ for some constant c_ε . They asked whether one could make this bound more precise and replace c_ε by a universal constant which would not depend on ε . We answer this question positively by proving that the chain $\{Z_t\}$ has cutoff at time $\frac{3}{2}n \log n$, with window of order n .

*Sorbonne Université, LPSM, France. E-mail: anna.ben-hamou@upmc.fr

†Microsoft Research, USA. E-mail: peres@microsoft.com

Theorem 1.1. *The chain $\{Z_t\}$ has total-variation cutoff at time $\frac{3}{2}n \log n$ with window n ,*

$$\lim_{\alpha \rightarrow +\infty} \liminf_{n \rightarrow +\infty} d\left(\frac{3}{2}n \log n - \alpha n\right) = 1,$$

and

$$\lim_{\alpha \rightarrow +\infty} \limsup_{n \rightarrow +\infty} d\left(\frac{3}{2}n \log n + \alpha n\right) = 0.$$

Motivation and related work

The chain $\{Z_t\}$ records the evolution of a single column in the following random walk on $\text{SL}_n(\mathbb{Z}_2)$, the group of invertible $n \times n$ matrices with coefficients in \mathbb{Z}_2 : at each step, the walk moves by picking an ordered pair of distinct rows uniformly at random and adding the first one to the other, modulo 2. This matrix random walk has received significant attention, both from group theoreticians and cryptologists. It was brought to our attention by Ron Rivest, who was mostly interested in computational mixing aspects, pertaining to *authentication protocols*. In cryptography, an authentication protocol is a scheme involving two parties, a *verifier* and a *prover*, the goal of the verifier being to certify the identity of the prover (*i.e.* to distinguish between an honest and a dishonest prover). A large family of authentication protocols, called *time-based authentication protocols*, is based on the time needed by the prover to answer a challenge. The authentication is successful if and only if the correct answer is provided fast enough. The following protocol was proposed by Sotiraki [7]. Starting from the identity matrix in $\text{SL}_n(\mathbb{Z}_2)$, the prover runs the above Markov chain driven by random row additions up to a certain time $t \in \mathbb{N}$. He makes the final matrix A_t public (this is called a *public key*), but only he knows the trajectory of the Markov chain. Then, whenever he wants to authenticate, the prover asks the verifier for a vector $x \in \{0, 1\}^n$. The challenge is to quickly compute $y = A_t x$. As the prover knows the chain's trajectory, he can apply to x the same row operations he has performed to create A_t and provide the correct answer in time t . On the other hand, if t is large enough, a dishonest party may not be able to distinguish, in polynomial time, A_t from a uniformly randomly chosen matrix (we say that the chain is computationally mixed), and its best solution would be to perform usual matrix-vector multiplication, which typically takes about n^2 operations. Hence, if the prover chooses t as the computational mixing time and if this time is shown to be much smaller than n^2 , the verifier will be able to distinguish honest provers from dishonest parties.

Now, it is reasonable to assume that if each column is mixed by time t , then the matrix A_t is computationally mixed and no dishonest party would be able to distinguish A_t from uniform (in polynomial time) and thus to answer the challenge in less than n^2 units of time, hereby motivating the study of the chain $\{Z_t\}$.

The question of determining the total-variation mixing time of the matrix walk is still largely open. Diaconis and Saloff-Coste [4] showed that the ℓ_2 -mixing time was $O(n^4)$, and the powerful results of Kassabov [5] yield the upper-bound $O(n^3)$, which is also the best known upper-bound in for total-variation mixing. By a simple counting argument, the total-variation mixing time can be lower bounded by $\Omega\left(\frac{n^2}{\log n}\right)$ (which is actually an estimate of the diameter of the underlying graph, see Andr n et al. [1], Christofides [2]).

Outline of the paper

Before proving Theorem 1.1, we first state some useful properties of the birth-and-death chain given by the Hamming weight of Z_t . In particular, we show that this

projected chain also has cutoff at $\frac{3}{2}n \log n$ (Section 2). Section 3 is then devoted to the proof of Theorem 1.1.

2 The Hamming weight

For a vertex $x \in \mathcal{X}$, we denote by $H(x)$ the Hamming weight of x , i.e.

$$H(x) = \sum_{i=1}^n x(i).$$

Consider the birth-and-death chain $H_t := H(Z_t)$, and denote by P_H , π_H , and $d_H(\cdot)$ its transition matrix, stationary distribution, and total-variation distance to equilibrium. For $1 \leq k \leq n$, we have

$$\begin{aligned} P_H(k, k+1) &= \frac{k(n-k)}{n(n-1)}, \\ P_H(k, k-1) &= \frac{k(k-1)}{n(n-1)}, \\ P_H(k, k) &= \frac{n-k}{n}, \end{aligned}$$

and

$$\pi_H(k) = \frac{\binom{n}{k}}{2^n - 1}.$$

The hitting time of state k is defined as

$$T_k = \min \{t \geq 0, H_t = k\}.$$

One standard result in birth-and-death chains is that, for $2 \leq \ell \leq n$,

$$\mathbb{E}_{\ell-1}(T_\ell) = \frac{1}{P_H(\ell, \ell-1)} \sum_{i=1}^{\ell-1} \frac{\pi_H(i)}{\pi_H(\ell)}, \tag{2.1}$$

(see for instance [6, Section 2.5]). The following lemma will be useful.

Lemma 2.1. *Let $0 < \beta < 1$ and $K = (1 - \beta)\frac{n}{2}$. Then there exist constants $a_\beta, b_\beta \in \mathbb{R}$ depending on β only such that*

$$\mathbb{E}_1(T_K) \leq n \log n + a_\beta n,$$

and

$$\text{Var}_1 T_K \leq b_\beta n^2.$$

Proof of Lemma 2.1. For $2 \leq k \leq K$, let $\mu_k = \mathbb{E}_{k-1} T_k$ and $v_k = \text{Var}_{k-1}(T_k)$. Resorting to (2.1), we have

$$\mu_k = \frac{\binom{n}{k-1}}{\binom{n-2}{k-2}} \sum_{i=1}^{k-1} \frac{\binom{n}{i}}{\binom{n}{k-1}} \leq \frac{\binom{n-1}{k-1}}{\binom{n-2}{k-2}} \sum_{i=1}^{k-1} \left(\frac{k-1}{n-k+2}\right)^{k-i-1} \leq \frac{n^2}{k(n-2k)} \tag{2.2}$$

Summing from 2 to K yields the desired bound on $\mathbb{E}_1 T_K$. Moving on to the variance, by independence of the successive hitting times, we have

$$\text{Var}_1 T_K = \sum_{k=1}^{K-1} v_{k+1}.$$

Hence, it is sufficient to show that there exists a constant $c_\beta > 0$ such that $v_{k+1} \leq \frac{c_\beta n^2}{k^2}$ for all $k \leq K$. To do so, we consider the following distributional identity for the hitting time T_{k+1} starting from k :

$$T_{k+1} = 1 + (1 - I)\tilde{T}_{k+1} + IJ(\hat{T}_k + \hat{T}_{k+1}),$$

where I is the indicator that the chain moves (i.e. that a one is picked as updating coordinate), J is the indicator that the chain decreases given that it moves (i.e. that the chosen one is added to another one), \tilde{T}_{k+1} and \hat{T}_{k+1} are copies of T_{k+1} , and \hat{T}_k is the hitting time of k starting from $k - 1$. All those variables may be assumed to be independent. We obtain the following induction relation:

$$\begin{aligned} v_{k+1} &= \frac{k-1}{n-1}(v_k + v_{k+1}) + \left(1 - \frac{k}{n}\right)\mu_{k+1}^2 + \frac{k-1}{n-1} \left(1 - \frac{k(k-1)}{n(n-1)}\right)(\mu_k + \mu_{k+1})^2 \\ &\leq \frac{k}{n}(v_k + v_{k+1}) + \mu_{k+1}^2 + \frac{k}{n}(\mu_k + \mu_{k+1})^2. \end{aligned}$$

Using the fact that for all $k \leq K$, we have $\mu_k \leq \frac{n}{\beta k}$ (which can be seen by inequality (2.2)), and after some simplification,

$$v_{k+1} \leq \frac{k}{n-k}v_k + \frac{3n^3}{\beta^2 k^2(n-k)} \leq \frac{k}{n-k}v_k + \frac{6n^2}{\beta^2 k^2}.$$

By induction and using that $v_2 \leq n^2$, we obtain that $v_{k+1} \leq \frac{c_\beta n^2}{k^2}$ for all $k \leq K$. □

The following proposition establishes cutoff for the chain $\{H_t\}$ and will be used in the next section to prove cutoff for the chain $\{Z_t\}$.

Proposition 2.2. *The chain H_t exhibits cutoff at time $\frac{3}{2}n \log n$ with window n .*

Proof. For the lower bound, we want to show that for $t = \frac{3}{2}n \log n - 2\alpha n$

$$d_H(t) \geq 1 - \varepsilon(\alpha),$$

where $\varepsilon(\alpha) \rightarrow 0$ as $\alpha \rightarrow +\infty$. Consider the chain started at $H_0 = 1$ and let $k = \frac{n}{2} - \alpha\sqrt{n}$ and $A = \{k, k + 1, \dots, n\}$. By definition of total-variation distance,

$$d_H(t) \geq \pi_H(A) - P_H^t(1, A) \geq \pi_H(A) - \mathbb{P}_1(T_k \leq t).$$

By the Central Limit Theorem, $\lim_{\alpha \rightarrow \infty} \lim_{n \rightarrow \infty} \pi_H(A) = 1$. Moving on to $\mathbb{P}_1(T_k \leq t)$, let us write

$$\mathbb{P}_1(T_k \leq t) = \mathbb{P}_1(T_{n/3} \leq n \log n - \alpha n) + \mathbb{P}_{n/3} \left(T_k \leq \frac{n \log n}{2} - \alpha n \right).$$

Note that $T_{n/3}$ is stochastically larger than $\sum_{i=1}^{n/3} G_i$, where $(G_i)_{i=1}^{n/3}$ are independent Geometric random variables with respective parameter i/n (this is because at each step, we need at least to pick a one to just move from the current position). By Chebyshev's Inequality,

$$\mathbb{P}_1(T_{n/3} \leq n \log n - \alpha n) = O\left(\frac{1}{\alpha^2}\right).$$

Now, starting from Hamming weight $n/3$ and up to time T_k , we may couple H_t with \tilde{H}_t , the Hamming weight of the standard lazy random walk on the hypercube (at each step,

pick a coordinate uniformly at random and randomize the bit at this coordinate), in such a way $T_k \geq S_k$, where $S_k = \inf\{t \geq 0, \tilde{H}_t = k\}$. It is known that S_k satisfies

$$\mathbb{P}_{n/3} \left(S_k \leq \frac{n \log n}{2} - \alpha n \right) \leq \varepsilon(\alpha),$$

with $\varepsilon(\alpha) \rightarrow 0$ as $\alpha \rightarrow +\infty$ (see for instance the proof of [6, Proposition 7.13]), which concludes the proof of the lower bound.

For the upper bound, letting $t = \frac{3}{2}n \log n + 2\alpha n$, we have

$$d_H(t) \leq \mathbb{P}_1 (T_{n/3} > n \log n + \alpha n) + \max_{k \geq n/3} d_H^{(k)} \left(\frac{n \log n}{2} + \alpha n \right). \quad (2.3)$$

Lemma 2.1 entails that $T_{n/3}$ concentrates well: $\mathbb{E}_1(T_{n/3}) = n \log n + cn$ for some absolute constant c , and $\text{Var}_1(T_{n/3}) = O(n^2)$. By Chebyshev's Inequality,

$$\mathbb{P}_1 (T_{n/3} > n \log n + \alpha n) = O\left(\frac{1}{\alpha^2}\right). \quad (2.4)$$

To control the second term in the right-hand side of (2.3), we use the coupling method (see Levin et al. [6, Corollary 5.3]). For all starting point $k \geq n/3$, we consider the following coupling between a chain H_t started at k and a chain H_t^π started from stationarity: at each step t , if H_t makes an actual move (a one is picked as updating bit in the underlying chain Z_t), we try "as much as possible" not to move H_t^π (picking a zero as updating bit). Conversely, when H_t does not move, we try "as much as possible" to move H_t^π , the goal being to increase the chance that the two chains do not cross each other (by moving at the same time). The chains stay together once they have met for the first time. We claim that the study of the coupling time can be reduced to the study of the first time when the chain started at $n/3$ reaches $n/2$. Indeed, when both chains have reached $n/2$, either they have met, or they have crossed each other. In this last situation, we know however that the expected time of their first return to $n/2$ is $O(\sqrt{n})$, so that $\mathbb{P}_{n/2} (T_{n/2}^+ > \sqrt{\alpha n}) = O(1/\sqrt{\alpha})$. Moreover, thanks to our coupling, during each of those excursions, the chains have positive probability to meet, so that after an additional time of order $\alpha\sqrt{n}$ we can guarantee that they have met with large probability. Moreover, as $\pi_H([2n/3, n]) = o(1)$, with high probability, $H_0^\pi \leq 2n/3$, and as starting from a larger Hamming weight can only speed up the chain, $\mathbb{P}_{2n/3}(T_{n/2} > t) \leq \mathbb{P}_{n/3}(T_{n/2} > t)$. We are thus left to prove that $\mathbb{P}_{n/3} (T_{n/2} > \frac{n \log n}{2} + \alpha n) \leq \varepsilon(\alpha)$, for a function ε tending to 0 at $+\infty$.

Starting from $H_0 = n/3$, we first argue that H_t will remain above $2n/7$ for a very long time. Namely, defining $\mathcal{G}_t = \{T_{2n/7} > t\}$, we have

$$\mathbb{P}_{n/3}(\mathcal{G}_{n^2}) = 1 - o(1). \quad (2.5)$$

This can easily be seen by considering $T_k^+ = \min\{t \geq 1, H_t = k\}$ and taking a union bound over the excursions around $k = n/3$ which visit $m = 2n/7$:

$$\mathbb{P}_k(T_m \leq n^2) \leq n^2 \mathbb{P}_k(T_m \leq T_k^+),$$

and

$$\mathbb{P}_k(T_m \leq T_k^+) = \frac{\mathbb{E}_k(T_k^+)}{\mathbb{E}_m(T_k) + \mathbb{E}_k(T_m)} \leq \frac{\mathbb{E}_k(T_k^+)}{\mathbb{E}_m(T_m^+)} = \frac{\pi_H(m)}{\pi_H(k)},$$

which decreases exponentially fast in n .

Our goal now will be to analyse the tail of $\tau = \inf\{t \geq 0, D_t \leq 0\}$, where

$$D_t = \frac{n}{2} - H_t.$$

Observe that

$$D_{t+1} - D_t = \begin{cases} 1 & \text{with probability } \frac{H_t(H_t-1)}{n(n-1)} \\ -1 & \text{with probability } \frac{H_t(n-H_t)}{n(n-1)} \\ 0 & \text{otherwise.} \end{cases} \quad (2.6)$$

We get

$$\mathbb{E}[D_{t+1} - D_t \mid D_t] = -\frac{2\left(\frac{n}{2} - D_t\right)(D_t + 1)}{n(n-1)} \leq -\frac{D_t}{n} + \frac{2D_t^2}{n(n-1)}. \quad (2.7)$$

Writing a similar recursion for the second moment of D_t gives

$$\mathbb{E}[D_{t+1}^2 - D_t^2 \mid D_t] = -\frac{4H_t D_t (D_t + 1/2)}{n(n-1)} + \frac{H_t}{n} \leq -\frac{4H_t D_t^2}{n^2} + 2.$$

On the event \mathcal{G}_t ,

$$\mathbb{E}[D_{t+1}^2 - D_t^2 \mid D_t] \leq -\frac{8D_t^2}{7n} + 2.$$

By induction, letting $\mathcal{D}_t = \mathbb{1}_{\mathcal{G}_t} D_t$ (and noticing that $\mathcal{G}_{t+1} \subset \mathcal{G}_t$), we get

$$\mathbb{E}[\mathcal{D}_t^2] \leq \mathbb{E}[\mathcal{D}_0^2] \left(1 - \frac{8}{7n}\right)^t + \frac{7n}{4} \leq \frac{n^2}{4} e^{-8t/7n} + 2n.$$

Plugging this back in (2.7),

$$\mathbb{E}[\mathcal{D}_{t+1}] \leq \left(1 - \frac{1}{n}\right) \mathbb{E}[\mathcal{D}_t] + e^{-8t/7n} + 4/n,$$

and by induction,

$$\mathbb{E}[\mathcal{D}_t] \leq an e^{-t/n} + b, \quad (2.8)$$

for absolute constants $a, b \geq 0$. Also, letting $\tau_\star = \inf\{t \geq 0, \mathcal{D}_t = 0\}$, we see by (2.6) that, provided $\tau_\star > t$, the process $\{\mathcal{D}_t\}$ is at least as likely to move downwards than to move upwards and that there exists a constant $\sigma^2 > 0$ such that $\text{Var}(\mathcal{D}_{t+1} \mid \mathcal{D}_t) \geq \sigma^2$ (this is because, on \mathcal{G}_t the probability to make a move at time t in larger than some positive absolute constant). By Levin et al. [6, Proposition 17.20], we know that for all $u > 0$ and $k \geq 0$,

$$\mathbb{P}_k(\tau_\star > u) \leq \frac{4k}{\sigma\sqrt{u}}. \quad (2.9)$$

Now take $H_0 = n/3$, $D_0 = n/6$, $s = \frac{1}{2}n \log n$ and $u = \alpha n$. We have

$$\mathbb{P}_{D_0}(\tau > s + u) \leq \mathbb{P}_{D_0}(\tau_\star > s + u) + \mathbb{P}_{H_0}(\mathcal{G}_{n^2}^c).$$

By equation (2.5), $\mathbb{P}_{H_0}(\mathcal{G}_{n^2}^c) = o(1)$, and, combining (2.9) and (2.8), we have

$$\mathbb{P}_{D_0}(\tau_\star > s + u) = \mathbb{E}_{D_0}[\mathbb{P}_{D_s}(\tau_\star > u)] \leq \mathbb{E}_{D_0}\left[\frac{4D_s}{\sigma\sqrt{u}}\right] = O\left(\frac{1}{\sqrt{\alpha}}\right),$$

which implies

$$\max_{k \geq n/3} d_H^{(k)}(s + u) = O\left(\frac{1}{\sqrt{\alpha}}\right), \quad (2.10)$$

and concludes the proof of the upper bound. \square

3 Proof of Theorem 1.1

First note that, as projections of chains can not increase total-variation distance, the lower bound on $d(t)$ readily follows from the lower bound on $d_H(t)$, as established in Proposition 2.2. Therefore, we only have to prove the upper bound.

Let $\mathcal{E} = \{x \in \mathcal{X}, H(x) \geq n/3\}$ and $\tau_{\mathcal{E}}$ be the hitting time of set \mathcal{E} . For all $t, s > 0$, we have

$$d(t+s) \leq \max_{x_0 \in \mathcal{X}} \mathbb{P}_{x_0}(\tau_{\mathcal{E}} > s) + \max_{x \in \mathcal{E}} d_x(t).$$

By (2.4), taking $s = n \log n + \alpha n$, we have $\max_{x_0 \in \mathcal{X}} \mathbb{P}_{x_0}(\tau_{\mathcal{E}} > s) = O(1/\alpha^2)$, so that our task comes down to showing that for all $x \in \mathcal{E}$,

$$d_x\left(\frac{n \log n}{2} + \alpha n\right) \leq \varepsilon(\alpha),$$

with $\varepsilon(\alpha) \rightarrow 0$ as $\alpha \rightarrow +\infty$. Let us fix $x \in \mathcal{E}$. Without loss of generality, we may assume that x is the vertex with $\bar{x} \geq n/3$ ones on the first \bar{x} coordinates, and $n - \bar{x}$ zeros on the last $n - \bar{x}$ coordinates. We denote by $\{Z_t\}$ the random walk started at $Z_0 = x$ and for a vertex $z \in \mathcal{X}$, we define a two-dimensional object $\mathbf{W}(z)$, keeping track of the number of ones within the first \bar{x} and last $n - \bar{x}$ coordinates of z , that is

$$\mathbf{W}(z) = \left(\sum_{i=1}^{\bar{x}} z(i), \sum_{i=\bar{x}+1}^n z(i) \right).$$

The projection of $\{Z_t\}_{t \geq 0}$ induced by \mathbf{W} will be denoted $\mathbf{W}_t = \mathbf{W}(Z_t) = (X_t, Y_t)$. We argue that the study of $\{Z_t\}_{t \geq 0}$ can be reduced to the study of $\{\mathbf{W}_t\}_{t \geq 0}$, and that, when coupling two chains distributed as \mathbf{W}_t , we can restrict ourselves to initial states with the same total Hamming weight. Indeed, letting $\nu_{\bar{x}}$ be the uniform distribution over $\{z \in \mathcal{X}, H(z) = \bar{x}\}$, by the triangle inequality

$$d_{x_0}(t) \leq \|\mathbb{P}_x(Z_t \in \cdot) - \mathbb{P}_{\nu_{\bar{x}}}(Z_t \in \cdot)\|_{\text{tv}} + \|\mathbb{P}_{\nu_{\bar{x}}}(Z_t \in \cdot) - \pi(\cdot)\|_{\text{tv}} \quad (3.1)$$

Starting from $\nu_{\bar{x}}$, the conditional distribution of Z_t given $\{H(Z_t) = h\}$ is uniform over $\{y \in \mathcal{X}, H(y) = h\}$. This entails

$$\|\mathbb{P}_{\nu_{\bar{x}}}(Z_t \in \cdot) - \pi(\cdot)\|_{\text{tv}} = \|\mathbb{P}_{\bar{x}}(H_t \in \cdot) - \pi_H(\cdot)\|_{\text{tv}}.$$

For $t = \frac{n \log n}{2} + \alpha n$, we know by (2.10) in the proof of Proposition 2.2 that $\|\mathbb{P}_{\bar{x}}(H_t \in \cdot) - \pi_H(\cdot)\|_{\text{tv}} + O(1/\sqrt{\alpha})$. As for the first term in the right-hand side of (3.1), note that if z and z' are two vertices such that $\mathbf{W}(z) = \mathbf{W}(z')$, then for all $t \geq 0$, $\mathbb{P}_x(Z_t = z) = \mathbb{P}_x(Z_t = z')$, and that for all $y \in \mathcal{X}$ such that $\mathbf{W}(y) = (k, \ell)$

$$\begin{aligned} \mathbb{P}_{\nu_{\bar{x}}}(Z_t = y) &= \sum_{\substack{i,j \\ i+j=\bar{x}}} \sum_{z, \mathbf{W}(z)=(i,j)} \frac{1}{\binom{n}{\bar{x}}} \mathbb{P}_z(Z_t = y) \\ &= \sum_{\substack{i,j \\ i+j=\bar{x}}} \frac{\binom{\bar{x}}{i} \binom{n-\bar{x}}{j}}{\binom{n}{\bar{x}}} \sum_{z, \mathbf{W}(z)=(i,j)} \frac{\mathbb{P}_z(Z_t = y)}{\binom{\bar{x}}{i} \binom{n-\bar{x}}{j}} \\ &= \sum_{\substack{i,j \\ i+j=\bar{x}}} \frac{\binom{\bar{x}}{i} \binom{n-\bar{x}}{j}}{\binom{n}{\bar{x}}} \frac{\mathbb{P}_{(i,j)}(\mathbf{W}_t = (k, \ell))}{\binom{\bar{x}}{k} \binom{n-\bar{x}}{\ell}}. \end{aligned}$$

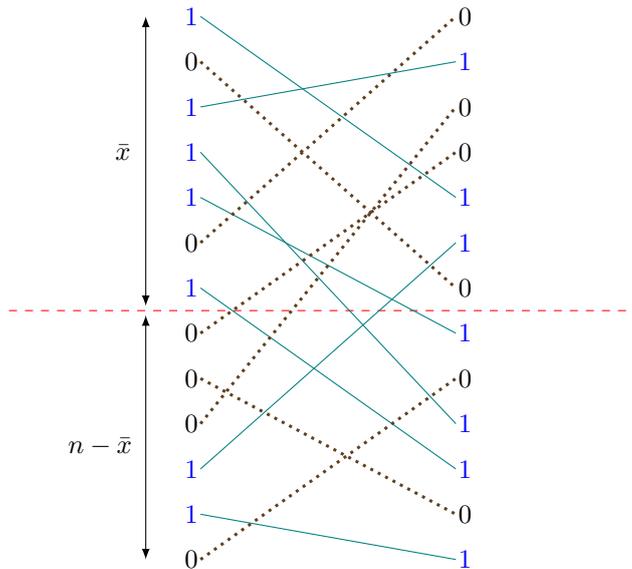


Figure 1: A pairing of coordinates of Z_t and \tilde{Z}_t .

Hence,

$$\|\mathbb{P}_x(Z_t \in \cdot) - \mathbb{P}_{\nu_{\bar{x}}}(Z_t \in \cdot)\|_{\text{tv}} \leq \max_{\substack{i,j \\ i+j=\bar{x}}} \|\mathbb{P}_{(\bar{x},0)}(\mathbf{W}_t \in \cdot) - \mathbb{P}_{(i,j)}(\mathbf{W}_t \in \cdot)\|_{\text{tv}}.$$

Now let $y \in \mathcal{E}$ such that $H(y) = \bar{x}$, and consider the chains Z_t, \tilde{Z}_t started at x and y respectively. Let $\mathbf{W}(Z_t) = (X_t, Y_t)$ and $\mathbf{W}(\tilde{Z}_t) = (\tilde{X}_t, \tilde{Y}_t)$. We couple Z_t and \tilde{Z}_t as follows: at each step t , provided $H(Z_t) = H(\tilde{Z}_t)$ and $\mathbf{W}(Z_t) \neq \mathbf{W}(\tilde{Z}_t)$, we consider a random permutation π_t which is such that $Z_t(i) = \tilde{Z}_t(\pi_t(i))$ for all $1 \leq i \leq n$, that is, we pair uniformly at random the ones (resp. the zeros) of Z_t with the ones (resp. the zeros) of \tilde{Z}_t (one such pairing of coordinates is depicted in Figure 1). If Z_t moves to Z_{t+1} by choosing the pair (i_t, j_t) and updating $Z_t(j_t)$ to $Z_t(j_t) + Z_t(i_t)$, then we move from \tilde{Z}_t to \tilde{Z}_{t+1} by updating $\tilde{Z}_t(\pi_t(j_t))$ to $\tilde{Z}_t(\pi_t(j_t)) + \tilde{Z}_t(\pi_t(i_t))$. Once $\mathbf{W}(Z_t) = \mathbf{W}(\tilde{Z}_t)$, the permutation π_t is chosen in such a way that the ones in the top (resp. in the bottom) in Z_t are matched with the ones in the top (resp. in the bottom) in \tilde{Z}_t , guaranteeing that from that time $\mathbf{W}(Z_t)$ and $\mathbf{W}(\tilde{Z}_t)$ remain equal. Note that this coupling ensures that for all $t \geq 0$, the Hamming weight of Z_t is equal to that of \tilde{Z}_t , and we may unequivocally denote it by H_t . In particular, coupling of the chains $\mathbf{W}(Z_t)$ and $\mathbf{W}(\tilde{Z}_t)$ occurs when X_t and \tilde{X}_t are matched. As $X_t \geq \tilde{X}_t$ for all $t \geq 0$, we may consider

$$\tau = \inf\{t \geq 0, \mathbf{D}_t = 0\},$$

where $\mathbf{D}_t = X_t - \tilde{X}_t$.

Before analyzing the behavior of $\{\mathbf{D}_t\}$, we first notice that the worst possible y for the coupling time satisfies $\mathbf{W}(y) = (\max\{0, 2\bar{x} - n\}, \min\{\bar{x}, n - \bar{x}\})$. We now fix y to be such a vertex, and show that, starting from x, y , the variables $\mathbf{W}(Z_t), \mathbf{W}(\tilde{Z}_t)$ remain “nice” for a very long time. More precisely, defining

$$\mathcal{B}_t = \bigcap_{s=0}^t \left\{ H_s \geq 2n/7, X_s \geq \frac{\bar{x}}{p}, \tilde{Y}_s \geq \frac{\min\{\bar{x}, n - \bar{x}\}}{p} \right\},$$

we claim that we can choose $p \geq 1$ fixed such that

$$\mathbb{P}_{x,y}(\mathcal{B}_{n^2}) = 1 - o(1). \tag{3.2}$$

Indeed, the fact that $\mathbb{P}_{n/3}(T_{2n/7} \leq n^2) = o(1)$ has already been established in the proof of Proposition 2.2 (equation (2.5)). Let us show, with the same kind of arguments, that $\mathbb{P}_{(\bar{x},0)}\left(\bigcup_{s=0}^{n^2}\{X_s < \bar{x}/p\}\right) = o(1)$. Letting $A = \{(\bar{x}/p, \ell), \ell = 0, \dots, n - \bar{x}\}$, $\pi_{\mathbf{W}}$ be the stationary distribution of \mathbf{W}_t , and $k_{\bar{x}} = \min\{\frac{\bar{x}}{2}, \frac{n-\bar{x}}{2}\}$, we have

$$\begin{aligned} \mathbb{P}_{(\bar{x},0)}(T_A \leq n^2) &\leq \mathbb{P}_{(\bar{x}/2, k_{\bar{x}})}(T_A \leq n^2) \leq n^2 \sum_{\ell=0}^{n-\bar{x}} \mathbb{P}_{(\bar{x}/2, k_{\bar{x}})}\left(T_{(\bar{x}/p, \ell)} \leq T_{(\bar{x}/2, k_{\bar{x}})}^+\right) \\ &\leq n^2 \sum_{\ell=0}^{n-\bar{x}} \frac{\pi_{\mathbf{W}}(\bar{x}/p, \ell)}{\pi_{\mathbf{W}}(\bar{x}/2, k_{\bar{x}})} = \frac{n^2 2^{n-\bar{x}} \binom{\bar{x}}{\bar{x}/p}}{\binom{\bar{x}}{\bar{x}/2} \binom{n-\bar{x}}{k_{\bar{x}}}}, \end{aligned}$$

and we can choose p large enough such that this quantity decreases exponentially fast in n . Similarly, starting from y , the value of \tilde{Y}_s will remain at a high level for a very long time, establishing (3.2).

Let us now turn to the analysis of $\{\mathbf{D}_t\}$. On the event $\{t < \tau\}$,

$$\mathbf{D}_{t+1} - \mathbf{D}_t = \begin{cases} 1 & \text{with probability } p_1^t \\ -1 & \text{with probability } p_{-1}^t \\ 0 & \text{otherwise,} \end{cases} \tag{3.3}$$

where

$$p_1^t = \frac{H_t}{n} \cdot \frac{n - H_t}{n - 1} \cdot \frac{\bar{x} - X_t}{n - H_t} \cdot \frac{n - \bar{x} - \tilde{Y}_t}{n - H_t} + \frac{H_t}{n} \cdot \frac{H_t - 1}{n - 1} \cdot \frac{Y_t}{H_t} \cdot \frac{\tilde{X}_t}{H_t},$$

and

$$p_{-1}^t = \frac{H_t}{n} \cdot \frac{n - H_t}{n - 1} \cdot \frac{\bar{x} - \tilde{X}_t}{n - H_t} \cdot \frac{n - \bar{x} - Y_t}{n - H_t} + \frac{H_t}{n} \cdot \frac{H_t - 1}{n - 1} \cdot \frac{X_t}{H_t} \cdot \frac{\tilde{Y}_t}{H_t}.$$

After computation, we get, on $\{t < \tau\}$,

$$\mathbb{E}\left[\mathbf{D}_{t+1} - \mathbf{D}_t \mid Z_t, \tilde{Z}_t\right] = -\frac{H_t \mathbf{D}_t}{n(n-1)} \left(1 + \frac{H_t - 1}{H_t}\right) \leq -\frac{\mathbf{D}_t}{n^2} (2H_t - 1) \tag{3.4}$$

From (3.4), it is not hard to see that the variable

$$M_t = \mathbb{1}_{\{\tau > t\}} \mathbf{D}_t \exp\left(\sum_{s=0}^{t-1} \frac{(2H_s - 1)}{n^2}\right)$$

is a super-martingale, which implies $\mathbb{E}_{x,y}[M_t] \leq \mathbb{E}_{x,y}[\mathbf{D}_0] \leq n$.

Now let $\tau_\star = \inf\{t \geq 0, \mathbb{1}_{\mathcal{B}_t} \mathbf{D}_t = 0\}$. By (3.3), we see that, provided $\{\tau_\star > t\}$, the process $\{\mathbb{1}_{\mathcal{B}_t} \mathbf{D}_t\}$ is a supermartingale ($p_{-1}^t \geq p_1^t$) and that there exists a constant $\sigma^2 > 0$ such that the conditional variance of its increments is larger than σ^2 (because on \mathcal{B}_t , the probability to make a move $p_{-1}^t + p_1^t$ is larger than some absolute constant). By Levin et al. [6, Proposition 17.20], for all $u > 0$ and $k \geq 0$,

$$\mathbb{P}_k(\tau_\star > u) \leq \frac{4k}{\sigma\sqrt{u}}. \tag{3.5}$$

Now take $t = \frac{n \log n}{2}$ and $u = \alpha n$. We have

$$\mathbb{P}_{x,y}(\tau > t + u) \leq \mathbb{P}_{x,y}(\mathcal{B}_{n^2}^c) + \mathbb{P}_{x,y}(\tau_\star > t + u).$$

By (3.2), we know that $\mathbb{P}_{x,y}(\mathcal{B}_{n^2}^c) = o(1)$. Also, considering the event

$$\mathcal{A}_{t-1} = \left\{ \sum_{s=0}^{t-1} H_s \geq \frac{n^2 \log n}{4} - \beta n^2 \right\},$$

and resorting to (3.5), we get

$$\begin{aligned} \mathbb{P}_{x,y}(\tau_* > t + u) &\leq \mathbb{E}_{x,y} \left[\mathbb{1}_{\{\tau_* > t\}} \mathbb{P}_{Z_t, \tilde{Z}_t}(\tau_* > u) \right] \\ &\leq \mathbb{P}_{x,y}(\{\tau_* > t\} \cap \mathcal{A}_{t-1}^c) + \mathbb{E}_{x,y} \left[\mathbb{1}_{\mathcal{A}_{t-1}} \mathbb{1}_{\{\tau_* > t\}} \frac{4\mathbf{D}_t}{\sigma\sqrt{u}} \right]. \end{aligned}$$

On the one hand, recalling the notation and results of Section 2 (in particular equation (2.8)), and applying Markov's Inequality,

$$\begin{aligned} \mathbb{P}_{x,y}(\{\tau_* > t\} \cap \mathcal{A}_{t-1}^c) &\leq \mathbb{P}_{x,y} \left(\sum_{s=0}^{t-1} \mathcal{D}_s > \beta n^2 \right) \\ &\leq \frac{1}{\beta n^2} \sum_{s=0}^{t-1} (an e^{-s/n} + b) = O\left(\frac{1}{\beta}\right). \end{aligned}$$

On the other hand,

$$\mathbb{E}_{x,y} \left[\mathbb{1}_{\{\tau_* > t\}} \mathbb{1}_{\mathcal{A}_{t-1}} \mathbf{D}_t \right] \leq \exp\left(-\frac{\log n}{2} + \frac{t}{n^2} + 2\beta\right) \mathbb{E}_{x,y} [M_t] = O(e^{2\beta} \sqrt{n}).$$

In the end, we get

$$\mathbb{P}_{x,y}(\tau > t + u) = O\left(\frac{1}{\beta} + \frac{e^{2\beta}}{\sqrt{\alpha}}\right).$$

Taking for instance $\beta = \frac{1}{5} \log \alpha$ concludes the proof of Theorem 1.1.

References

- [1] D. Andr n, L. Hellstr m, and K. Markstr m. On the complexity of matrix reduction over finite fields. *Advances in applied mathematics*, 39(4):428–452, 2007. MR-2356430
- [2] D. Christofides. The asymptotic complexity of matrix reduction over finite fields. *arXiv preprint arXiv:1406.5826*, 2014.
- [3] F. R. K. Chung and R. L. Graham. Stratified random walks on the n-cube. *Random Structures and Algorithms*, 11(3):199–222, 1997. MR-1609505
- [4] P. Diaconis and L. Saloff-Coste. Walks on generating sets of abelian groups. *Probability theory and related fields*, 105(3):393–421, 1996. MR-1425868
- [5] M. Kassabov. Kazhdan constants for $SL_n(\mathbb{Z})$. *International Journal of Algebra and Computation*, 15(05n06):971–995, 2005. MR-2197816
- [6] D. Levin, Y. Peres, and E. Wilmer. *Markov chains and mixing times*. AMS Bookstore, 2009. MR-2466937
- [7] A. Sotiraki. *Authentication protocol using trapdoored matrices*. PhD thesis, Massachusetts Institute of Technology, 2016.

Acknowledgments. We thank Persi Diaconis and Ohad Feldheim for helpful discussions and references. We thank Ryokichi Tanaka and Alex Lin Zhai for helpful comments and corrections. We also thank Ron Rivest for suggesting this problem to us.