Harm Nooitgedagt

Two convergence limits of Markov chains: Cutoff and Metastability.

Master thesis, defended on Augustus 31, 2010

Thesis advisor: Prof. dr. W.T.F. den Hollander

Mastertrack: Applied Mathematics

Mathematisch Instituut, Universiteit Leiden

1 Introduction

We start this master thesis by looking at card shuffles, in particular the top-to-random shuffle, which will be discussed in detail in section 3. Via coupling we can prove that for an n-card deck the distance to a uniform configuration of the deck in the limit as $n \to \infty$ is reached after approximately $n \log n + cn$ shuffles, with c a large constant. In other words, after $n \log n + cn$ top-to-random shuffles the deck is thoroughly shuffled. We also prove that if we only apply $n \log n - cn$ top-to-random shuffles to the n-card deck, then the shuffled cards still closely resemble the original configuration. So the configuration drops from "being close to the original configuration" to"close to uniform" at approximately $n \log n$ shuffles with a width of order n.

This drop is called *cutoff* and forms one of the subjects of this master thesis. A question we would like to answer is: What conditions do we need on the state space and on the transition kernel of a Markov chain so that it shows cutoff? To answer this question we start by searching for other examples. Diaconis has done a lot of work in this area, see; [4], [5] and [6]. He gives an explanation for the cutoff phenomenon, discusses the top-to-random shuffle and the riffle shuffle and applies his theory about cutoff to a series of other problems. The preliminaries to this phenomenon and a description of what cutoff exactly is in a mathematical sense is given in section 2. The top-to-random and the riffle shuffle are discussed in detail in section 3 and another explanation of cutoff via eigenvalues and eigenvectors is given in section 4.

The second subject of this master thesis is called *metastability*, which is in complete contrast with cutoff. In cutoff, the Markov chain starts from a given state and stays in the vicinity of this state until after a specific amount of time it becomes uniformly distributed over the whole state space. But if the Markov chain has a metastable state, i.e., a deep well, then after an exponentially distributed time, it will climb out of this well and move on to a deeper well. As the name "well" suggests, it is easy to fall in but hard to climb out. If a Markov chain sits in a metastable state consisting of a local minimum, it has to climb out of this "well" before it can topple into a deeper well concentrated around the global minimum, which will take an exponentially distributed time and therefor is not sharp as in cutoff.

As for cutoff we would like to answer the question: What conditions do we need on the state space and on the transition kernel of a Markov chain so that it shows metastability? Again we start by looking at an example. In section 5 we will study an example of an one dimensional random walk in a double well. After that we will discuss some real-world examples where metastability plays a key role, namely the Glauber and Kawasaki dynamics, which we will discuss in section 6 and in section 7. In these sections we will explain how these dynamics work by explaining the ideas behind the theorems stated in [3].

Finally, in the last part of this thesis, part III, we will give some insight on how the cutoff and metastability are linked.

Contents

1 Introduction **3**

I The Cutoff Phenomena 7

2 Prelimenaries **9**

3 Random shuffles **10**
3.1 Top-to-random shuffle 10
3.2 GSR-shuffle 13

4 Bounds on total variation **17**
4.1 Example without cutoff 20
4.2 Example with cutoff 22

II Metastability 27

5 1-Dimensional Example of Metastability **29**

6 A Ferromagnet and Glauber Dynamics **37**

7 A Supersaturated Gas and Kawasaki Dynamics **43**

III Conclusion 49

Part I
The Cutoff Phenomena

2 Prelimenaries

Definition 2.1 *The total variation distance between two distributions λ and μ on the same state space Ω is defined as*

$$\|\mu - \lambda\|_{TV} = \max_{A \subset \Omega} |\mu(A) - \lambda(A)|.$$

The total variation distance is a natural way to quantify how much two distributions differ from each other. There are many other definitions, that serve the same purpose. From basic Markov theory we know that if we have an irreducible and aperiodic Markov chain on a finite state space Ω with transition kernel P, then

$$\lim_{t \to \infty} P^t(x, \cdot) = \pi(\cdot) \quad \forall x \in \Omega,$$

where π denotes the stationary distribution and t runs through the integers $\mathbb{N}_0$.
In terms of the total variation distance this becomes the following theorem

Theorem 2.2 (Markov Chain Convergence Theorem) *Suppose that P is an irreducible and aperiodic transition kernel, with stationary distribution π. Then there exists constants $\alpha \in (0,1)$ and $C \in (0, \infty)$ such that*

$$\max_{x \in \Omega} \left\| P^t(x, \cdot) - \pi \right\|_{TV} \leq C\alpha^t \quad \forall t \in \mathbb{N}_0.$$

We speak of a *cutoff phenomenon* when the total variation distance drops from being close to 1 to being close to 0 within a time interval, whose width is small compared to the location of its center. Notice that the Markov chain convergence theorem (MCCT) above does not yet give sharp bounds for this phenomenon.

Before we go into further detail we need to standardize some notation. First define the *maximal distance* between $P^t(x, \cdot)$ and π as

$$d(t) := \max_{x \in \Omega} \left\| P^t(x, \cdot) - \pi \right\|_{TV}.$$

For $\varepsilon \in (0,1)$ we define the *mixing time* as

$$t_{mix}(\varepsilon) := \min\{t : d(t) \leq \varepsilon\},$$

and we take

$$t_{mix} := t_{mix}(1/4).$$

With this definition we can define the cutoff phenomenon.

Definition 2.3 *Suppose we have a sequence of Markov chains indexed by $n \in \mathbb{N}$ with corresponding mixing times $t_{mix}^{(n)}(\varepsilon)$. This sequence has a cutoff if, for all $\varepsilon > 0$ we have,*

$$\lim_{n \to \infty} \frac{t_{mix}^{(n)}(\varepsilon)}{t_{mix}^{(n)}(1-\varepsilon)} = 1.$$

In terms of the total variation distance to stationarity we have the following characterization of the cutoff phenomenon.

Corollary 2.4 *Let $t_{mix}^{(n)}$ and d_n be the mixing time and distance to stationary, respectively, for the n-th chain in a sequence of Markov chains. The sequence has a cutoff if and only if*

$$\lim_{n \to \infty} d_n(ct_{mix}^{(n)}) = \begin{cases} 1 & \text{if } 0 < c < 1, \\ 0 & \text{if } c > 1. \end{cases}$$

Such a sequence of mixing times are also called the *threshold times* of a sequence of Markov chains. Furthermore remark that we can replace $t_{mix}^{(n)}$ with $t_{mix}^{(n)}(\varepsilon)$ and $\varepsilon \in (0, 1/2)$. [1]

[1] The proofs of Theorem 2.2 and Corollary 2.4 can be found in [7].

3 Random shuffles

One of the examples where the cutoff phenomenon occurs is in card shuffles. In this section we will discuss two shuffles, the top-to-random shuffle and the riffle-shuffle. Before going to the details we first need some definitions.

Definition 3.1 *We call a stopping time τ, possibly depending on the starting position x, a* strong uniform time *with respect to a Markov chain $X = (X_t)_{t\in\mathbb{N}_0}$ on Ω with stationary distribution π if*

$$\mathbb{P}_x(\tau = t,\ X_\tau = y) = \mathbb{P}_x(\tau = t)\pi(y) \quad \forall t \in \mathbb{N}_0, y \in \Omega.$$

Now we can state a lemma that gives an upper bound for the distance between the Markov chain and its stationary distribution.

Lemma 3.2 *Let $X = (X_t)_{t\in\mathbb{N}_0}$ be an irreducible and aperiodic Markov chain on Ω with stationary distribution π. If τ is a strong uniform time for X. Then*

$$d(t) = \max_{x\in\Omega} \|\mathbb{P}_x(X_t \in \cdot) - \pi(\cdot)\|_{TV} \leq \max_{x\in\Omega} \mathbb{P}_x(\tau > t) \quad \forall t \in \mathbb{N}_0.$$

Proof: For $t \in \mathbb{N}_0$, $x \in \Omega$ and $A \subset \Omega$ we have

$$\begin{aligned}
\mathbb{P}_x(X_t \in A) &= \mathbb{P}_x(X_t \in A,\ \tau > t) + \sum_{i\leq t} \mathbb{P}_x(X_t \in A,\ \tau = i), \\
&= \mathbb{P}_x(X_t \in A|\tau > t)\mathbb{P}_x(\tau > t) + \pi(A)\sum_{i\leq t} \mathbb{P}_x(\tau = i), \\
&= \{\mathbb{P}_x(X_t \in A|\tau > t) - \pi(A)\}\,\mathbb{P}_x(\tau > t) + \pi(A).
\end{aligned}$$

Since $|\mathbb{P}_x(X_t \in A|\tau > t) - \pi(A)| \leq 1$ the claim follows by taking the maximum over $A \subset \Omega$ and using the definition of $d(t)$.

□

3.1 Top-to-random shuffle

Consider a deck of n cards, labeled $1, 2, \ldots, n$. An arrangement of the deck is an element from the group of permutations on n elements, denoted as $\mathcal{S}_n$, also called the symmetric group. The top card is represented by the first element of the arrangement and the bottom card by the last. A *shuffle* of the deck is represented by applying to the deck a permutation drawn from $\mathcal{S}_n$. A *random shuffle* is a shuffle applied to the deck drawn according to some probability distribution on $\mathcal{S}_n$. Repeatedly shuffling the deck according to one distribution, say μ, is equivalent to running a random walk on $\mathcal{S}_n$, with independent increments according to μ. If the support of μ is equal to $\mathcal{S}_n$ and $\mu(\mathrm{id}) > 0$, where id denotes the identity permutation, then this random walk is irreducible and aperiodic. Furthermore, the stationary distribution π is uniform on $\mathcal{S}_n$.

In the top-to-random shuffle we take the top card of the deck and insert it back into the deck randomly, i.e., uniformly at any of the n locations including the top position itself. The cards above the chosen position all move upward by one position. We need about $n \log n$ shuffles to get the deck close to random. To show why, we follow the bottom card of the deck. Let τ_1 be the first time that a card is inserted below the original bottom card. Since the probability that a shuffle puts a card at the bottom is equal to $1/n$, this happens after approximately n shuffles. Next, let τ_2 be the first time after τ_1 that a second card is placed below the original bottom card. This happens with probability $2/n$, so it takes approximately $n/2$ more shuffles. The two cards under the original bottom card are equally likely to be in any order. In a similar way we can define times $\tau_1 < \tau_2 < \ldots < \tau_{n-1}$.

At time τ_{n-1} the card originally from the bottom comes on top. By induction it follows that the bottom $n-1$ cards are equally likely to be in any of the $(n-1)!$ possible arrangements. One final shuffle, at time $\tau_n = \tau_{n-1} + 1$, puts the original bottom card back into the deck randomly at one of the n positions. Hence the deck is equally likely to be in any of the $n!$ arrangements in $\mathcal{S}_n$ at time τ_n.

When the original bottom card is at position i it takes approximately n/i shuffles to place a card below it. Hence we have that

$$\mathbb{E}(\tau_i - \tau_{i-1}) = \frac{n}{i}$$

and since $\tau_n = \sum_{i=1}^{n}(\tau_i - \tau_{i-1})$, with $\tau_0 = 0$, we have

$$\begin{aligned}\mathbb{E}(\tau_n) &= \mathbb{E}\left(\sum_{i=1}^{n}(\tau_i - \tau_{i-1})\right) = \sum_{i=1}^{n}\mathbb{E}(\tau_i - \tau_{i-1}) \\ &= \sum_{i=1}^{n}\frac{n}{i} \sim n\log n \quad \text{as } n \to \infty.\end{aligned}$$

Remark that τ_n is strong uniform time. The next theorem makes the cutoff precise.

Theorem 3.3 *For the top-to-random shuffle and any $\varepsilon > 0$, there exists a constant α_0 such that*
(i) $d_n(n\log n + cn) \le e^{-c}$ for $c \ge 0$, $n \ge 2$.
(ii) $d_n(n\log n - \alpha n) \ge 1 - \varepsilon$ for $\alpha > \alpha_0$.

To prove Theorem 3.3 we will compare the top-to-random shuffle with an auxiliary random variable V_n, which represents the number of draws needed to draw all n balls, with replacement, at least once from an urn that contains n balls. For $i = 0, 1 \ldots, n$, let V_i denote the first time, i.e., the number of draws needed, until i distinct balls have been drawn. Remark that

$$\tau_i - \tau_{i-1} \stackrel{D}{=} V_{n-(i-1)} - V_{n-i} \stackrel{D}{=} \text{Geo}\left(\frac{i}{n}\right), \quad i = 1, \ldots, n,$$

are independent and that

$$\tau_n = \sum_{i=1}^{n}(\tau_i - \tau_{i-1}) \stackrel{D}{=} \sum_{i=1}^{n}(V_{n-(i-1)} - V_{n-i}) = V_n.$$

Proof of Theorem 3.3 Label the balls $1, \ldots, n$ and let A_i be the event that ball i is not drawn in the first $n\log n + cn$ draws. Then

$$\begin{aligned}\mathbb{P}(\tau_n > n\log n + cn) &= \mathbb{P}(V_n > n\log n + cn) \\ &= \mathbb{P}(\cup_{i=1}^{n} A_i) \le \sum_{i=1}^{n}\mathbb{P}(A_i) \\ &= n(1 - \frac{1}{n})^{n\log n + cn} \\ &\le n e^{-\log n - c} = e^{-c},\end{aligned}$$

where the second inequality follows from the fact that $1 - x \le e^{-x}$ for $0 < x \le 1$. Since τ_n is a strong stationary time, we can combine this result with Lemma 3.2 to get Theorem 3.3(i).

To prove Theorem 3.3(ii) assume that the original deck is in the arrangement id and consider the events

$$B_j := \{\text{the original bottom } j \text{ cards are in their original relative order}\}, \quad j = 1, \ldots, n.$$

Let τ'_j denote the number of shuffles required to move the jth card from the bottom to the top, and let τ_{ji} denote the number of shuffles required to put i cards below the jth card. Since

$$\tau_{j(i+1)} - \tau_{ji} \stackrel{D}{=} \mathrm{Geo}((i+j)/n)$$

we have

$$\tau'_j = \sum_{i=0}^{n-j-1} (\tau_{j(i+1)} - \tau_{ji}) \stackrel{D}{=} \sum_{i=j}^{n-1} (V_{n-(i-1)} - V_{n-i}).$$

Since $\mathbb{E}(V_{n-(i-1)} - V_{n-i}) = n/i$ and $\mathrm{Var}(V_{n-(i-1)} - V_{n-i}) = (n^2(1-i/n))/i^2 < n^2/i^2$ we have

$$\begin{aligned}
\mathbb{E}(\tau'_j) &= \mathbb{E}\left(\sum_{i=j}^{n-1} (V_{n-i} - V_{n-(i+1)})\right) \\
&= \sum_{i=j}^{n-1} \mathbb{E}(V_{n-i} - V_{n-(i+1)}) \\
&= \sum_{i=j}^{n-1} \frac{n}{i} \geq n(\log n - \log j - 1)
\end{aligned}$$

and

$$\mathrm{Var}(\tau'_j) = \sum_{i=j}^{n-1} \frac{n^2}{i^2} < \frac{n^2}{j-1}, \quad j \geq 2.$$

Combining these bounds with Chebyshev's inequality we get

$$\begin{aligned}
\mathbb{P}(\tau'_j < n\log n - \alpha n) &\leq \mathbb{P}(\tau'_j - \mathbb{E}(\tau'_j) < -n(\alpha - \log j - 1)) \\
&\leq \mathbb{P}(|\tau'_j - \mathbb{E}(\tau'_j)| > n(\alpha - \log j - 1)) \\
&\leq \frac{\mathrm{Var}(\tau'_j)}{n^2(\alpha - \log j - 1)^2} \\
&< \frac{1}{(j-1)(\alpha - \log j - 1)^2} \\
&\leq \frac{1}{j-1}, \quad \alpha \geq \log j + 2.
\end{aligned}$$

Take $t_n(\alpha) = n\log n - \alpha n$. Then

$$\mathbb{P}(X_{t_n(\alpha)} \in B_j) \geq \mathbb{P}(\tau'_j \geq t_n(\alpha)) = 1 - \mathbb{P}(\tau'_j < t_n(\alpha)) > 1 - \frac{1}{j-1}.$$

Together with the observation that $\pi(B_j) = 1/j! \leq 1/(j-1)$, we get

$$d_n(t_n(\alpha)) = \left\|\mathbb{P}(X_{t_n(\alpha)} \in \cdot) - \pi(\cdot)\right\|_{TV} \geq \mathbb{P}(X_{t_n(\alpha)} \in B_j) - \pi(B_j) > 1 - \frac{2}{j-1}.$$

Now take $j = e^{\alpha-2}$, so that $n \geq e^{\alpha-2}$, and define

$$g(\alpha) := \frac{2}{e^{\alpha-2} - 1}.$$

Then for all $\varepsilon > 0$ there exists an $\alpha_0 := \alpha_0(\varepsilon)$ with $g(\alpha_0) = 1 - \varepsilon$. Hence, for all $\alpha \geq \alpha_0$ we have

$$d_n(n\log n - \alpha n) > g(\alpha) \geq 1 - \varepsilon,$$

which proves Theorem 3.3.(ii).

□

Combining Theorem 3.3 with Corollary 2.4 we see that the top-to-random shuffle has a cutoff with threshold time $t^n_{mix} = n\log n$, $n \in \mathbb{N}_0$.

3.2 GSR-shuffle

The shuffle most commonly used by card players is called the riffle shuffle, where the dealer cuts the deck into two packets and then riffles them together. Such a shuffle can be modeled mathematically relatively easy. First note that the sizes off the two packets can vary. Furthermore, since most dealers are not professionals, the dealer drops a varying number of cards from each stack until the cards form a whole deck again. Together Gilbert and Shannon gave a mathematical description of this shuffle in 1955 and so did Reeds, independently 1981. Henceforth we shall call this shuffle the GSR-shuffle, modeled as follows:

Let M be a binomial$(n, 1/2)$ random variable. Split the deck into the top M cards, say the left packet, and the bottom $n - M$ cards, the right packet. Drop cards one by one either from the left or the right packet, until all cards form a single packet again. A card from the left packet is dropped with probability $A/(A + B)$ and from the right packet with probability $B/(A + B)$, where A and B denote the size of the left, respectively, the right packet before each drop.

As for the top-to-random shuffle we can model a sequence of riffle-shuffles as a Markov chain $X = (X_t)_{t\in\mathbb{N}_0}$ on the group of permutations $\mathcal{S}_n$. Start from $X_0 = \mathrm{id}$ and choose $\sigma \in \mathcal{S}_n$ with probability,

$$\mathbb{P}_{\mathrm{id}}(X_t = \sigma(X_{t-1})) = \begin{cases} (n+1)/2^n & , \text{ if } \sigma = \mathrm{id} \\ 1/2^n & , \text{ if } \sigma \text{ contains exactly two rising sequences} \\ 0 & , \text{ otherwise} \end{cases},$$

where a *rising sequence* is a maximal set of consecutive cards that occur in their original relative order. Dave Bayer and Persi Diaconis [6] showed in 1992, that the GSR-shuffle has a cutoff at $t^n_{mix} = (3/2)\log n$. In particular they proved the following theorem.

Theorem 3.4 *If n cards are shuffled t times with $t = (3/2)\log_2(nc)$, then for large n*

$$d_n(t) = \|\mathbb{P}(X_t \in \cdot) - \pi(\cdot)\|_{TV} = 1 - 2\Phi\left(-\frac{1}{4c\sqrt{3}}\right) + \mathcal{O}\left(\frac{1}{n^{1/4}}\right)$$

with

$$\Phi(x) = \frac{1}{\sqrt{2\pi}}\int_{-\infty}^{x} e^{-t^2/2}dt.$$

The function $1 - 2\Phi\left(-\frac{1}{4c\sqrt{3}}\right)$ has the following asymptotic behavior:

$$1 - 2\Phi\left(-\frac{1}{4c\sqrt{3}}\right) \sim \frac{1}{2c\sqrt{6\pi}} \quad \text{as } c \to \infty,$$
$$1 - 2\Phi\left(-\frac{1}{4c\sqrt{3}}\right) \sim 1 - \frac{4c\sqrt{3}}{\sqrt{2\pi}}\exp\left\{-\frac{1}{2}\left(-\frac{1}{4c\sqrt{3}}\right)^2\right\} \quad \text{as } c \to 0.$$

Since the drop has a width of order 1 around $(3/2)\log_2 n$, the GSR-shuffle has a cutoff.

The proof of Theorem 3.4 is beyond the scope of this thesis, since it is highly technical and uses a fair amount of advanced group theory. We will, however, highlight the most important definitions, lemma's and ideas needed to complete it.

First of all, the GSR-shuffle is a special case of a more general form of a riffle shuffle. In the general case we perform an a-shuffle, which is done as follows. Divide the deck into a packets, packets may be empty, and then interleaving them together. Each card at the bottom of a packet is being dropped with probability A_i/A, where A_i stands for the number of cards in packet $i = 1, \ldots, a$ before a card is being dropped and $A = \sum_{i=1}^{a} A_i$. The following descriptions of the a-shuffle plays an important role:

Geometric description

Place n points, $x_1 < x_2 < \cdots < x_n$, representing the cards, uniformly and independently on the line (0,1). For a positive integer a, the map $a \mapsto ax \pmod 1$ maps $[0,1]$ onto itself and preserves measure. This map rearranges the points x_i, which represents an a-shuffle, and so gives a probability measure on the symmetric group.

Entropy description

All possible ways of cutting the deck into a packets and the subsequent interleaving of the packets are equally likely. Empty packets are allowed.

Inverse description

All possible ways of pulling a deck back apart into a packets are equally likely. Empty packets are allowed. This shuffle is generated by labeling all the cards independently and uniformly with a number $0, 1, \ldots, a-1$, pulling the cards out of the deck into a packets while respecting there relative ordering, and reassembling the deck by putting the packets on top of each other, such that the cards labeled with a 0 come on top, followed by the cards labeled with a 1, etc., until the cards labeled with a $a-1$ make up the last cards.

Sequential description

Choose integers $j_1, j_2, \ldots, j_a$ such that $\sum_{i=1}^{a} j_i = n$ according to the multinomial distribution

$$\mathbb{P}(j_1, j_2, \ldots, j_a) = \binom{n}{j_1 j_2 \ldots j_a} \frac{1}{a^n}.$$

Given j_i, cut off the top j_1 cards, the next j_2 cards and so on, producing a or fewer packets. Shuffle the first two packets together by a GSR-shuffle. Shuffle this combined packet with the next and so on, until we have one packet.

The next lemma links these descriptions together.

Lemma 3.5 *The four descriptions above generate the same permutation distribution. Moreover, in each model an a-shuffle followed by an b-shuffle is equivalent to an ab-shuffle.*

The permutation distribution above is also called the GSR measure.

Theorem 3.6 *The probability that an a-shuffle will result in the permutation σ is*

$$\frac{\binom{a+n-r}{n}}{a^n},$$

where r is the number of rising sequences in σ.

Corollary 3.7 *If a deck of n cards is given a sequence of m shuffles of types $a_1, \ldots, a_m$, then the probability that the deck is in the arrangement σ is given by*

$$\frac{\binom{a+n-r}{n}}{a^n},$$

where $a = a_1 \times a_2 \times \cdots \times a_m$ and r is the number of rising sequences in σ.

Corollary 3.8 *Let a Markov chain on the symmetric group begin at the identity and proceed by successive independent a-shuffles chosen from the GSR measure. Then $R(\sigma)$, the number of rising sequences, forms a Markov chain.*

With these results we can state the two propositions that form the body of the proof of Theorem 3.4. The first one decomposes the GSR-measure into an exponential form, such that we can bound the probability of a permutation σ with r rising sequences by $1/n!$. This is crucial, since it determines on which set the total variation is achieved.

Proposition 3.9 *Let* $Q^m(r) = \binom{2^m+n-r}{n}/2^{mn}$ *be the probability of a permutation with* r *rising sequences after* m *shuffles drawn from the GSR-distribution. Let* $r = n/2 + h$*, with* $-n/2 + 1 \leq h \leq n/2$*. Let* $m = \log_2(n^{3/2}c)$ *with* $0 < c < \infty$ *fixed. Then*

$$Q^m(r) = \frac{1}{n!} \exp\left\{ \frac{1}{c\sqrt{n}} \left(-h + \frac{1}{2} + \mathcal{O}\left(\frac{h}{n}\right)\right) - \frac{1}{24c^2} - \frac{1}{2}\left(\frac{h}{cn}\right)^2 + \mathcal{O}\left(\frac{1}{n}\right)\right\}.$$

Proposition 3.10 *With the same notation as in Proposition 3.9, let* h^* *be an integer such that* $Q^m(n/2 + h) \geq 1/n!$ *if and only if* $h \leq h^*$*. Then, for any fixed* c*, as* $n \to \infty$

$$h^* = -\frac{-\sqrt{n}}{24c^2} + \frac{1}{12c^2} + B + \mathcal{O}\left(\frac{1}{\sqrt{n}}\right),$$

where $-1 \leq B \leq 1$*.*

Via the definition of total variation and the decomposition of the GSR-measure, Proposition 3.9, gives the following equation for the total variation,

$$d_n(t) = \left\|Q^t - \pi\right\|_{TV} = \max_{A \subset \{0,1,\ldots,r\}} \{Q^t(A) - \pi(A)\} = \sum_{-n/2 < h \leq h^*} R_{nh} \left(Q^t\left(\frac{n}{2} + h\right) - \frac{1}{n!}\right),$$

where R_{nh} stands for the number of permutations in $\mathcal{S}_n$ with $h + n/2$ rising sequences. This gives the right bound thanks to Proposition 3.10. Use of the central limit theorem and a standard large deviation bound, combined with a lot of calculus, finally completes the proof of Theorem 3.4.

4 Bounds on total variation

To explain when cutoff happens we will look at the eigenvalues and eigenfunctions (or eigenvectors) of the transition kernel of the Markov chain that we are examining. With these we can find other means to bound the total variation distance and the mixing times. In order to do so, we need the *separation distance* which is defined as

$$s(t) := \max_{x\in\Omega} s_x(t), \quad \text{with} \quad s_x(t) := \max_{y\in\Omega}\left[1 - \frac{\mathbb{P}_x(X_t = y)}{\pi(y)}\right]$$

We have the following Lemma.

Lemma 4.1 *Let $X = (X_t)_{t\in\mathbb{N}_0}$ be an irreducible Markov chain with strong stationary time τ. Then*

$$\|\mathbb{P}_x(X_t \in \cdot) - \pi(\cdot)\|_{TV} \leq s_x(t) \leq \mathbb{P}_x(\tau > t).$$

Proof. To prove the first inequality, we fix $x \in \Omega$ and estimate

$$\begin{aligned}
\|\mathbb{P}_x(X_t \in \cdot) - \pi(\cdot)\|_{TV} &= \sum_{\substack{y\in\Omega \\ \mathbb{P}_x(X_t=y)\leq\pi(y)}} (\pi(y) - \mathbb{P}_x(X_t = y)) \\
&= \sum_{\substack{y\in\Omega \\ \mathbb{P}_x(X_t=y)\leq\pi(y)}} \pi(y)\left(1 - \frac{\mathbb{P}_x(X_t=y)}{\pi(y)}\right) \\
&\leq \max_{y\in\Omega}\left[1 - \frac{\mathbb{P}_x(X_t=y)}{\pi(y)}\right] = s_x(t).
\end{aligned}$$

To prove the second inequality, we fix $x, y \in \Omega$ and make use of the fact that

$$\mathbb{P}_x(X_t = y) = \mathbb{P}_x(X_t = y, \tau > t) + \mathbb{P}_x(X_t = y, \tau \leq t),$$

to get

$$1 - \frac{\mathbb{P}_x(X_t=y)}{\pi(y)} \leq 1 - \frac{\mathbb{P}_x(X_t=y,\tau\leq t)}{\pi(y)} = 1 - \frac{\mathbb{P}_x(\tau\leq t)\pi(y)}{\pi(y)} = \mathbb{P}_x(\tau > t).$$

□

Let $\langle\cdot,\cdot\rangle$ denote the inner product on $\mathbb{R}^\Omega$, which is given by $\langle f, g\rangle = \sum_{x\in\Omega} f(x)g(x)$. Since we will be working with Markov chains with a stationary distribution, we also need a weighted inner product on $\mathbb{R}^\Omega$ given by $\langle f, g\rangle_\pi = \sum_{x\in\Omega} f(x)g(x)\pi(x)$. The use of this weighted inner product will become clear in the proof of the following lemma.

Lemma 4.2 *Given an irreducible and reversible Markov chain X with transition kernel P, the following holds*

(i) *The inner product space $(\mathbb{R}^\Omega, \langle\cdot,\cdot\rangle_\pi)$ has an orthonormal basis of real-valued eigenfunctions $\{f_j\}_{j=1}^n$ corresponding to real eigenvalues $\{\lambda_j\}_{j=1}^n$, with $n = |\Omega|$, ordered as $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_n$.*

(ii) *The matrix P can be decomposed as*

$$\frac{P^t(x,y)}{\pi(y)} = \frac{\mathbb{P}_x(X_t=y)}{\pi(y)} = \sum_{j=1}^n f_j(x)f_j(y)\lambda_j^t.$$

(iii) *The eigenfunction f_1 corresponding to the eigenvalue $\lambda_1 = 1$ can be taken to be the constant vector $\mathbb{1}$, so that*

$$\frac{P^t(x,y)}{\pi(y)} = \frac{\mathbb{P}_x(X_t = y)}{\pi(y)} = 1 + \sum_{j=2}^{|\Omega|} f_j(x) f_j(y) \lambda_j^t. \tag{4.1}$$

Proof.[2] (i): Given a reversible Markov chain X with transition kernel P, recall that

$$\pi(x)P(x,y) = \pi(y)P(y,x), \quad x, y \in \Omega.$$

Let D_π be a diagonal matrix such that $D_\pi(x,x) = \pi(x)$ and $D_\pi(x,y) = 0$ if $x \neq y$, and take $A = D_\pi^{1/2} P D_\pi^{-1/2}$, so that $A(x,y) = \sqrt{\pi(x)}/\sqrt{\pi(y)} P(x,y)$ for all $x, y \in \Omega$. Since

$$\begin{aligned} A(x,y) = \frac{\sqrt{\pi(x)}}{\sqrt{\pi(y)}} P(x,y) &= \frac{\pi(x)}{\sqrt{\pi(x)\pi(y)}} P(x,y) \\ &= \frac{\pi(y)}{\sqrt{\pi(x)\pi(y)}} P(y,x) = \frac{\sqrt{\pi(y)}}{\sqrt{\pi(x)}} P(y,x) = A(y,x), \end{aligned}$$

A is a symmetric matrix and so by the spectral theory for symmetric matrices we know that A has real-valued eigenfunctions, say $\{\varphi_j\}_{j=1}^n$, corresponding to real-valued eigenvalues, say $\{\lambda_j\}_{j=1}^n$, which form an orthonormal basis, say $\mathcal{B}$, for the inner product space $(\mathbb{R}^\Omega, \langle \cdot, \cdot \rangle)$. Take $f_j = D_\pi^{-1/2} \varphi_j$ for $j = 1, \ldots, n$. Then,

$$P f_j = P D_\pi^{-1/2} \varphi_j = D_\pi^{-1/2} A D_\pi^{1/2} D_\pi^{-1/2} \varphi_j = D_\pi^{-1/2} A \varphi_j = \lambda_j D_\pi^{-1/2} \varphi_j = \lambda_j f_j.$$

Let $\delta_{i,j} : \mathcal{B} \times \mathcal{B} \to \{0,1\}$ denote the Dirac delta function, i.e. $\delta_{i,j}(\varphi_i, \varphi_j) = 1$ if and only if $i = j$. Since the original eigenfunctions of A are orthonormal in $(\mathbb{R}^\Omega, \langle \cdot, \cdot \rangle)$, we have

$$\begin{aligned} \delta_{i,j}(\varphi_i, \varphi_j) = \langle \varphi_i, \varphi_j \rangle &= \langle D_\pi^{1/2} f_i, D_\pi^{1/2} f_j \rangle \\ &= \sum_{x \in \Omega} \sqrt{\pi(x)} f_i(x) \sqrt{\pi(x)} f_j(x) \\ &= \sum_{x \in \Omega} f_i(x) f_j(x) \pi(x) = \langle f_i, f_j \rangle_\pi. \end{aligned}$$

Thus the inner product space $(\mathbb{R}^\Omega, \langle \cdot, \cdot \rangle_\pi)$ has a real-valued orthonormal basis $\{f_j\}_{j=1}^n$ with corresponding eigenvalues $\{\lambda_j\}_{j=1}^n$.

(ii). Remark that $\sqrt{\pi}$ is an eigenfunction of A with corresponding eigenvalue $\lambda_1 = 1$, since

$$\begin{aligned} \left(\sqrt{\pi} A\right)(x) &= \sum_{y \in \Omega} \sqrt{\pi(y)} \frac{\sqrt{\pi(y)}}{\sqrt{\pi(x)}} P(y,x) = \sum_{y \in \Omega} \frac{\pi(y)}{\sqrt{\pi(x)}} P(y,x) \\ &= \sum_{y \in \Omega} \sqrt{\pi(x)} P(x,y) = \sqrt{\pi(x)}. \end{aligned}$$

This implies that $D_\pi^{-1/2} \sqrt{\pi} = \mathbb{1}$ is an eigenvector of P. Let δ_y be the following function

$$\delta_y(x) = \begin{cases} 1, & \text{if } x = y, \\ 0, & \text{if } x \neq y. \end{cases}$$

[2] See [7], for the original proof.

In the inner product space, $(\mathbb{R}^\Omega, \langle\cdot,\cdot\rangle_\pi)$ with orthonormal bases $\{f_j\}_{j=1}^n$, we can decompose this function in the following way:

$$\delta_y = \sum_{j=1}^n \langle \delta_y, f_j\rangle_\pi f_j = \sum_{j=1}^n \left[\sum_{x\in\Omega} \delta_y(x) f_j(x)\pi(x)\right] f_j = \sum_{j=1}^n \pi(y) f_j(y) f_j.$$

Using this result, we get

$$P^t(x,y) = \left(P^t\delta_y\right)(x) = \left(\sum_{j=1}^n \pi(y) f_j(y) P^t f_j\right)(x) = \sum_{j=1}^n \pi(y) f_j(y)\lambda_j^t f_j(x),$$

and thus we have that

$$\frac{P^t(x,y)}{\pi(y)} = \sum_{j=1}^n f_j(x) f_j(y)\lambda_j^t.$$

(iii). From the fact that $f_1 = \mathbb{1}$, we get

$$\frac{P^t(x,y)}{\pi(y)} = 1 + \sum_{j=2}^{|\Omega|} f_j(x) f_j(y)\lambda_j^t.$$

□

We need to state one more definition before we can state the next theorem. For $p \geq 0$, the weighted $\ell^p(\pi)$-norm on $\mathbb{R}^\Omega$ is defined as

$$\|f\|_{\pi,p} := \left[\sum_{x\in\Omega} |f(x)|^p \pi(x)\right]^{1/p}.$$

Note that for $p = 2$, $\|f\|_2 = \sqrt{\langle f, f\rangle_\pi}$.

Theorem 4.3 *Let X be a Markov chain with a reversible transition kernel P with eigenvalues $\lambda_1 = 1 \geq \lambda_2 \geq \ldots \lambda_n \geq -1$ corresponding to the eigenfunctions $\{f_j\}_{j=1}^n$ and orthonormal with respect to the inner product $\langle\cdot,\cdot\rangle_\pi$. Then*

$$4\left\|\mathbb{P}_x(X_t \in \cdot) - \pi(\cdot)\right\|_{TV}^2 \leq \left\|\frac{\mathbb{P}_x(X_t\in\cdot)}{\pi(\cdot)} - 1\right\|_2^2 = \sum_{j=2}^{|\Omega|} f_j(x)^2\lambda_j^{2t}.$$

Furthermore if the chain X is also transitive then we have

$$4\left\|\mathbb{P}_x(X_t \in \cdot) - \pi(\cdot)\right\|_{TV}^2 \leq \sum_{j=2}^{|\Omega|} \lambda_j^{2t}.$$

Proof: Using Lemma 4.2.(iii) we get

$$\begin{aligned}\left\|\frac{\mathbb{P}_x(X_t\in\cdot)}{\pi(\cdot)} - 1\right\|_2^2 &= \left\|\sum_{j=2}^{|\Omega|} \lambda_j^t f_j(x) f_j\right\|_2^2 \\ &= \langle \sum_{j=2}^{|\Omega|} \lambda_j^t f_j(x) f_j, \sum_{j=2}^{|\Omega|} \lambda_j^t f_j(x) f_j\rangle_\pi \\ &= \sum_{j=2}^{|\Omega|} f_j(x)^2\lambda_j^{2t}. \end{aligned} \tag{4.2}$$

Furthermore, using the fact that $\|\mu - \nu\|_{TV} = \frac{1}{2}\sum_{x\in\Omega} |\mu(x) - \nu(x)|$ we get

$$\begin{aligned}\|\mathbb{P}_x(X_t \in \cdot) - \pi(\cdot)\|_{TV} &= \frac{1}{2}\sum_{y\in\Omega}\left|\frac{\mathbb{P}_x(X_t = y)}{\pi(y)} - 1\right|\pi(y) \\ &= \frac{1}{2}\left\|\frac{\mathbb{P}_x(X_t \in \cdot)}{\pi(\cdot)} - 1\right\|_1.\end{aligned}$$

Since $p \mapsto \|f\|_p$ is non-decreasing, we get

$$\begin{aligned}4\,\|\mathbb{P}_x(X_t \in \cdot) - \pi(\cdot)\|_{TV}^2 &= \left\|\frac{\mathbb{P}_x(X_t \in \cdot)}{\pi(\cdot)} - 1\right\|_1^2 \\ &\leq \left\|\frac{\mathbb{P}_x(X_t \in \cdot)}{\pi(\cdot)} - 1\right\|_2^2\end{aligned}$$

Combine both estimates to give the first bound.

If the chain is also transitive, then for each pair $(x, y) \in \Omega^2$ there exists a bijection $\varphi_{(x,y)} : \Omega \mapsto \Omega$, such that $\varphi(x) = y$ and $P(z, w) = P(\varphi(z), \varphi(w))$ for all $z, w \in \Omega$. Since the chain is also reversible we have that π is uniform on Ω. From this it follows that the left-hand site of (4.2) does not depend on x. So for any $x_0 \in \Omega$ we have

$$\left\|\frac{\mathbb{P}_{x_0}(X_t \in \cdot)}{\pi(\cdot)} - 1\right\|_2^2 = \sum_{j=2}^{|\Omega|} f_j(x)^2\lambda_j^{2t}.$$

Summing over x on both sites then gives

$$\begin{aligned}|\Omega|\left\|\frac{\mathbb{P}_{x_0}(X_t \in \cdot)}{\pi(\cdot)} - 1\right\|_2^2 &= |\Omega|\sum_{x\in\Omega}\left(\pi(x)\sum_{j=2}^{|\Omega|} f_j(x)^2\lambda_j^{2t}\right) \\ &= |\Omega|\sum_{j=2}^{|\Omega|}\left(\sum_{x\in\Omega} f_j(x)^2\pi(x)\right)\lambda_j^{2t} = |\Omega|\sum_{j=2}^{|\Omega|}\lambda_j^{2t}.\end{aligned}$$

Dividing both sites with $|\Omega|$ gives the second result

□

4.1 Example without cutoff

Let us look at the lazy random walk $X = (X_t)_{t\in\mathbb{N}_0}$ on the integers modulo n, starting at $X_0 = 0$. Where at each unit of time the walker can take one step up, one step down or stand still with equal probability, this yields the following transition kernel

$$P(i, j) = \begin{cases} 1/3 & \text{if } j = i + 1 \mod n, \\ 1/3 & \text{if } j = i - 1 \mod n, \\ 1/3 & \text{if } j = i, \\ 0 & \text{else,} \end{cases} \qquad \forall i, j \in \mathbb{Z}_n.$$

This random walk can be viewed as a lazy random walk on the n-element cyclic group. It is immediate that this random walk is irreducible and reversible and so the stationary distribution is uniform. We will use this representation to find the eigenvalues and eigenfunctions of the transition kernel. With those we can apply Theorem 4.3 to find an upper bound on the total variation time. A lower bound will be given by comparing this walk with a lazy random walk on the integers. From these bounds it will be clear that the lazy random walk on the integers modulo n does not have a cutoff.

Lemma 4.4 *For the lazy random walk $X = (X_t)_{t\in\mathbb{N}_o}$ on the integers modulo n, we have the following upper bound on the total variation distance*

$$d\left(\frac{3}{4\pi^2}cn^2\right) \leq e^{-c} \quad c > 0.$$

Proof. Let $\omega = e^{2\pi i/n}$ and let $W_n = \{\omega, \omega^2, \ldots, \omega^{n-1}, 1\}$ be the set of n-th roots of unity which form a n-gon inscribed in the unit circle. Remark that $(W_n, \cdot)$, with $\cdot$ denoting the multiplication, is the n-cyclic group since

$$\omega^j \cdot \omega^k = \omega^{j+k} = \omega^{j+k \mod n},$$

and the group is generated by ω. Now remark that for all eigenfunctions f with corresponding eigenvalue λ the following must hold

$$\lambda f(\omega^k) = Pf(\omega^k) = \frac{f(\omega^{k-1}) + f(\omega^k) + f(\omega^{k+1})}{3}.$$

Fix k and for $j = 0, 1, \ldots, n-1$ define $f_j(\omega^k) = \omega^{jk}$. Then

$$\begin{aligned}
\lambda_j f_j(\omega^k) &= \frac{f_j(\omega^{k-1}) + f_j(\omega^k) + f_j(\omega^{k+1})}{3} \\
&= \frac{\omega^{jk-j} + \omega^{jk} + \omega^{jk+j}}{3} \\
&= \omega^{jk}\frac{1 + \omega^j + \omega^{-j}}{3} \\
&= f_j(\omega^k)\frac{1 + \omega^j + \omega^{-j}}{3}.
\end{aligned}$$

So for $j = 0, 1, \ldots, n-1$, f_j is an eigenfunction for P with corresponding eigenvalue

$$\lambda_j = \frac{1}{3}\left(1 + \omega^j + \omega^{-j}\right) = \frac{1}{3}\left(1 + 2\cos(2\pi j/n)\right).$$

Note that $\lambda_1 > 1 - 4\pi^2/3n^2 \geq \lambda_i$, with $i = 2, \ldots, n-1$ and that f_j is bounded for all j. By Theorem 4.3 we therefor get

$$\|\mathbb{P}_x(X_t \in \cdot) - \pi(\cdot)\|_{TV} \leq \frac{1}{2}\sqrt{\sum_{j=1}^{n-1} f_j(x)^2\lambda_j^{2t}} < \frac{1}{2}\lambda_1^t \underbrace{\sqrt{\sum_{j=1}^{n-1} f_j(x)^2}}_{\leq 1} \leq \frac{1}{2}\lambda_1^t.$$

Looking only at the lead term and letting n grow, as Diaconis suggest in [4], we get

$$\|\mathbb{P}_x(X_t \in \cdot) - \pi(\cdot)\|_{TV} \leq \frac{1}{2}\left(1 - \frac{4\pi^2}{3n^2}\right)^t \leq \frac{1}{2}e^{-\frac{4\pi^2}{3n^2}t}.$$

If $t = \frac{3}{4\pi^2}cn^2$ then this tends to e^{-c}.

□

So we can conclude that the mixing time, t_{mix}, for the lazy random walk on the integers modulo n is of the order n^2. The next lemma shows that the the lower bound of the total variation distance is also of the order n^2.

Lemma 4.5 *For the lazy random walk $X = (X_t)_{t\in\mathbb{N}_o}$ on the integers modulo n, there exist a $\alpha_0 > 0$, so that for the total variation distance the following holds*

$$d\left(\alpha n^2\right) \geq 1/4 \qquad \forall \alpha < \alpha_0.$$

Proof. To prove this statement we shall couple $X = (X_t)_{t\in\mathbb{N}_0}$ to $S = (S_t)_{t\in\mathbb{N}_0}$, the same lazy random walk on the integers without the modulo count, starting from $S_0 = 0$ until time τ, with

$$\tau := \min\{t : |S_t| = \lfloor n/2 \rfloor\}.$$

For S_t we have

$$\mathbb{E}(S_t) = 0 \qquad \mathrm{Var}(S_t) = \frac{2}{3}t. \tag{4.3}$$

Using (4.3) and the Chebyshev's inequality we find

$$\begin{aligned}\mathbb{P}\left(\sup_{t\leq\alpha n^2} |X_t| > n/4\right) &= \mathbb{P}\left(\sup_{t\leq\alpha n^2} |S_t| > n/4\right) \leq \mathbb{P}(|S_t| > n/4) \\ &= \mathbb{P}(|S_{\alpha n^2} - \mathbb{E}(S_{\alpha n^2})| > n/4 - \mathbb{E}(S_{\alpha n^2})) \\ &\leq 16\frac{\mathrm{Var}(S_{\alpha n^2})}{n^2} = \frac{32\alpha}{3}.\end{aligned}$$

For $\alpha \leq \alpha_0 = 3/128$, we have $\mathbb{P}(\sup_{t\leq\alpha n^2} |X_t| > n/2) \leq 1/4$. If we now take $A := \{k : |k| \geq n/4\}$, i.e. the set of integers between $-n/2$ and $n/2$, such that the absolute value is greater or equal to $n/2$, then $\pi(A) \geq 1/2$. So

$$\begin{aligned}\|\mathbb{P}(X_{\alpha n^2} \in \cdot) - \pi(\cdot)\|_{TV} &\geq |\pi(A) - \mathbb{P}(X_{\alpha n^2} \in A)| \\ &\geq \pi(A) - \mathbb{P}(X_{\alpha n^2} > n/4) \geq \frac{1}{4}.\end{aligned}$$

□

Intuitively it is clear that the lazy random walk on the integers modulo n does not have a cutoff phenomena. This also follows directly from corollary 2.4, since the requirements do not hold for all c as we have shown in the previous lemma's.

4.2 Example with cutoff

Let us look at the Ehrenfest Urn introduced by P. and T. Ehrenfest. In this model we look at two urns and n balls. Initially all balls are in urn 2. At each unit of time one of the n balls is chosen uniformly, via a Bernoulli distribution with $p = 1/2$, the ball is placed in either urn 1 or in urn 2. It is intuitively clear that after some amount of time any ball is equally likely to be in any of the two urns. Now, the state of this Markov chain is determined by the number of balls in urn 1. So we get a lazy random walk on the set $\{0, 1, \ldots, n\}$ with the following transition probabilities:

$$P(i,j) = \begin{cases} \frac{1}{2}\frac{n-i}{n} & \text{if } j = i+1, \\ \frac{1}{2} & \text{if } j = i, \\ \frac{1}{2}\frac{i}{n} & \text{if } j = i-1, \end{cases} \qquad \text{for } i \in \{0, 1 \ldots, n\}.$$

The stationary distribution of this model is the binomial distribution with parameters n and $p = 1/2$. To evaluate the eigenfunctions and eigenvalues we will look at the lazy random walk on the n-dimensional hypercube. This is a random walk $X = (X_t)_{t\in\mathbb{N}_0}$ on the set $\Omega = \mathbb{Z}_2^n = \{0,1\}^n$ starting from the point $X_0 = (0, 0, \ldots, 0)$, with increments $(0, 0, \ldots, 0)$, $e_1 = (1, 0, \ldots, 0), \ldots, e_n(0, 0, \ldots, 1)$, such that

$$\mathbb{P}(X_{t+1} = X_t) = \frac{1}{2} \quad \text{and} \quad \mathbb{P}(X_{t+1} = X_t + e_i) = \frac{1}{2n} \quad \text{for all } i \in \{1, \ldots, n\}.$$

Let $W(\mathbf{x}) = \sum_{i=1}^n x_i$ denote the Hamming weight of $\mathbf{x} = (x_1, x_2, \ldots, x_n) \in \mathbb{Z}_2^n$. Then it is not difficult to see that the Ehrenfest random walk is a projection of the lazy random walk on the n-dimensional hypercube. First we will give the lower bound on the total variation distance, then the upper bound and finally we will show that the total variation of the Ehrenfest Urn walk and the lazy random walk on the n dimensional hypercube are indeed the same.

Proposition 4.6 *Let μ and ν be two probability distributions on Ω, and let f be a real-valued function on Ω. If*

$$|\mathbb{E}_\mu(f) - \mathbb{E}_\nu(f)| \geq r\sigma, \tag{4.4}$$

where $\sigma^2 = (\mathrm{Var}_\mu(f) + \mathrm{Var}_\nu(f))/2$, then

$$\|\mu - \nu\|_{TV} \geq 1 - \frac{4}{4 + r^2}.$$

Lemma 4.7 *Consider the coupon collecting problem with n distinct coupon types, and let $I_j(t)$ be the indicator of the event that the j-th coupon has not been collected by time t. Let $R_t = \sum_j^n I_j(t)$ be the number of coupon types* not *collected by time t. The random variables $I_j(t)$ are negatively correlated, with $p = (1 - 1/n)^t$ for $t \geq 0$*

$$\mathbb{E}(R_t) = np, \qquad \mathrm{Var}(R_t) \leq np(1-p).$$

With these we can state the lower bound. [3]

Proposition 4.8 *For the lazy random walk on the n-dimensional hypercube*

$$d\left(\frac{1}{2}n\log n - \alpha n\right) \geq 1 - 4e^{-2\alpha+1}.$$

Proof. Let X be as defined above and let W be the Hamming weight function. We will use Proposition 4.6 to bound the total variation of the lazy random walk on the n-dimensional hypercube from below. Since X is a reversible and transitive random walk, we know that π is uniform on $\mathbb{Z}_2^n$. Under π the random variable $W := W(X)$ is binomial with parameters n and $p = 1/2$ and thus

$$\mathbb{E}_\pi(W_t) = \frac{n}{2} \qquad \mathrm{Var}_\pi(W_t) = \frac{n}{4}. \tag{4.5}$$

Now let R_t be the number of coordinates that where not updated at time t. When starting from $\mathbf{0}$ the conditional distribution of W_t given R_t is binomial with parameters $n - R_t$ and $p = 1/2$. So

$$\mathbb{E}_{\mathbf{0}}(W_t|R_t) = \frac{n - R_t}{2} \qquad \mathrm{Var}_{\mathbf{0}}(W_t|R_t) = \frac{n - R_t}{4},$$

and using Lemma 4.7 we get

$$\mathbb{E}_{\mathbf{0}}(W_t) = \mathbb{E}\left(\mathbb{E}_{\mathbf{0}}(W|R_t)\right) = \frac{n - \mathbb{E}_{\mathbf{0}}(R_t)}{2} = \frac{n}{2}\left(1 - \left(1 - \frac{1}{n}\right)^t\right). \tag{4.6}$$

Using the identity

$$\mathrm{Var}(Y) = \mathrm{Var}(\mathbb{E}(Y|Z)) + \mathbb{E}(\mathrm{Var}(Y|Z)),$$

we get

$$\begin{aligned}
\mathrm{Var}_{\mathbf{0}}(W_t) &= \mathrm{Var}_{\mathbf{0}}(\mathbb{E}_{\mathbf{0}}(W_t|R_t)) + \mathbb{E}_{\mathbf{0}}(\mathrm{Var}_{\mathbf{0}}(W_t|R_t)) \\
&= \mathrm{Var}_{\mathbf{0}}\left(\frac{n - R_t}{2}\right) + \mathbb{E}_{\mathbf{0}}\left(\frac{n - R_t}{4}\right) \\
&= \frac{n}{4} + \frac{1}{4}\left(\mathrm{Var}_{\mathbf{0}}(R_t) - \mathbb{E}_{\mathbf{0}}(R_t)\right) \leq \frac{n}{4},
\end{aligned}$$

[3] The proofs of proposition 4.6 and lemma 4.7 can be found in [7]

where the last inequality follows from the fact that $I_j(t)$ are negatively related so that $\mathbb{E}(R_t) \geq \mathrm{Var}(R_t)$. Take

$$\sigma = \sqrt{\max\{\mathrm{Var}_\pi(W_t), \mathrm{Var}_{\mathbf{0}}(W_t)\}} = \frac{\sqrt{n}}{2}.$$

Then, by (4.5) and (4.6), we get

$$\begin{aligned}
|\mathbb{E}_{\mathbf{0}}(W_t) - \mathbb{E}_\pi(W_t)| &= \left| \frac{n}{2}\left(1 - \left(1 - \frac{1}{n}\right)^t\right) - \frac{n}{2} \right| \\
&= \frac{n}{2}\left(1 - \frac{1}{n}\right)^t = \sigma\sqrt{n}\left(1 - \frac{1}{n}\right)^t \\
&= \sigma e^{t\log\left(1-\frac{1}{n}\right) + \frac{1}{2}\log n}.
\end{aligned}$$

Using the fact that $\log(1-x) \leq -x - x^2$ for $0 \leq x \leq 1/2$, we get

$$\geq \sigma e^{-\frac{t}{n}\left(\frac{n+1}{n}\right) + \frac{1}{2}\log n}.$$

Applying Proposition 4.6 with $r = e^{-\frac{t}{n}\left(\frac{n+1}{n}\right) + \frac{1}{2}\log n}$, we obtain

$$\begin{aligned}
\|\mathbb{P}_{\mathbf{0}}(X_t \in \cdot) - \pi(\cdot)\|_{TV} &\geq 1 - \frac{4}{4 + e^{-\frac{2t}{n}\left(\frac{n+1}{n}\right) + \frac{1}{2}\log n}} \\
&\geq 1 - 4e^{\frac{2t}{n}\left(\frac{n+1}{n}\right) - \log n}.
\end{aligned}$$

For $\alpha \leq n\log n$ we have

$$(n+1)\left(\frac{1}{2}n\log n - \alpha n\right) \leq n\left(\frac{1}{2}n\log n - \alpha n\right) + \frac{1}{2}n^2.$$

If we take

$$t_n = \frac{n}{n+1}\left(\frac{1}{2}n\log n - \left(\alpha - \frac{1}{2}\right)n\right),$$

then

$$t_n > \frac{1}{2}n\log n - \alpha n.$$

So

$$d(t_n) = \|\mathbb{P}_{\mathbf{0}}(X_{t_n} \in \cdot) - \pi(\cdot)\|_{TV} \geq 1 - 4e^{-2\alpha+1}.$$

□

To prove the upper bound we will look at the eigenvalues and eigenfunctions of the random walk on the n-dimensional hypercube. For this purpose we will look at the product of n irreducible Markov chains. For $j = 1, 2, \ldots, n$, let P_j be an irreducible transition matrix on the state space Ω_j, with stationary distribution π_j. In the next lemma we look at the product Markov chain on $\widetilde{\Omega} := \Omega_1 \times \Omega_2 \times \cdots \times \Omega_n$. The chain selects a coordinate i according to some distribution w on $\{1, 2, \ldots, n\}$ and then only moves according to P_i on the coordinate i. Let $\mathbf{x} = (x_1, x_2, \ldots, x_n) \in \widetilde{\Omega}$. Then the transition matrix $\widetilde{P}$ of this chain is defined as

$$\widetilde{P}(\mathbf{x}, \mathbf{y}) := \sum_{j=1}^{n} P_j(x_j, y_j) \prod_{\substack{i=1 \\ i \neq j}}^{n} \mathbb{1}_{\{x_i = y_i\}}.$$

Lemma 4.9 *For $j = 1, 2, \ldots, n$, let P_j be an irreducible transition matrix on the state space Ω_j, with stationary distribution π_j. Let $\widetilde{P}$ be the transition kernel on $\widetilde{\Omega}$ as defined above. Then*

(i) *The function* $\varphi := \varphi^{(1)} \otimes \ldots \otimes \varphi^{(n)}$ *is an eigenfunction of the transition matrix* $\widetilde{P}$*, with eigenvalue* $\sum_{j=1}^{n} w_j \lambda^{(j)}$*.*

(ii) *Suppose that for each* j*,* $\mathcal{B}_j$ *is an orthogonal basis in* $\ell^2(\pi_j)$*. Then the collection*

$$\widetilde{\mathcal{B}} := \left\{ \bigotimes_{j=1}^{n} \varphi^{(j)} \ : \ \varphi^{(j)} \in \mathcal{B}_j; \ j = 1, \ldots, n \right\},$$

is a basis for $\ell^2(\pi_1 \times \cdots \times \pi_n)$*.*

We will now start to look at the lazy random walk on the n-dimensional hypercube and will use Lemma 4.9 to determine its eigenfunctions and its eigenvalues. For every $j = 1, 2, \ldots, n$ we have the transition kernel

$$P_j = \frac{1}{2} \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix}.$$

Straightforward computations show that P_j has eigenvalues $\lambda_0^{(j)} = 1$ and $\lambda_1^{(j)} = 0$ with corresponding eigenfunctions $\varphi_0^{(j)} = (1/\sqrt{2})(1,1)^t$ and $\varphi_1^{(j)} = (1/\sqrt{2})(1,-1)^t$. By Lemma 4.9 it is clear that the lazy random walk on the n-dimensional hypercube has 2^n eigenvalues and eigenfunctions. Which can be labeled by $J \in \mathbb{Z}_2^n$ as follows

$$\lambda_J = \sum_{i=1}^{n} \pi(i) \lambda_{J_i}^{(i)} = \frac{1}{n} \sum_{i=1}^{n} \mathbb{1}_{\{J_i = 0\}} = \frac{n - W(J)}{n} \quad \text{and} \quad \varphi_J = \bigotimes_{i=1}^{n} \varphi_{J_i}^{(i)}.$$

Note that the eigenfunctions φ_j correspond to the Jth columns of the Hadamard matrix H_{2^n} and thus form an orthonormal basis for $\widetilde{P}$.

Proposition 4.10 *For the lazy random walk on the n-dimensional hypercube*

$$d\left(\frac{1}{2} n \log n + \alpha n\right) \leq \frac{1}{\sqrt{2}} e^{-\alpha}.$$

Proof. Remark that the lazy random walk on the n-dimensional hypercube is transitive. Using Theorem 4.3 with the eigenvalues we have found above, we therefor have

$$\begin{aligned} 4 \left\| \mathbb{P}_0(X_t \in \cdot) - \pi(\cdot) \right\|_{TV}^2 &\leq \sum_{J \in \mathbb{Z}_2^n \setminus \{\mathbf{0}\}} \lambda_J^{2t} = \sum_{j=1}^{n} \left(1 - \frac{j}{n}\right)^{2t} \binom{n}{j} \leq \sum_{j=1}^{n} e^{-2jt/n} \binom{n}{j} \\ &= (1 - e^{-2t/n})^n - 1. \end{aligned}$$

Taking $t = \frac{1}{2} n \log n + \alpha n$, we get

$$4 \left\| \mathbb{P}_0(X_t \in \cdot) - \pi(\cdot) \right\|_{TV}^2 \leq \left(1 - \frac{1}{n} e^{-2\alpha}\right) - 1 \leq e^{e^{-2\alpha}} - 1 \leq 2e^{-2\alpha}.$$

Dividing by 4 and taking the square on both sides, we get

$$\left\| \mathbb{P}_0(X_t \in \cdot) - \pi(\cdot) \right\|_{TV} \leq \frac{1}{\sqrt{2}} e^{-\alpha}.$$

□

Now it remains to prove that the total variation of the Ehrenfest Urn is the same as the total variation of the lazy random walk on the n-dimensional hypercube.

Theorem 4.11 *Let $(X_t)_{t\in\mathbb{N}_0}$ be the lazy random walk on the n-dimensional hypercube with stationary distribution π and let $(W_t)_{t\in\mathbb{N}_0}$ be the Ehrenfest urn walk with stationary distribution π_W. Then*

$$\|X_t - \pi\|_{TV} = \|W_t - \pi_W\|_{TV}\,.$$

Proof. Let W be the Hamming weight function as defined above and for $w \in \{1,\dots,n\}$ let $\Omega_w := \{x \in \mathbb{Z}_2^n : W(x) = w\}$ be the set of all $x \in \mathbb{Z}_2^n$ with w bits set to one. Remark that for all $x, y \in \Omega_w$, $\mathbb{P}_0(X_t = x) = \mathbb{P}_0(X_t = y)$ and $\pi(x) = \pi(y)$, so that

$$\sum_{x\in\Omega_w} |\mathbb{P}_0(X_t = x) - \pi(x)| = \left|\sum_{x\in\Omega_w} \mathbb{P}_0(X_t = x) - \pi(x)\right| = |\mathbb{P}_0(W_t = w) - \pi_W(w)|\,.$$

Summing over $w \in \{0,1,\dots,n\}$ and dividing by 2, we obtain

$$\|X_t - \pi\|_{TV} = \|W_t - \pi_W\|_{TV}\,.$$

□

So the Ehrenfest Urn has a cutoff at $\frac{1}{2}n\log n$ with of a size of factor n. We have proven this via the eigenvalues of the lazy random walk on the n-dimensional hypercube.

Part II
Metastability

5 1-Dimensional Example of Metastability

In this part of the thesis we will explain what metastability is by looking at some non-equilibrium phenomena studied in physics. Experiments done to study these phenomena are theoretically backed up by calculations done on simplified mathematical models. As a warmup we will look at a 1-dimensional Random Walk (1dRW) that shows metastable behavior. Theorems for this simple example carry over to more complex higher-dimensional models. Though the latter are a lot harder to handle, the basic ideas are the same.

Consider a thermodynamic system in equilibrium at a point P in phase space, which is typically determined by parameters such as temperature, pressure, density and external fields. In equilibrium, the system is stable and has a low energy level. Due to external influences, such as a change in the phase-space parameters, the parameters at point P can shift to a point $P' \neq P$. If the change is gradual and sufficient energy is put into the system, it will gradually move towards its new equilibrium P'. However if the change is (almost) instantaneous and a small amount of energy is put into the system, then the system stays in the old equilibrium P for a long time, which we now call a metastable state, before it moves to the new equilibrium P' under the influence of random fluctuations. Since the system has to overcome some energy barrier when making this transition, it takes a long time for the random fluctuations to achieve the crossover.

The transition only occurs after the system creates a sufficiently large droplet of the new phase inside the old phase, called the "critical droplet". This droplet is created by the system in a manner that requires the minimal amount of energy possible, which is in contrast with an immediate transition that requires high amounts of energy.

This phenomenon occurs in the following examples:

(1) a wrongly magnetized ferromagnet,

(2) a supersaturated gas.

The ferromagnet will be discussed in section 6, the supersaturated gas, in section 7. Now we will first discuss the toy-example of a 1dRW.

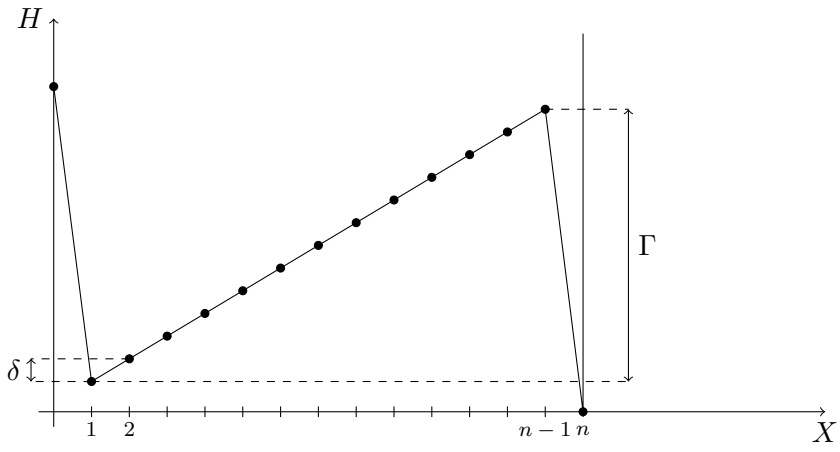

Figure 1: Hamiltonian for a 1d-RW on a double well.

Let $X = (X_t)_{t\in\mathbb{N}}$ be a random walk on the state space $\Omega = \{0, 1, \ldots, n\}$ that is reversible with respect to the Gibbs measure μ, given by

$$\mu(x) = \frac{e^{-\beta H(x)}}{Z}, \quad x \in \Omega,$$

where $Z = \sum_{x\in\Omega} e^{-\beta H(x)}$ and where H is the Hamiltonian for the double-well model as shown in figure 1. The Hamiltonian gives the energy level for each element $x \in \Omega$ for the double-well

model. The following transition probabilities for the random walk are given by the Metropolis rule

$$P(x,y) = \begin{cases} \frac{1}{2}e^{-\beta[H(y)-H(x)]_+} & \text{if } |x-y|=1, \\ 0 & \text{if } |x-y|>1, \\ 1-\sum_{y:|x-y|=1}\frac{1}{2}e^{-\beta[H(y)-H(x)]_+} & \text{if } x=y, \end{cases}$$

where we put $H(-1) = H(n+1) = \infty$, to ensure that the random walk stays in Ω, and $[f(x)]_+ = \max\{f(x), 0\}$. We can think of β as the inverse of the temperature. If the temperature is low, i.e., near the absolute 0-temperature, then β is large and transitions against the drift are unlikely. On the other hand, at high temperatures transitions become more likely as we expect an easy motion.

If we take a closer look at the measure μ, then we see that it is concentrated on the state of minimal energy, state n. After a very long time the Markov chain will be in this state with a probability converging to one as β tends to infinity. The other states have a weight that is exponentially smaller in β, of which the state 1 has the largest weight. This state is the local minimum and will play the metastable part is this example.

As $\beta \to \infty$, the 1d-RW closely resembles the 1d-RW at zero temperature, with absorbing states 1 and n. Let $X^0 = (X_t^0)_{t\in\mathbb{N}}$ denote this random walk, which has the following transition probabilities:

$$\begin{aligned} &P^{(0)}(0,1) = P^{(0)}(1,1) = P^{(0)}(n,n) = 1, \\ &P^{(0)}(1,n-1) = P^{(0)}(1,0) = P^{(0)}(1,2) = 0, \\ &P^{(0)}(i,i-1) = P^{(0)}(i,i) = 1/2, \quad \text{for } i = 2,\dots,n-2, \\ &P^{(0)}(n-1,n-2) = P^{(0)}(n-1,n) = 1/2, \end{aligned}$$

where $P^{(0)}(x,y)$ stands for the transition probability from x to y at zero temperature. The state $n-1$ obviously represents a saddle point, and the states 1 and n are absorbing. When we run X^0, it has a drift towards 1 at any state $x \in [1, n-1]$ and a drift to the right at state 0, towards 1 and at state $n-1$ towards n. This can be used to give certain lower bounds for hitting times needed later on.

The following theorem gives some insight on the behavior of the 1d-RW.

Theorem 5.1 (1) *Uniformly in $x \in \Omega$,*

$$\mathbb{P}_x(\tau_{\{1,n\}} > e^{\varepsilon\beta}) = SES, \quad \forall \varepsilon > 0,$$

where SES means super exponentially small *in β. Moreover, for any $x < n-1$,*

$$\lim_{\beta\to\infty} \mathbb{P}_x(\tau_1 < \tau_n) = 1.$$

(2) *If $\Gamma = H(n-1) - H(1)$ uniformly in $x < n-1$,*

$$\lim_{\beta\to\infty} \mathbb{P}_x(e^{\beta(\Gamma-\varepsilon)} < \tau_n < e^{\beta(\Gamma+\varepsilon)}) = 1, \quad \forall \varepsilon > 0.$$

(3) *For any $x < n-1$,*

$$\lim_{\beta\to\infty} \frac{1}{\beta} \ln \mathbb{E}_x(\tau_n) = \Gamma.$$

(4) *Let $\tau_{\{0,n-1\}}$ be the first exit time from the interval $[1, n-2]$. Then for any x in this interval,*

$$\mathbb{P}_x(X_{\tau_{\{0,n-1\}}} = 0) \le e^{\beta[H(0)-H(n-1)-\varepsilon]}, \quad \forall \varepsilon > 0.$$

(5) *For any* $x < n-1$,

$$\lim_{\beta\to\infty} \mathbb{P}_x(\tau_n > t\mathbb{E}_x(\tau_n)) = e^{-t}.$$

Proof: For the proofs of (1)-(5) we need the following abbreviation: $\delta := H(2) - H(1)$.

(1) Let $\{\phi_t(x)\}_{t\in\mathbb{N}}$ be a sequence given by

$$\begin{aligned}
&\phi_t(x) = \max\{x-t, 1\} && \text{if } x = 2, \ldots, n-1,\\
&\phi_t(x) = x && \text{if } x = 1, n,\\
&\phi_t(0) = \min\{t, 1\}.
\end{aligned}$$

Note that this sequence follows a decreasing path, with respect to the energy, as in the 0-temperature walk. The transitions $x \to x$ are taken out and the path moves from $n-1$ to $n-2$ instead of also having probability $1/2$ to move to n. Now, consider the set D_n, the set of all decreasing paths (with respect to the energy) starting from any $x \in \Omega$ and following the decreasing sequence $\{\phi_t(x)\}_{t\in\mathbb{N}}$ for n time steps and hence reaching the absorbing states 1 or n, i.e.,

$$D_n := \{X_t = \phi_t(X_0) \text{ for } t = 0, 1, \ldots, n\}.$$

Remark that

$$\mathbb{P}_x(\{X_0, X_1, \ldots, X_n\} \in D_n) \geq 2^{-n},$$

since the transition probabilities $P(i, i-1)$, $i = 2, \ldots, n-1$, $P(1,1)$, $P(0,1)$ and $P(n,n)$ are all $\geq 1/2$.

The probability that we need more then $e^{\varepsilon\beta}$ steps to reach either state 1 or n can be bounded from above by the probability that in $e^{\varepsilon\beta}/n$ intervals of length n each interval is not an element of D_n. So

$$\begin{aligned}
\mathbb{P}_x(\tau_{\{1,n\}} > e^{\varepsilon\beta}) &\leq (1 - \mathbb{P}_x(\{X_0, X_1, \ldots, X_n\} \in D_n))^{\frac{e^{\varepsilon\beta}}{n}}\\
&\leq (1 - 2^{-n})^{\frac{e^{\varepsilon\beta}}{n}} \leq \exp\left[-2^{-n}\frac{e^{\delta\beta}}{n}\right] = SES.
\end{aligned}$$

For the second statement, consider a gambler who at each bet makes by one unit a bet and increases his fortune by one unit with probability p and decreases his fortune with probability $q = 1-p$. The gambler starts with a fortune k and he stops betting when he is either broke or reaches a fixed fortune N. We can view this as a random walk $S = (S_t)_{t\in\mathbb{N}}$ on $\mathbb{Z}$, with increments ± 1. We have that $\mathbb{P}(S_{t+1} - S_t = 1 | S_0, \ldots, S_t) = p$ and $\mathbb{P}(S_{t+1} - S_t = -1 | S_0, \ldots, S_t) = 1 - p = q$. We will use martingales to compute the following probability

$$\mathbb{P}_k(\tau_0 < \tau_n), \quad k = 0, \ldots, N,$$

i.e., is the probability that the gambler goes bankrupt before attaining a fortune of N units. Let $M = (M_t)_{t\in\mathbb{N}}$, with $M_t = (q/p)^{S_t}$. M is clearly adapted to the natural filtration $\{\mathcal{F}_t\}_{t\in\mathbb{N}}$, with $\mathcal{F}_t = \sigma(S_0, \ldots, S_t)$. Furthermore $\mathbb{E}(|M_t|) < \infty$ for all $t \in \mathbb{N}$, and we have

$$\begin{aligned}
\mathbb{E}(M_{t+1} | S_0, \ldots, S_t) &= \mathbb{E}((q/p)^{S_{t+1}} | S_0, \ldots, S_t)\\
&= \mathbb{E}((q/p)^{S_t} (q/p)^{S_{t+1} - S_t} | S_0, \ldots, S_t)\\
&= (q/p)^{S_t} (p\,(q/p) + q\,(p/q))\\
&= (q/p)^{S_t} = M_t,
\end{aligned}$$

so that M is indeed a martingale. Now define $\tau = \tau_0 \wedge \tau_N$. Since $\mathbb{E}(M_\tau) < \infty$, we can use the optional stopping theorem, so that $M_{\tau\wedge t}$ is a martingale and $\mathbb{E}(M_\tau) = \mathbb{E}(M_0)$. We get

$$\mathbb{E}_k(M_\tau) = \mathbb{E}_k((q/p)^{S_\tau}) = (q/p)^k.$$

On the other hand we have

$$\mathbb{E}_k(M_\tau) = \mathbb{P}_k(\tau_0 < \tau_N) + (1 - \mathbb{P}_k(\tau_0 < \tau_N))\,(q/p)^N\,.$$

Combining these two results we find that

$$\mathbb{P}_k(\tau_0 < \tau_N) = \frac{(q/p)^N - (q/p)^k}{(q/p)^N - 1}. \tag{5.1}$$

For our original problem we had the transition probabilities, for $i \in 2, \ldots, n-2$,

$$P(i, i-1) = \frac{1}{2}, \quad P(i,i) = \frac{1}{2} - \frac{1}{2}e^{-\delta\beta} \quad \text{and} \quad P(i, i+1) = \frac{1}{2}e^{-\delta\beta}.$$

Conditioning on the event $\{X_{t+1} - X_t = \pm 1\}$ for all $t \in \mathbb{N}$, we get the following probabilities:

$$\mathbb{P}_k(X_{t+1} = i+1 | X_t = i) = \frac{e^{-\delta\beta}}{1 + e^{-\delta\beta}} = p$$

and

$$\mathbb{P}_k(X_{t+1} = i-1 | X_t = i) = \frac{1}{1 + e^{-\delta\beta}} = q.$$

Using this p and q for equation (5.1) we get

$$\mathbb{P}_k(\tau_0 < \tau_n) = \frac{e^{\delta\beta N} - e^{\delta\beta k}}{e^{\delta\beta N} - 1},$$

and so we have

$$\lim_{\beta\to\infty} \mathbb{P}_k(\tau_0 < \tau_N) = 1 \quad \text{for } k = 0, \ldots, n-1.$$

We can translate our original problem to the case of the gambler's ruin by translating state 1 to a fortune of size 0 and translating state $n-1$ to a fortune of state N. Then we have

$$\lim_{\beta\to\infty} \mathbb{P}_x(\tau_1 < \tau_{n-1}) = 1 \quad \text{for } x < n-2,$$

and since we must reach state $n-1$ before reaching state n it follows immediately that

$$\lim_{\beta\to\infty} \mathbb{P}_x(\tau_1 < \tau_n) = 1 \quad \text{for } x < n-2.$$

(2) We will show that the probability that τ_n is larger than $e^{\beta(\Gamma+\varepsilon)}$ is SES and that the probability that τ_n is smaller than $e^{\beta(\Gamma-\varepsilon)}$ tends to 0 as $\beta \to \infty$. To prove the first statement, define the set

$$R_n := \{X_t = \max\{X_0 + t, n\} \text{ for } t = 0, 1, \ldots, n\},$$

i.e., the set of all paths of length $n+1$ that start from any point $x \in \Omega$, move up each time step, and when they reach state n stay in state n. If X moves along a path in R_n, then the event $\{\tau_n \le n\}$ is realized. The probability that X moves along a path in R_n can be bounded from below as follows:

$$\begin{aligned}\mathbb{P}_x(\{X_0, X_1, \ldots, X_n\} \in R_n) &\ge \prod_{i=0}^{n-1} P(i, i+1) \ge \left(\frac{1}{2}\right)^n e^{\beta\Gamma} P(0,1) \\ &\ge e^{-\beta(\Gamma+\varepsilon/2)},\end{aligned}$$

where the last inequality holds if β is large enough. We now have

$$\begin{aligned}\mathbb{P}_x(\tau_n > e^{\beta(\Gamma+\varepsilon)}) &\leq (1-\mathbb{P}_x(\{X_0, X_1, \ldots, X_n\} \in R_n))^{e^{\beta(\Gamma+\varepsilon)}/2n} \\ &\leq \left(1 - e^{-\beta(\Gamma+\varepsilon/2)}\right)^{e^{\beta(\Gamma+\varepsilon)}/2n} \leq e^{\left(e^{-\beta(\Gamma+\varepsilon/2)}e^{\beta(\Gamma+\varepsilon)}\right)/2n} \\ &= e^{e^{-\beta\varepsilon/2}/2n} = SES. \end{aligned} \tag{5.2}$$

To prove the other bound, remark that

$$\mathbb{P}_x(\tau_n < e^{\beta(\Gamma-\epsilon)}) \leq \mathbb{P}_x(\tau_1 < \tau_n < e^{\beta(\Gamma-\epsilon)}) + \mathbb{P}_x(\tau_1 > \tau_n), \tag{5.3}$$

and that

$$\mathbb{P}_x(\tau_1 < \tau_n < e^{\beta(\Gamma-\epsilon)}) \leq \mathbb{P}_1(\tau_n < e^{\beta(\Gamma-\epsilon)})$$

We have the following inequality

$$\mathbb{P}_1(\tau_n < e^{\beta(\Gamma-\epsilon)}) \leq \mathbb{P}_1(\tau_{n-1} < e^{\beta(\Gamma-\epsilon)}) = \sum_{t=1}^{e^{\beta(\Gamma-\epsilon)}} \sum_{\substack{\omega = (\omega_0 = 1, \ldots, \omega_t = n-1) \\ \omega_i \neq n-1, i = 1, \ldots, t-1}} P(\omega).$$

Since the chain is time reversible, we have

$$\begin{aligned} P(\omega) &:= P(\omega_0, \omega_1)P(\omega_1, \omega_2)\cdots P(\omega_{t-1}, \omega_t) \\ &= \frac{\mu(n-1)}{\mu(1)} P(\omega_t, \omega_{t-1})\cdots P(\omega_2, \omega_1)P(\omega_1, \omega_0) =: \frac{\mu(n-1)}{\mu(1)} P(\overleftarrow{\omega}), \end{aligned}$$

which yields

$$\mathbb{P}_1(e^{\beta(\Gamma-\epsilon)}) \leq \frac{\mu(n-1)}{\mu(1)} \sum_{t=1}^{e^{\beta(\Gamma-\epsilon)}} \sum_{\substack{\omega = (\omega_0 = 1, \ldots, \omega_t = n-1) \\ \omega_i \neq n-1, i = 1, \ldots, t-1}} P(\overleftarrow{\omega}).$$

The summation leaves out the paths that start from state $n-1$ and arrive at state 1 at time t but do not visit the states $n-1$ or n in between, and so the sum is smaller than 1. This gives us

$$\mathbb{P}_1(e^{\beta(\Gamma-\epsilon)}) < \frac{\mu(n-1)}{\mu(1)} e^{\beta(\Gamma-\epsilon)} = e^{\beta(\Gamma-\epsilon)} e^{-\beta(H(n-1)+H(1))} = e^{-\beta\varepsilon}. \tag{5.4}$$

Now using part (1) and equation (5.4), we see that both parts in the right hand side of equation (5.3) tend to 0 as $\beta \to \infty$. This together with inequality (5.2) gives the desired result

$$\mathbb{P}_x(e^{\beta(\Gamma-\epsilon)} < \tau_n < e^{\beta(\Gamma+\epsilon)}) = 1.$$

(3) For a discrete time Markov chain there are two ways to bound the expectation time of the hitting time τ_n. On one hand we have

$$\begin{aligned}\mathbb{E}_x(\tau_n) = \sum_{t=1}^{\infty} t\mathbb{P}_x(\tau_n = t) &> \sum_{t=e^{\beta(\Gamma-\epsilon)}}^{e^{\beta(\Gamma+\epsilon)}} t\mathbb{P}_x(\tau_n = t) \\ &> e^{\beta(\Gamma-\epsilon)}\mathbb{P}_x(e^{\beta(\Gamma-\epsilon)} < \tau_n < e^{\beta(\Gamma+\epsilon)}),\end{aligned}$$

and, on the other hand, we have

$$\mathbb{E}_x(\tau_n) = \sum_{t=1}^{\infty} \mathbb{P}_x(\tau_n > t) < e^{\beta(\Gamma+\varepsilon)} + \sum_{t > e^{\beta(\Gamma+\varepsilon)}} \mathbb{P}_x(\tau_n > t).$$

Letting $\beta \to \infty$, we see that part (2) of this theorem yields

$$e^{\beta(\Gamma-\varepsilon)} < \mathbb{E}_x(\tau_n) < e^{\beta(\Gamma+\varepsilon)}.$$

Since $\varepsilon > 0$ is arbitrary, therefor

$$\lim_{\beta\to\infty} \frac{1}{\beta} \ln \mathbb{E}_x(\tau_n) = \Gamma.$$

(4) Straightforward computation shows

$$\begin{aligned}
\mathbb{P}_x(\tau_{\{0,n-1\}} = 0) &= \sum_{t=1}^{\infty} \mathbb{P}_x(\{X_t = 0\} \cap \{\tau_{\{0,n-1\}} = t\}) \\
&\leq \sum_{t'=0}^{\infty} \mathbb{P}_x(\{X_{t'} = 1\} \cap \{\tau_{\{0,n-1\}} > t'\}) P(1,0) \\
&\leq e^{-\beta(H(0)-H(1))} \mathbb{E}_x(\tau_{\{0,n-1\}}) \\
&\leq e^{-\beta(H(0)-H(1))} \mathbb{E}_x(\tau_{n-1}) \\
&\leq e^{-\beta(H(0)-H(1))} \mathbb{E}_x(\tau_n) \\
&\leq e^{-\beta(H(0)-H(1))} e^{\beta(H(n-1)-H(1)+\varepsilon)} \\
&= e^{-\beta(H(n)-H(n-1)-\varepsilon)},
\end{aligned}$$

where the last equation follows from part (3).

(5) Let N be the number of failed attempts by the random walk to move from 1 to n, i.e., all attempts starting from state 1 and returning to state 1 before reaching state n. Call p the escape probability. Then $N \simeq \text{Geo}(p)$. Now remark that τ_n conditioned on the event that $\{\tau_1 < \tau_n\}$ can be written as

$$\tau_n = \tau_1 + \sum_{i=1}^{N} \tau_1^{(i)} + \tau_n',$$

where τ_1, $\tau_1^{(i)}$ and τ_n' are defined as

$$\begin{aligned}
\tau_1 &:= \inf\{t \geq 0 : X_t = 1 | X_0 = x\}, \\
\tau_1^{(i)} &:= \inf\{t > 0 : X_t = 1 \text{ and } X_s \neq n \text{ for } s \leq t | X_0 = 1\}, \\
\tau_n' &:= \inf\{t \geq 0 : X_t = n \text{ and } X_s \neq 1 \text{ for } s \leq t | X_0 = 1\}
\end{aligned}$$

The $\tau_1^{(i)}$ are clearly i.i.d. hitting times and represent all failed attempts to leave the metastable state. Easy computation shows that, as β tends to ∞, we have

$$\lim_{\beta\to\infty} \mathbb{E}_x(\tau_1) = \frac{3}{2}|x-1|, \qquad \lim_{\beta\to\infty} \mathbb{E}_1(\tau_n') = \frac{1}{2}(3n-4), \qquad \lim_{\beta\to\infty} \mathbb{E}_1(\tau_1^{(i)}) = 1.$$

Since all times are independently distributed, this yields via part (3) and the "Law of Large Numbers" that

$$\tau_n \overset{LLN}{\sim} N \quad \text{and} \quad \mathbb{E}(N) = \mathbb{E}_x(\tau_n) = e^{\beta\Gamma}.$$

Let $p := \mathbb{E}(\tau_n)^{-1}$. Since $p \ll 1$, we have

$$\mathbb{P}_1(N = k) = p(1-p)^k \sim pe^{-kp}, \quad k = 0, 1, \ldots,$$

and hence $N \sim \text{Exp}(p)$. Now, we get

$$\lim_{\beta\to\infty} \mathbb{P}_x(\tau_n > t\mathbb{E}_x(\tau_n)) = \lim_{\beta\to\infty} \mathbb{P}_x(N > t/p) = e^{-p\cdot t/p} = e^{-t}, \quad t > 0.$$

□

Theorem 5.1 gives the following characterization of metastability:

- If the dynamics X_t is in the vicinity of the local minimum, i.e., situated inside the well surrounding state 1, then it will move according to the drift towards the local minimum.
- After a time, that is exponentially large in β and dependent on the energy needed to reach the saddle point, Γ, a fluctuation occurs against the drift, after which the dynamics will move towards the global minimum.
- If we look at the dynamics at time steps of length $\mathbb{E}(\tau_n)$, i.e., the expected time in which a motion against the drift occurs, then the probability that the dynamics is in the global minimum is exponentially distributed with mean 1.

6 A Ferromagnet and Glauber Dynamics

Consider a Ferromagnet under an adjustable external magnetic field at low temperature T. Like most systems, it seeks a configuration that has a low energy level. We can imagine such a magnet as a system of spins, either pointing up or down, which prefers to be parallel to the external magnetic field. These spins behave as tiny magnets themselves, so they have the tendency to align and when they do their spin interaction energy is negative, while when they do not align their interaction energy is positive. Each individual spin has a negative energy if it aligns with the external magnetic field and positive energy when it does not.

Starting with a strong negative external magnetic field the system will have minimal energy when all spins are pointing down, so if the system is stable most spins will point this way. If we change the external magnetic field rapidly to weakly positive en preserve this new field the system would have minimal energy when all spins prefer to be pointing up. But due to the sudden change and the interaction energy between the parallel spins the system will stay in the old phase, which we now call a metastable state.

Every now and then, due to a random fluctuation, a spin flips, loosing energy because it aligns with the external magnetic field, but gaining energy because it no longer aligns with its neighbors. If the net result is positive, then the spin will have the tendency to flip back. After a random fluctuation a group of spins will flip up, loosing energy because this group of spins now align with the magnetic field, but gaining energy because the spins on the boundary of the group are not parallel to the spins outside the group. However, the net energy per spin is lower then when a single spin was flipped. If the net energy per spin that can be added by flipping a spin on the outside of the boundary of this group is increasing, then the spins will still have the tendency to flip back and we say that the group has a *subcritical* size. Eventually a group of spins will flip up so that the net energy per spin is decreasing when more spins are flipped up at the boundary. We say that such a group of spins has a *supercritical* size. The spins around the boundary of this group will now have the tendency to flip up, until a majority of the spins will point up and the system will finally be positively magnetized.

The same phenomenon occurs if we change the magnetization of the external field from positive to weakly negative and the system started with minimal energy, i.e., a configuration where most spins where pointing up.

We can model this system as follows. Let $\Lambda \subset \mathbb{Z}^2$ be a large finite box, with periodic boundary conditions. With each site $x \in \Lambda$ we associate an **Ising**-spin variable $\sigma(x) \in \{-1, 1\}$, indicating whether the spin at x is up or down corresponding, respectively, to $\sigma(x) = 1$ and $\sigma(x) = -1$. A spin configuration $\sigma = \{\sigma(x) : x \in \Lambda\}$ is an element of the configuration space $\mathcal{X} := \{-1, 1\}^\Lambda$. With each configuration $\sigma \in \mathcal{X}$ we associate an energy, as described above, by the following Hamiltonian

$$H(\sigma) := -\frac{J}{2} \sum_{(x,y)\in\Lambda^*} \sigma(x)\sigma(y) - \frac{h}{2} \sum_{x\in\Lambda} \sigma(x), \tag{6.1}$$

where Λ^* denotes the set of neighboring spins, $J > 0$ is the *ferromagnetic pair potential* acting between neighboring pair of spins and $h > 0$ is the *magnetic field* acting on the spins.

Let β denote the inverse temperature, i.e. $\beta = T^{-1}$. Via the Metropolis algorithm we get a stochastic dynamics on $\mathcal{X}$ given by

$$\sigma \longrightarrow \sigma' \text{ at rate } e^{-\beta[H(\sigma')-H(\sigma)]_+},$$

for all σ' obtained from σ by flipping exactly one spin. This stochastic dynamics is called the Glauber dynamics on a torus at inverse temperature $\beta > 0$. The Gibbs measure associated with H is

$$\mu_\beta(\sigma) = \frac{e^{-\beta H(\sigma)}}{Z},$$

where $Z = \sum_{\sigma \in \mathcal{X}} e^{-\beta H(\sigma)}$. The Gibbs measure μ_β is the reversible equilibrium of the Glauber Dynamics. We are interested in the metastable regime

$$h \in (0, 2J), \quad \beta \to \infty.$$

Two important configuration in this dynamics are: (1) the configuration where all spins are pointing down, i.e., equal to minus one, called the metastable state; (2) the configuration where all spins are pointing up, i.e., equal to plus one, called the equilibrium state. Let us denote these to configurations as

$$\boxminus := \{\sigma(x) = -1, \text{for all } x \in \Lambda\}$$
$$\boxplus := \{\sigma(x) = 1, \text{for all } x \in \Lambda\}$$

Since the Glauber dynamics is aperiodic and reversible, it will eventually converge to the Gibbs measure μ_β. As in the 1dRW, it will concentrate around the configuration with minimal energy, i.e., configuration $\boxplus$. After a very long time the Markov chain will be in this configuration with probability 1 as β tends to infinity. The other configurations again have a weight that is exponentially smaller in β, of which configuration $\boxminus$ has the largest weight. This configuration is a local minimum and will play the role of metastable state in the case of the Ferromagnet.

Starting from $\boxminus$, after a exponentially large time a random fluctuation makes a group of spins flip up. If the group is large enough, then the energy per spin added is decreasing and the dynamics will move to the global minimum, $\boxplus$. Since the Ferromagnet lives on a finite but large box $\Lambda \subset \mathbb{Z}^2$, the most convenient way of keeping the boundary of such a group small is a square. So we will look at configurations $\sigma_{l\times l} \in \mathcal{X}$, where $\sigma_{l\times l}$ is the configuration that contains an $l \times l$ square of spins pointing up, i.e., a square of ones, and all other spins pointing down, i.e., are minus ones. The energy of such a configuration, called E, in comparison to the metastable state $\boxminus$ is given by

$$E(\sigma_{l\times l}) := H(\boxminus) - H(\sigma_{l\times l}) = 4Jl - hl^2. \tag{6.2}$$

Simple calculation shows that the maximum energy of an $l \times l$ square of spins pointing up in a sea of spins pointing down is reached at $l = 2J/h$. To avoid duplicity, assume that $J/2h$ is not an integer, and take $l_c := \lceil J/2h \rceil$. As illustrated in figure 2, if $l < l_c$, then from a configuration $\sigma_{l\times l}$ the system will return to the metastable configuration $\boxminus$ again; such a configuration we call *subcritical* If $l > l_c$, then from a configuration $\sigma_{l\times l}$ the system will move to the equilibrium configuration $\boxplus$; such a configuration we call *supercritical*. In both cases this happens because the systems seeks a configuration of minimal energy.

Let $\mathbb{P}_\sigma$ be the law of the Glauber dynamics $(\sigma_t)_{t\in\mathbb{N}}$ on $\mathcal{X}$ starting from $\sigma_0 = \sigma$. Before we can state the theorems equivalent to Theorem 5.1 we need a few extra definitions. Let

$$\Gamma := \min_{\omega:\boxminus\to\boxplus} \max_{\sigma\in\omega} [H(\sigma) - H(\boxminus)]$$

be the *communication height* between $\boxminus$ and $\boxplus$, where the minimum runs over all *admissible* paths ω connecting $\boxminus$ and $\boxplus$, i.e., paths allowed by the Glauber dynamics, and the maximum runs over all configurations σ encountered along ω. Let

$$\mathcal{S} := \{\zeta \in \mathcal{X} : \exists \omega : \boxminus \to \boxplus, \zeta \in \omega : \max_{\sigma\in\omega} E(\sigma) = E(\zeta) = \Gamma\}$$

be the *communication level set* between $\boxminus$ and $\boxplus$, containing all configurations from all admissible paths between $\boxminus$ and $\boxplus$ that have an energy level equal to the communication height and equal to the maximum energy level of the path.

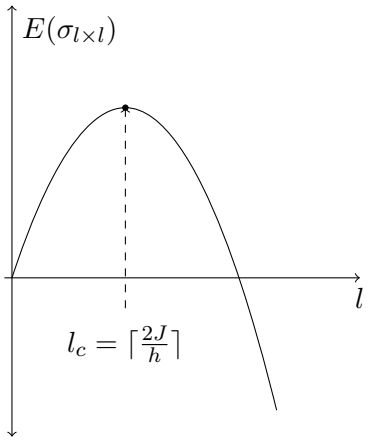

Figure 2: Plot of $E(\sigma_{l\times l})$.

Theorem 6.1 .

(i) $\lim_{\beta\to\infty} \mathbb{P}_{\boxminus}(e^{\beta(\Gamma-\delta)} < \tau_{\boxplus} < e^{\beta(\Gamma+\delta)} \mid \tau_{\boxplus} < \tau_{\boxminus}) = 1 \quad$ *for all* $\delta > 0$.

(ii) $\lim_{\beta\to\infty} \mathbb{P}_{\boxminus}(\tau_{\mathcal{S}} < \tau_{\boxplus} \mid \tau_{\boxplus} < \tau_{\boxminus}) = 1$.

Theorem 6.1.(i) shows that the required time to reach the equilibrium $\boxplus$ when starting the dynamics from $\boxminus$ takes an exponential time with $\Gamma = H(\mathcal{S})$ in the exponent of the transition time. Theorem 6.1.(ii) tells us that the communication level set $\mathcal{S}$ is a gate, i.e., all transitions from $\boxminus$ to $\boxplus$ must pass through $\mathcal{S}$. The bounds are sharp as $\beta \to \infty$.

To compute Γ and to refine the geometry of the relevant configurations in $\mathcal{S}$, define:

Definition 6.2 .

(i) *Let* $C = \bar{C} \cup \tilde{C}$, *with*

- $\bar{C}$ *is the set of all configurations containing a single* $l_c \times (l_c - 1)$ *square of spins pointing up, with a protuberance at a side of length* l_c *also pointing up, in a sea of spins pointing down.*
- $\tilde{C}$ *is the set of all configurations containing a single* $l_c \times (l_c - 1)$ *square of spins pointing up, with a protuberance at a side of length* $l_c - 1$ *also pointing up, in a sea of spins pointing down.*

(ii) *Let* $\Gamma^* = E(C) = 4J - h(l_c(l_c - l) + 1)$.

Theorem 6.3 .

(i) $C \subseteq \mathcal{S}$.

(ii) $\Gamma = \Gamma^*$.

(iii) $\lim_{\beta\to\infty} \mathbb{P}_{\boxminus}(\tau_{\bar{C}} < \tau_{\boxplus} \mid \tau_{\boxplus} < \tau_{\boxminus}) = 1$.

Theorem 6.3.(i) shows us that C is a proper subset of $\mathcal{S}$. Theorem 6.3.(ii) gives us a representation of Γ in terms of the parameters of the dynamics. By Theorem 6.3.(iii) we know that the true gate is made up from the configurations contained in $\bar{C}$, whereas the configurations in $\tilde{C}$ are *dead-ends*. The configurations in $\bar{C}$ form the critical droplets, because:

- *Adding* a bar of length l to a droplet costs $2J - h$ energy; adding a protuberance costs $2J$ energy since this creates two extra pairs of anti-parallel spins and gains h energy because an extra spin aligns with the field. The rest of the bar is added "downhill", i.e., extra spins do not create more anti-parallel pairs of spins, they only lower the energy cost of the droplet by h for each spin flipped up.
- *Removing* a bar of length l from the droplet costs $(l-1)h$ energy; removing $l-1$ spins costs h energy per spin, but gains no energy since no anti-parallel pairs are removed from the droplet. Removing the protuberance is "downhill" since we gain $2J$ energy for removing two anti-parallel pairs of spins and only lose h energy by flipping the protuberance down, since $h < 2J$.
- The two costs match when $l = 2J/h$. The configurations in $\bar{C}$ are obtained by first creating a 1×1 square, then successively adding bars of length $1, 2, \ldots, l_c - 1$ following a sequence of growing squares,

  $$1 \times 1, 1 \times 2, 2 \times 2, 2 \times 3, \ldots, (l_c - 1) \times (l_c - 1), (l_c - 1) \times l_c,$$

  and finally adding a protuberance, thereby reaching the "top of the hill". The rest of the process is "downhill".

We now state two theorems, giving further refinements.

Theorem 6.4 .

(i) $C \subsetneq \mathcal{S}$

(ii) $\lim_{\beta\to\infty} \mathbb{P}_{\boxminus}(\sigma_{\tau_{\bar{C}}} = \sigma | \tau_{\bar{C}} < \tau_{\boxminus}) = \frac{1}{|\bar{C}|}$ *for all* $\sigma \in \bar{C}$.

Theorem 6.4.(i) shows that $\mathcal{S}$ contains more configurations than C. An example of such a configuration is obtained by picking any configuration in C, flipping up a spin next to the protuberance and the quasi-square and flipping down one of the corners of the quasi-square. Clearly, Theorem 6.4.(ii) shows that the entrance distribution of $\bar{C}$ is uniform. This fact follows from symmetry arguments.

Theorem 6.5 .

(i) $\lim_{\beta\to\infty} \mathbb{E}_{\boxminus}(\tau_{\boxplus}) = Ke^{\beta\Gamma^*}(1 + O(1))$ *with* $K = \frac{3}{4(2l_c-1)}\frac{1}{|\Lambda|}$.

(ii) $\lim_{\beta\to\infty} \mathbb{P}_{\boxminus}(\tau_{\boxplus} > t\mathbb{E}_{\boxminus}(\tau_{\boxplus})) = (1 + O(1))\, e^{-t(1+O(1))}$ *for* $t \geq 0$.

Theorem 6.5.(i) gives a sharp asymptotics for the average magnetization time. The expectation is made up of two components. Namely, the average time it takes to reach a configuration of $\bar{C}$;

$$\frac{1}{|\bar{C}|} e^{\beta\Gamma^*}(1 + O(1)) \text{ as } \beta \to \infty,$$

while 1 over the probability to leave such a configuration into the direction of $\boxplus$ rather than $\boxminus$, is given by

$$\pi(l_c)^{-1} = \frac{3l_c}{2l_c - 1} \text{ as } \beta \to \infty.$$

Theorem 6.5.(ii) shows that the transition time is exponentially distributed, with each unsuccessful attempt to create a critical droplet ending with a return to the configuration $\boxminus$ where the system starts from scratch. This is typical for "success after many unsuccessful attempts".

Theorems 6.1-6.5 are equivalent to Theorem 5.1. In this section we have only given a sketch of the heuristics behind the proofs, whereas in the previous section we gave the complete proof. It

is way beyond the scope of this master thesis to dive into the technical details of these proofs. In three dimensions or higher, the proofs become even more complex. A path from the metastable phase to the stable phase is realized by the creation of a critical droplet after some random fluctuations. the construction of this critical droplet with the help of isoperimetric inequalities is a delicate matter.

7 A Supersaturated Gas and Kawasaki Dynamics

We consider a container with gas. When the density of the gas is high, say supersaturated, it has the tendency to condense and form droplets, e.g. rain consists of condensed water molecules. For this particular example of metastability, we examen a closed system containing gas molecules. We rapidly change the density of the molecules from a low density, such that molecules will have the tendency to move apart, to a slightly supersaturated density, such that the molecules will prefer to clump together.

As in the previous two examples (the Ferromagnet and the random walk on the double well), the system prefers to be in a state of minimal energy. Each molecule in the system adds some energy, whereas molecules that clump together lose some energy, due to van der Waals forces, here called *interaction energy*.

Under the influence of the Kawasaki dynamics the molecules will gradually clump together, but at low temperature this can take a very long time as we explain next. At some moment two molecules will meet and clump together, losing some energy, but not enough for another molecule to attach itself before the two molecule break up again. After another random fluctuation three molecules will clump together, having a slightly lesser energy level than in the case of two molecules, but once again before a next molecule can join they will break up. At a certain moment a large random fluctuation takes place and a large group of molecules clump together such that the total interaction energy is lower than the energy needed to put each molecule into the system. Now the group molecules is large enough to stick together so that a next molecule has the to attach itself to this droplet. Such a group of molecules we call a *supercritical droplet*, while a group of molecules that does not have the tendency to grow we call a *subcritical droplet*. Molecules will keep attaching themselves to the supercritical droplet, so that the gas will condensate.

We can model this process as follows. Let $\Lambda \subset \mathbb{Z}^2$ be a large finite box. With each site $x \in \Lambda$ we associate an occupation bit, $\eta(x) \in \{0, 1\}$, indicating whether a molecule is located at site x or not, represented, respectively, by $\eta(x) = 1$ and $\eta(x) = 0$. A gas configuration $\eta = \{\eta(x) : x \in \Lambda\}$ is an element of the configuration space $\mathcal{X} = \{0, 1\}^\Lambda$.

With each configuration $\eta \in \mathcal{X}$ we associate an energy level as described by the following Hamiltonian

$$H(\eta) := -U \sum_{(x,y)\in\Lambda^{-*}} \eta(x)\eta(y) + \Delta \sum_{x\in\Lambda} \eta(x),$$

where Λ^- is the set of all sites of Λ minus the boundary of Λ, denoted by $\partial\Lambda$, i.e., $\Lambda^- := \Lambda \setminus \partial\Lambda$ and Λ^{-*} is the set of all neighboring pairs in Λ^-. U denotes the *binding energy* between to neighboring molecules and Δ denotes the *activation energy* of single molecules.

Again, let β denote the inverse temperature, i.e., $\beta = T^{-1}$. Via the Metropolis algorithm we get a stochastic dynamics on $\mathcal{X}$, given by

$$\eta \longrightarrow \eta' \quad \text{at rate} \quad e^{-\beta[H(\eta')-H(\eta)]_+},$$

for all η' can be reached from η by either of the following two transitions:

1) A molecule is swapped between two neighboring sites, i.e., there exists a unique element $(x, y) \in \Lambda^*$ such that $\eta(x) \neq \eta(y)$, $\eta'(x) = \eta(y)$ and $\eta'(y) = \eta(x)$.
2) A molecule is created or removed at the boundary, i.e., there exists an unique element $x \in \partial\Lambda$ such that $\eta'(x) \neq \eta(x)$.

This corresponds to "motion" of molecules in Λ, respectively, the "creation" or "annihilation" of molecules at the boundary $\partial\Lambda$. We can imagine Λ to be a box in some greater gas reservoir, so that the creation and annihilation of particles at the boundary corresponds, respectively, to the hopping of particles in and out of the box. This is called the *Kawasaki dynamics* at inverse temperature $\beta > 0$.

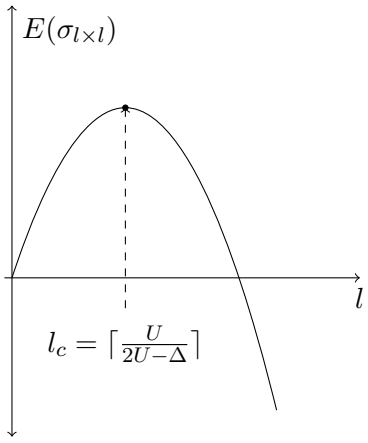

Figure 3: Plot of $E(\sigma_{l\times l})$.

The Gibbs measure associated with H is

$$\mu_\beta(\eta) := \frac{e^{-\beta H(\eta)}}{Z_\beta} \quad \forall \eta \in \mathcal{X},$$

with $Z_\beta = \sum_{\eta\in\mathcal{X}} e^{-\beta H(\eta)}$, which is the reversible equilibrium of the Kawasaki dynamic with density rate $\rho = e^{-\Delta\beta}$, i.e., the rate at which particles are created at the boundary $\partial\Lambda$. We are interested in the metastable regime

$$U > 0, \quad \Delta \in (U, 2U), \beta \to \infty.$$

As in the case of the Ferromagnet, we see that an energy efficient group has a form in which with as few particles as possible we have as many neighboring pairs as possible. Again, such a configuration is achieved when the particles form a square. With each square of size $l \times l$ of ones we associate the energy level

$$E(\eta_{l\times l}) := H(\eta_{l\times l}) = -U(2l(l-1)) + \Delta l^2.$$

Easy computation shows that a square of size $l = U/(2U-\Delta)$ has the highest energy level. For duplicity reasons assume that $U/(2U-\Delta)$ is not an integer and take $l_c = \lceil U/(2U-\Delta) \rceil$. This also explains why the parameters are chosen, such that $\Delta \in (U, 2U)$. As illustrated in figure 3, if a square has sides larger than l_c, then its energy is decreasing as it becomes larger and so it has a supercritical size, while a square has sides smaller than l_c, then its energy is increasing as it becomes larger and has a subcritical size.

Let $\mathbb{P}_\eta$ be the law of the Kawasaki dynamics $(\eta_t)_{t\geq 0}$ on $\mathcal{X}$ starting from $\eta_0 = \eta$. Before we can state the theorems equivalent to Theorems 5.1 and 6.1-6.5, we need the following definitions. Let

$$\Gamma := \min_{\omega:\square\to\blacksquare} \max_{\eta\in\omega} [H(\eta)]$$

be the *communication height* between $\square$ and $\blacksquare$, where the minimum runs over all *admissible* paths ω connecting $\square$ and $\blacksquare$, i.e., paths allowed by the Kawasaki dynamics, and the maximum runs over all configurations η encountered along ω. Let

$$\mathcal{S} := \{\zeta \in \mathcal{X} : \exists \omega : \square \to \blacksquare, \zeta \in \omega : \max_{\eta\in\omega} H(\eta) = H(\zeta) = \Gamma\}$$

be the *communication level set* between $\square$ and $\blacksquare$, containing all configurations from all admissible paths between $\square$ and $\blacksquare$ that have a energy equal to the communication height and to the maximum energy level of its path.

Theorem 7.1 .

(i) $\lim_{\beta\to\infty}\ \mathbb{P}_{\square}(e^{\beta(\Gamma-\delta)} < \tau_{\blacksquare} < e^{\beta(\Gamma+\delta)}) = 1,\ \forall\delta > 0.$

(ii) $\lim_{\beta\to\infty}\ \mathbb{P}_{\square}(\tau_{\mathcal{S}} < \tau_{\blacksquare} \mid \tau_{\blacksquare} < \tau_{\square}) = 1.$

Theorem 7.1.(i) shows that the required time to reach the equilibrium ■ is exponential with Γ in the exponent of the transition time. Theorem 7.1.(ii) shows that the communication level set $\mathcal{S}$ is a gate, i.e., all transitions from □ to ■ must pass through at least one configuration in $\mathcal{S}$. The bounds are sharp as $\beta \to \infty$.

As in the previous section, we can refine these statements by specifying what the configurations that act as gates look like and by computing Γ.

Definition 7.2 .

(i) *Let $\mathcal{Q} = \bar{\mathcal{Q}} \cup \tilde{\mathcal{Q}}$, with*

- *$\bar{\mathcal{Q}}$ is the set of all configurations containing a single $l_c \times (l_c - 1)$ square of particles, with a protuberance at a side of length l_c in an otherwise empty Λ.*
- *$\tilde{\mathcal{Q}}$ is the set of all configurations containing a single $l_c \times (l_c - 1)$ square of particles, with a protuberance at a side of length $l_c - 1$ in an otherwise empty Λ.*

(ii) *Let $\mathcal{D} \supset \mathcal{Q}$ be the set of configurations that can be reached from $\mathcal{Q}$ via a U-path, where a U-path is defined as $\omega : \eta \to \eta'$, with $\eta \in \mathcal{Q}$ and $\eta' \in \mathcal{D}$ such that*

$$\begin{aligned} &|\zeta| = |\eta| = |\eta'| \quad \text{for all } \zeta \in \omega, \\ &H(\eta) = H(\eta'), \\ &\max_{\zeta\in\omega} H(\zeta) \le H(\eta) + U, \end{aligned}$$

i.e., all configurations contain the same number of particles, the starting configuration and the end configuration in a U-path have the same energy level, and this path corresponds to the movement of a particle around the border of the droplet.

(iii) *Let $\mathcal{C} := \mathcal{D} + \{$a free particle in $\Lambda\}$.*

(iv) *Let*

$$\begin{aligned} \Gamma^* &= E(C) = H(\mathcal{D}) + \Delta = H(\mathcal{Q}) + \Delta \\ &= U[(l_c-1)^2 + l_c(l_c-2) + 1] + \Delta[l_c(l_c-1)+2] = 2U[l_c^2 - 2l_c + 1] + \Delta[l_c^2 - l_c + 2]. \end{aligned}$$

Theorem 7.3 .

(i) $C \subseteq \mathcal{S}$.

(ii) $\Gamma = \Gamma^*$.

(iii) $\lim_{\beta\to\infty}\ \mathbb{P}_{\square}(\tau_{\bar{C}} < \tau_{\blacksquare} \mid \tau_{\blacksquare} < \tau_{\square}) = 1.$

Theorem 7.3.(i) shows that C is a subset of $\mathcal{S}$. Theorem 7.3.(ii) gives us a representation of Γ in terms of the parameters of the dynamics. By Theorem 7.3.(iii) we know that the true gate consists of those configurations that are in C. In fact, the set C forms the minimal gate for the transition. Thus, C plays the role of the set of *critical* droplets, where Γ^* is the energy associated with such a droplet. These configurations are critical because:

- *Adding* a bar of length l to a droplet costs $2\Delta - U$ energy; a new particle is created at cost Δ, after which it moves towards the droplet at cost 0 and attaches itself at gain U. A new particle must be created at the boundary at cost Δ. Since traveling towards the droplet is free and attaching itself next to the protuberance now gains $2U$ energy, the rest of the process is "downhill", since $2U > \Delta$.

- *Removing* a bar of length l from the droplet costs $(l-2)(2U-\Delta)+2U$ energy; removing $l-2$ particles costs $2U-\Delta$ energy per particle, since for each particle that detaches itself the configuration loses $2U$ energy, after which it moves at cost 0 to the boundary and is removed at gain Δ. The $(l-1)$-th particle only needs to detach itself at cost $2U$, the rest of the process is again "downhill", since $2\Delta > U$.
- The two costs match when $l = U/(2U-\Delta)$. The configurations in $\bar{C}$ are obtained by first creating a 1×1 square, then successively adding bars of length $1, 2, \ldots, l_c - 1$ following a sequence of growing squares,

$$1\times 1, 1\times 2, 2\times 2, 2\times 3, \ldots, (l_c-1)\times(l_c-1), (l_c-1)\times l_c,$$

  and finally adding a protuberance, thereby reaching the "top of the hill".
- The configurations in $\mathcal{D}$ are those configurations the dynamics can reach from $\mathcal{Q}$ before the arrival of a new particle in the system. A U-path is completed in a time of order $e^{U\beta}$, whereas the creation of a new particle takes a time of order $e^{\Delta\beta}$, and $e^{U\beta} \ll e^{\Delta\beta}$. The arrival of the new particle moves the dynamics into C and completes the formation of the critical droplet.

Theorem 7.4 *$\mathcal{D} = \bar{\mathcal{D}} \cup \tilde{\mathcal{D}}$ with*

- *$\bar{\mathcal{D}}$ the set of configurations where the particles form a $(l_c-2)\times(l_c-2)$ square with four bars attached to the four sides of length $\bar{k}_i$ satisfying*

$$1 \le \bar{k}_i \le l_c - 1, \quad \sum_{i=1}^{4} \bar{k}_i = 3l_c - 3.$$

- *$\tilde{\mathcal{D}}$ the set of configurations where the particles form a $(l_c-3)\times(l_c-1)$ square with four bars attached to the four sides of length $\tilde{k}_i$ satisfying*

$$1 \le \tilde{k}_i \le l_c - 1, \quad \sum_{i=1}^{4} \tilde{k}_i = 3l_c - 2.$$

Remark that $\bar{\mathcal{Q}} \subset \bar{\mathcal{D}}$, $\tilde{\mathcal{Q}} \subset \tilde{\mathcal{D}}$, and that the configurations in $\mathcal{D}$ are precisely those configurations that are reachable via an U-path from any configuration in Q, i.e., a particle at cost U slides around at the boundary of the droplet towards the protuberance, where it reattaches itself at gain U.

Theorem 7.5 .

(i) $C \subsetneq \mathcal{S}$.

(ii) $\lim_{\beta\to\infty} \mathbb{P}_{\square}(\tau_{\mathcal{Q}} < \tau_C < \tau_{\blacksquare} \mid \tau_{\blacksquare} < \tau_{\square}) = 1$.

(iii) $\lim_{\beta\to\infty} \mathbb{P}_{\square}(\eta_{\tau_{C^-}} = \eta | \tau_C < \tau_{\square}) = \frac{1}{|\mathcal{D}|}$ *for all* $\eta \in \mathcal{D}$,

where τ_{C^-} is the time prior to τ_C.

Theorem 7.5.(i) shows that $\mathcal{S}$ contains more configurations than C. An example of such a configuration is obtained by picking any configuration in $\mathcal{Q}$, adding an extra particle at the boundary of the box, so that the energy of the system is equal to $H(C)$, and moving it next to the protuberance but without attaching it. Then let the protuberance detaches itself from the droplet and simultaneously attach itself to the free particle. Since the energy of the system does not change, this configuration is an element of $\mathcal{S}$. Theorem 7.5.(ii) shows that the dynamics must pass a configuration in $\mathcal{Q}$, and possibly through $C\setminus\mathcal{Q}$, before reaching C. Clearly, Theorem 7.5.(iii) shows that the entrance distribution of $\bar{C}$ is uniform, i.e., all configurations of $\mathcal{D}$ are equally likely before the extra particle is created.

Theorem 7.6 .

(i) *There exists a* $K > 0$ *such that*

$$\lim_{\beta\to\infty} \mathbb{E}_{\square}(\tau_{\blacksquare}) = Ke^{\beta\Gamma^*}\,(1+O(1))\,.$$

(ii) $\lim_{\beta\to\infty} \mathbb{P}_{\square}(\tau_{\blacksquare} > t\mathbb{E}_{\square}(\tau_{\blacksquare})) = (1+O(1))\,e^{-t(1+O(1))}$ *for* $t \geq 0$.

Theorem 7.6.(i) gives a sharp asymptotics for the average magnetization time. As in the case of the Ferromagnet, K is a constant derived from the probability that once the top of the hill is reached the dynamics exits towards the equilibrium and, from the geometry of the configurations in C. For the Kawasaki dynamics, K is a lot harder to determine than for the Glauber dynamics and hence we do not state further details. [4]

Theorem 7.6.(ii) shows that the transition time is exponentially distributed with each unsuccessful attempt to create a critical droplet ending with a return to the configuration $\square$, from which the system starts anew.

Theorems 7.1-7.6 are equivalent to Theorem 5.1 and Theorems 6.1-6.5. Again, we only gave a sketch of the proofs and the whole proofs are even more difficult than in the case of the Glauber dynamics, due to the difficulties that arise when describing the transition gates. In three dimensions a detailed description of the set $\mathcal{D}$ is not yet available. Again, isoperimetric problems play a key role in the metastable behavior of the system.

[4] For a detailed description of K see [3].

Part III
Conclusion

In this master thesis we have studied two possible types of convergence that can occur in with Markov chains, both being asymptotic in the limit as a certain parameter diverges. In Markov chains that show cutoff behavior, this parameter is given by the dimension, n, of the configuration or sample space. In Markov chains that show metastable behavior, this parameter is given by the inverse of the temperature, β.

Though both phenomena are asymptotic, they are very different in nature. This is best explained by looking at the random variable that corresponds to the time it takes the Markov chains come close to equilibrium. Let $U^{(n)}$ be family of random variables, that characterizes the mixing times for a family of Markov chains that shows cutoff. Then

$$\lim_{n\to\infty} \frac{U^{(n)}}{\mathbb{E}(U^{(n)})} = 1 \quad \text{in probability.}$$

This is equivalent to the statement that $\lim_{n\to\infty} \mathbb{P}(U^{(n)} > c\mathbb{E}(U^{(n)})) = 1$ for $c < 1$ and $= 0$ for $c > 1$, which corresponds to definition of cutoff, i.e., Definition 2.3. Let $V^{(\beta)}$ be a family of random variables that characterizes the times needed to leave the metastable state and reach the stable state for a family of Markov chains that show metastability. Then

$$\lim_{n\to\infty} \frac{V^{(\beta)}}{\mathbb{E}(V^{(\beta)})} = \exp(1) \quad \text{in distribution.}$$

This is equivalent to the statement that $\lim_{\beta\to\infty} \mathbb{P}(V^{(\beta)} > t\mathbb{E}(V^{(\beta)})) = e^{-t}$ for $t \geq 0$ as explained in Theorem 5.1, 6.5 and 7.6. Thus we see that cutoff is a deterministic asymptotics, whereas metastability is a random asymptotics. This dichotomy corresponds exactly to what we would expect from cutoff and metastability.

One of the targets of this master thesis was to explain what characterizes cutoff and metastability. For a family of Markov processes that show cutoff we have the following two characterizations:

1. The eigenvalues corresponding to the transition kernels of the Markov chains are ordered $1 = \lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_k > -1$ and the second eigenvalue has a high multiplicity.
2. There exists a *strong stationary time* t_{mix}, with the help of which we can bound the total variation distance to equilibrium from above. Near t_{mix}, the total variation drop from being close to one to being close to zero and the width of this drop is of a lower order than t_{mix}.

Both characteristics only give us an upper bound for the total variation distance. A lower bound of the same order can often be found by choosing some subset of the configuration space and showing that the bound holds at this subset.

From our examples it is clear that the state space does not have any influence on whether or not a family of Markov chains shows cutoff behavior. On the other hand, if a family of Markov processes shows metastable behavior, then we expect that the energy landscape, defined by the Hamiltonian H on the configuration space, is made up of hills and valleys, which represents a group of configuration that have a high energy, respectively, a low energy level. The process is stable as soon as it reaches the global minimum, which lies in the deepest valley. The time needed to leave the local minima, lying in higher valleys, is exponentially distributed according to the parameters of H.

Though cutoff was studied by P. Diaconis as early as the 1980's, the subject has not reached widespread attention. On reason may be that cutoff studies need a lot of algebra or combinatorics, as we have seen in the example of the GSR-shuffle. Later work off Peres gives a conjecture saying that a necessary and sufficient condition for cutoff is that the product of the mixing time and the spectral gap tends to infinity. In a paper of J. Barrera, O. Bertoncini, and R. Fernández [1], cutoff and metastability are linked because of this conjecture. They argue that both phenomena are naturally intertwined. Reaching a stable state from a metastable state happens after

some random fluctuation. When this happens, the transition is almost instantaneous, because the Markov chain is over the "top of the hill" and runs down to the bottom of the valley. The latter corresponds precisely to cutoff. An explanation is provided via an example especially constructed for this purpose, namely a certain death-birth process on a finite interval.

In this master thesis we have discussed both cutoff and metastability at the hand of some concrete examples. These examples give insight in to what drives both phenomena.

References

[1] J. Barrera, O. Bertoncini, and R. Fernández. Abrupt convergence and escape behavior for birth and death chains. *Jour. of Stat. Phys.*, pages 595–623, 2019.

[2] F. den Hollander. Coupling. Notes from a course given in the Fall and Winter of 2010 at the University of Leiden.

[3] F. den Hollander. Metastability under stochastic dynamics. *Stoch. Proc. Appl.*, 114:1–26, 2004.

[4] P. Diaconis. The cutoff phenomenon in finite markov chains. *Proc. Nath. Acad. Sci. USA*, 93:1659–1664, 1996.

[5] P. Diaconis and D. Aldous. Shuffling cards and stopping times. *Amer. Math'l Monthly*, 5:333–348, 1986.

[6] P. Diaconis and D. Bayer. Trailing the dovetail shuffle to its lair. *Ann Appl. Prob.*, 2:294–313, 1992.

[7] D.A. Levin, Y. Peres, and E.L. Wilmer. *Markov Chains and Mixing Times.* American Mathematical Society, 2009.

[8] E. Olivieri and E. Scoppola. An introduction to metastability through random walks. *Brazilian Prob. Stat.*, 2010.

[9] D. Williams. *Probability with Martingales.* Cambridge University Press, Cambridge, 1991.