Randomised Algorithms

Lecture 4: Markov Chains and Mixing Times

Thomas Sauerwald (tms41@cam.ac.uk)

Lent 2022

UNIVERSITY OF CAMBRIDGE

Outline

Applications of Markov Chains in Computer Science

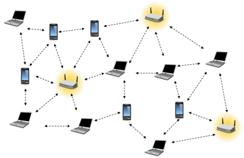

Broadcasting

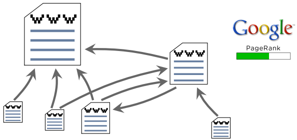

Ranking Websites

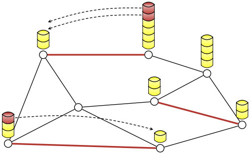

Load Balancing

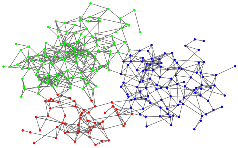

Clustering

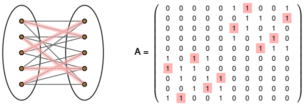

Sampling and Optimisation

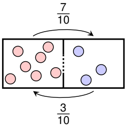

Particle Processes

Markov Chains

Markov Chain (Discrete Time and State, Time Homogeneous)

We say that $(X_t)_{t=0}^{\infty}$ is a Markov Chain on State Space Ω with Initial Distribution μ and Transition Matrix P if:

1. For any $x \in \Omega$, $\mathbf{P}[X_0 = x] = \mu(x)$.
2. The Markov Property holds: for all $t \geq 0$ and any $x_0, \ldots, x_{t+1} \in \Omega$,

$$\mathbf{P}\left[X_{t+1} = x_{t+1} \,\middle|\, X_t = x_t, \ldots, X_0 = x_0\right] = \mathbf{P}\left[X_{t+1} = x_{t+1} \,\middle|\, X_t = x_t\right]$$
$$:= P(x_t, x_{t+1}).$$

From the definition one can deduce that (check!)

- For all $t, x_0, x_1, \ldots, x_t \in \Omega$,

$$\begin{aligned}&\mathbf{P}[X_t = x_t, X_{t-1} = x_{t-1}, \ldots, X_0 = x_0] \\ &= \mu(x_0) \cdot P(x_0, x_1) \cdot \ldots \cdot P(x_{t-2}, x_{t-1}) \cdot P(x_{t-1}, x_t).\end{aligned}$$

- For all $0 \leq t_1 < t_2, x \in \Omega$,

$$\mathbf{P}[X_{t_2} = x] = \sum_{y \in \Omega} \mathbf{P}[X_{t_2} = x \mid X_{t_1} = y] \cdot \mathbf{P}[X_{t_1} = y].$$

What does a Markov Chain Look Like?

Example: the carbohydrate served with lunch in the college cafeteria.

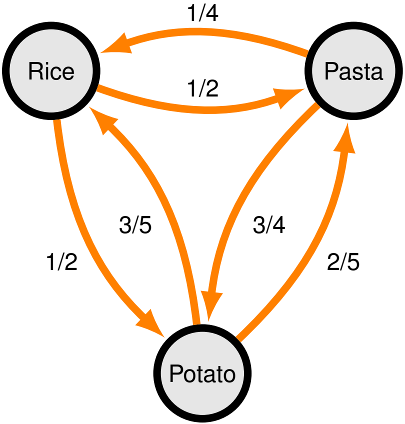

This has transition matrix:

$$
\begin{array}{c}
\begin{array}{ccc} \text{Rice} & \text{Pasta} & \text{Potato} \end{array} \\
P = \begin{bmatrix} 0 & 1/2 & 1/2 \\ 1/4 & 0 & 3/4 \\ 3/5 & 2/5 & 0 \end{bmatrix} \begin{array}{l} \text{Rice} \\ \text{Pasta} \\ \text{Potato} \end{array}
\end{array}
$$

Transition Matrices and Distributions

The Transition Matrix P of a Markov chain (μ, P) on $\Omega = \{1, \ldots n\}$ is given by

$$P = \begin{pmatrix} P(1,1) & \ldots & P(1,n) \\ \vdots & \ddots & \vdots \\ P(n,1) & \ldots & P(n,n) \end{pmatrix}.$$

- $\rho^t = (\rho^t(1), \rho^t(2), \ldots, \rho^t(n))$: state vector at time t (row vector).
- Multiplying ρ^t by P corresponds to advancing the chain one step:

$$\rho^t(y) = \sum_{j \in \Omega} \rho^{t-1}(x) \cdot P(x,y) \qquad \text{and thus} \qquad \rho^t = \rho^{t-1} \cdot P.$$

- The Markov Property and line above imply that for any $t \geq 0$

$$\rho^t = \rho \cdot P^{t-1} \qquad \text{and thus} \qquad P^t(x,y) = \mathbf{P}\left[\, X_t = y \mid X_0 = x \,\right].$$

Thus $\rho^t(x) = (\mu P^t)(x)$ and so $\rho^t = \mu P^t = (\mu P^t(1), \mu P^t(2), \ldots, \mu P^t(n))$.

- Everything boils down to deterministic vector/matrix computations

$\Rightarrow$ can replace ρ by any (load) vector and view P as a balancing matrix!

Stopping and Hitting Times

A non-negative integer random variable τ is a stopping time for $(X_t)_{t \geq 0}$ if for every $s \geq 0$ the event $\{\tau = s\}$ depends only on $X_0, \ldots, X_s$.

Example - College Carbs Stopping times:

✓ "We had rice yesterday" $\rightsquigarrow$ $\tau := \min\{t \geq 1 \colon X_{t-1} = \text{"rice"}\}$

× "We are having pasta next Thursday"

For two states $x, y \in \Omega$ we call $h(x, y)$ the hitting time of y from x:

$$h(x, y) := \mathbf{E}_x[\tau_y] = \mathbf{E}\left[\,\tau_y \mid X_0 = x\,\right] \quad \text{where } \tau_y = \min\{t \geq 1 : X_t = y\}.$$

Some distinguish between $\tau_y^+ = \min\{t \geq 1 \colon X_t = y\}$ and $\tau_y = \min\{t \geq 0 \colon X_t = y\}$

A Useful Identity

Hitting times are the solution to a set of linear equations:

$$h(x, y) \overset{\text{Markov Prop.}}{=} 1 + \sum_{z \in \Omega \setminus \{y\}} P(x, z) \cdot h(z, y) \qquad \forall x \neq y \in \Omega.$$

Outline

Irreducible Markov Chains

A Markov Chain is irreducible if for every state $x \in \Omega$ there is an integer $k \geq 0$ such that $P^k(x, x) > 0$.

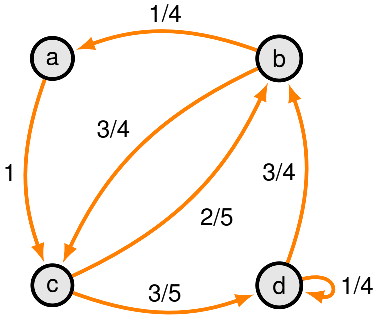

✓ irreducible

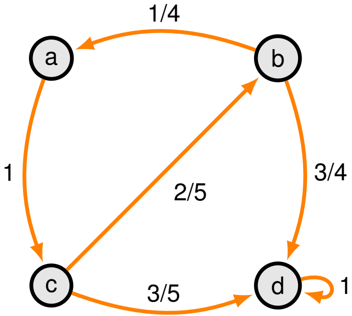

× not-irreducible (thus reducible)

Finite Hitting Time Theorem

For any states x and y of a finite irreducible Markov Chain $h(x, y) < \infty$.

Stationary Distribution

A probability distribution $\pi = (\pi(1), \ldots, \pi(n))$ is the stationary distribution of a Markov Chain if $\pi P = \pi$ (π is a left eigenvector with eigenvalue 1)

College carbs example:

$$\underset{\pi}{\left(\frac{4}{13}, \frac{4}{13}, \frac{5}{13}\right)} \cdot \underset{P}{\begin{pmatrix} 0 & 1/2 & 1/2 \\ 1/4 & 0 & 3/4 \\ 3/5 & 2/5 & 0 \end{pmatrix}} = \underset{\pi}{\left(\frac{4}{13}, \frac{4}{13}, \frac{5}{13}\right)}$$

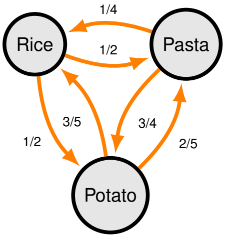

- A Markov Chain reaches stationary distribution if $\rho^t = \pi$ for some t.
- If reached, then it persists: If $\rho^t = \pi$ then $\rho^{t+k} = \pi$ for all $k \geq 0$.

Existence and Uniqueness of a Positive Stationary Distribution

Let P be finite, irreducible M.C., then there exists a unique probability distribution π on Ω such that $\pi = \pi P$ and $\pi(x) = 1/h(x,x) > 0$, $\forall x \in \Omega$.

Periodicity

- A Markov Chain is aperiodic if for all $x \in \Omega$, $\gcd\{t \geq 1 : P^t_{x,x} > 0\} = 1$.
- Otherwise we say it is periodic.

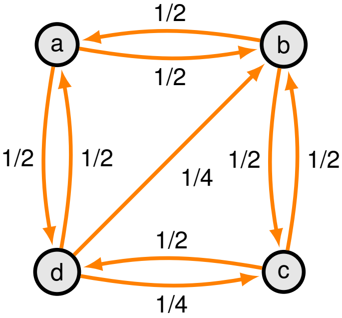

✓ Aperiodic

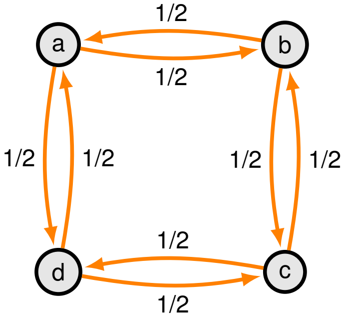

× Periodic

Exercise: Which of the two chains (if any) are aperiodic?

Convergence Theorem

Ergodic $=$ Irreducible $+$ Aperiodic

Convergence Theorem

Let P be any finite, irreducible, aperiodic Markov Chain with stationary distribution π. Then for any $x, y \in \Omega$,

$$\lim_{t \to \infty} P_{x,y}^t = \pi_y.$$

- mentioned before: For finite irreducible M.C.'s π exists, is unique and

$$\pi_y = \frac{1}{h(y,y)} > 0.$$

- We will prove a simpler version of the Convergence Theorem after introducing Spectral Graph Theory.

Convergence to Stationarity (Example)

- Markov Chain: stays put with $1/2$ and moves left (or right) w.p. $1/4$
- At step t the value at vertex $x \in \{1, 2, \ldots, 12\}$ is $P^t(1, x)$.

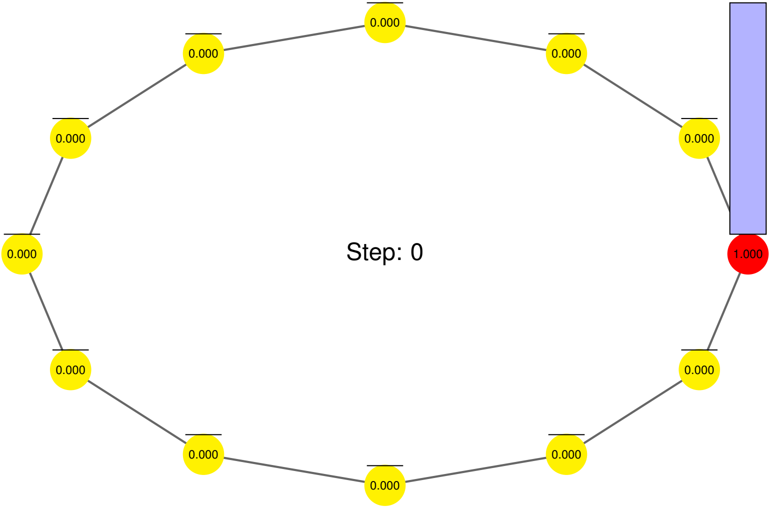

Convergence to Stationarity (Example)

- Markov Chain: stays put with $1/2$ and moves left (or right) w.p. $1/4$
- At step t the value at vertex $x \in \{1, 2, \ldots, 12\}$ is $P^t(1, x)$.

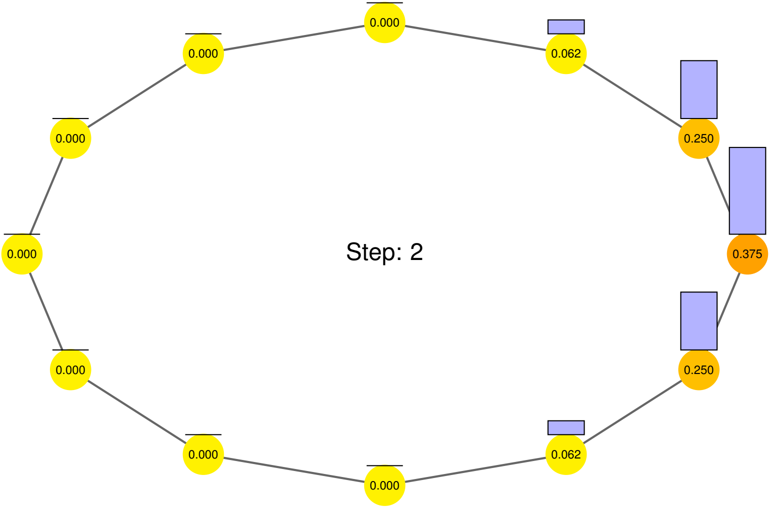

Convergence to Stationarity (Example)

- Markov Chain: stays put with $1/2$ and moves left (or right) w.p. $1/4$
- At step t the value at vertex $x \in \{1, 2, \ldots, 12\}$ is $P^t(1, x)$.

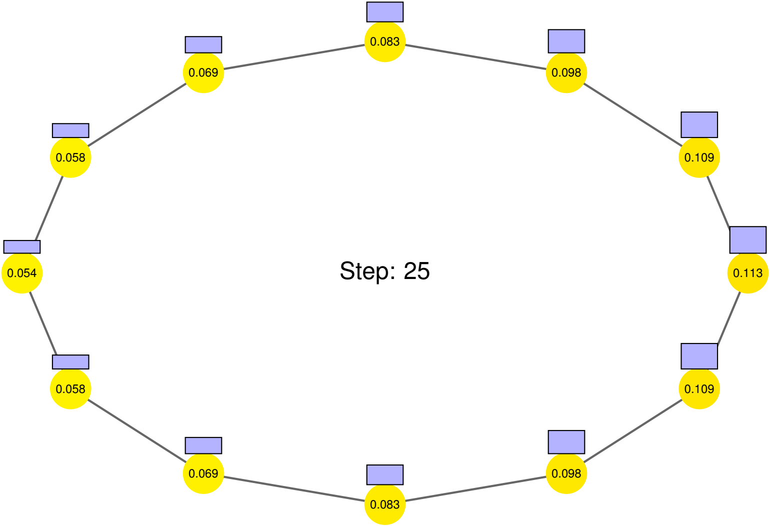

Convergence to Stationarity (Example)

- Markov Chain: stays put with $1/2$ and moves left (or right) w.p. $1/4$
- At step t the value at vertex $x \in \{1, 2, \ldots, 12\}$ is $P^t(1, x)$.

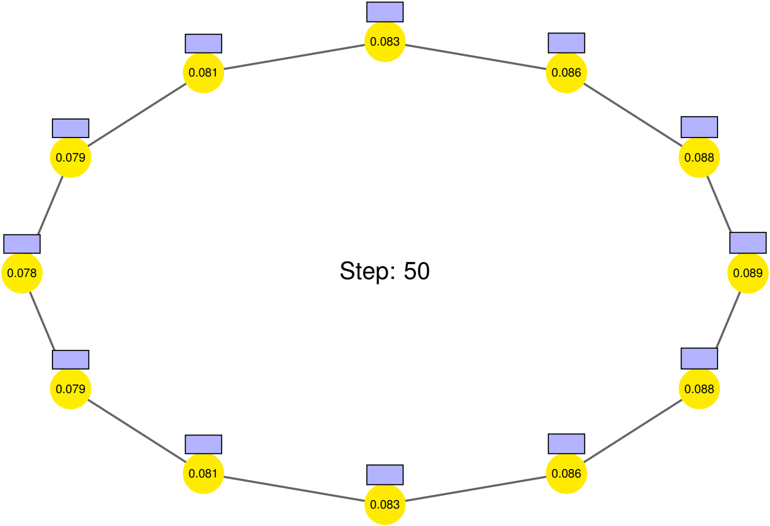

Outline

How Similar are Two Probability Measures?

Loaded Dice

- You are presented three loaded (unfair) dice A, B, C:

x	1	2	3	4	5	6
$\mathbf{P}[A = x]$	1/3	1/12	1/12	1/12	1/12	1/3
$\mathbf{P}[B = x]$	1/4	1/8	1/8	1/8	1/8	1/4
$\mathbf{P}[C = x]$	1/6	1/6	1/8	1/8	1/8	9/24

- Question 1: Which dice is the least fair? Most of you choose A. Why?
- Question 2: Which dice is the most fair? Dice B and C seem "fairer" than A but which is fairest?

We need a formal "fairness measure" to compare probability distributions!

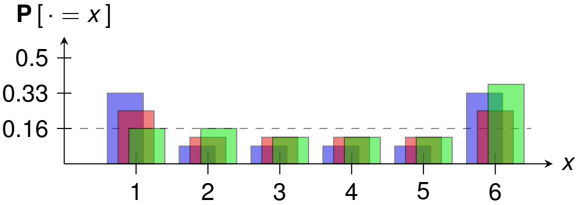

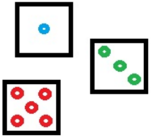

Total Variation Distance

The Total Variation Distance between two probability distributions μ and η on a countable state space Ω is given by

$$\|\mu - \eta\|_{tv} = \frac{1}{2} \sum_{\omega \in \Omega} |\mu(\omega) - \eta(\omega)|.$$

Loaded Dice: let $D = \mathit{Unif}\{1,2,3,4,5,6\}$ be the law of a fair dice:

$$\|D - A\|_{tv} = \frac{1}{2}\left(2\left|\frac{1}{6} - \frac{1}{3}\right| + 4\left|\frac{1}{6} - \frac{1}{12}\right|\right) = \frac{1}{3}$$

$$\|D - B\|_{tv} = \frac{1}{2}\left(2\left|\frac{1}{6} - \frac{1}{4}\right| + 4\left|\frac{1}{6} - \frac{1}{8}\right|\right) = \frac{1}{6}$$

$$\|D - C\|_{tv} = \frac{1}{2}\left(3\left|\frac{1}{6} - \frac{1}{8}\right| + \left|\frac{1}{6} - \frac{9}{24}\right|\right) = \frac{1}{6}.$$

Thus

$$\|D - B\|_{tv} = \|D - C\|_{tv} \qquad \text{and} \qquad \|D - B\|_{tv}, \|D - C\|_{tv} < \|D - A\|_{tv}.$$

So A is the least "fair" however B and C are equally "fair" (in TV distance).

TV Distances and Markov Chains

Let P be a finite Markov Chain with stationary distribution π.

- Let μ be a prob. vector on Ω (might be just one vertex) and $t \geq 0$. Then

$$P_\mu^t := \mathbf{P}\left[X_t = \cdot \mid X_0 \sim \mu \right],$$

  is a probability measure on Ω.

- For any μ,

$$\left\| P_\mu^t - \pi \right\|_{tv} \leq \max_{x \in \Omega} \left\| P_x^t - \pi \right\|_{tv}.$$

Convergence Theorem (Implication for TV Distance)

For any finite, irreducible, aperiodic Markov Chain

$$\lim_{t \to \infty} \max_{x \in \Omega} \left\| P_x^t - \pi \right\|_{tv} = 0.$$

We will prove a similar result later after introducing spectral techniques!

Mixing Time of a Markov Chain

Convergence Theorem: "Nice" Markov Chains converge to stationarity.

Question: How fast do they converge?

Mixing Time

The Mixing time $\tau_x(\epsilon)$ of a finite Markov Chain P with stationary distribution π is defined as

$$\tau_x(\epsilon) = \min \left\{ t\colon \left\| P_x^t - \pi \right\|_{tv} \leq \epsilon \right\},$$

and,

$$\tau(\epsilon) = \max_x \tau_x(\epsilon).$$

- This is how long we need to wait until we are "ε-close" to stationarity
- We often take $\varepsilon = 1/4$, indeed let $t_{mix} := \tau(1/4)$

Outline

What is Card Shuffling?

Source: wikipedia

Here we will focus on one shuffling scheme which is easy to analyse.

How long does it take to shuffle a deck of 52 cards?

How quickly do we converge to the uniform distribution over all $n!$ permutations?

His research revealed a lot of beautiful connections between Markov Chains and Algebra.

Persi Diaconis (Professor of Statistics and former Magician)

Source: `www.soundcloud.com`

The Card Shuffling Markov Chain

TopToRandomShuffle (Input: A pile of n cards)

```
For t = 1, 2, ...
    Pick i ∈ {1, 2, ..., n} uniformly at random
    Take the top card and insert it behind the i-th card
```

This is a slightly informal definition, so let us look at a small example...

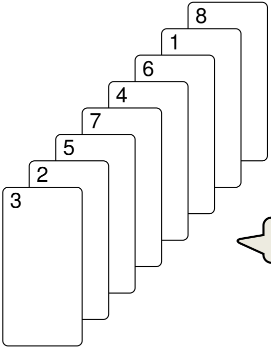

We will focus on this "small" set of cards ($n = 8$)

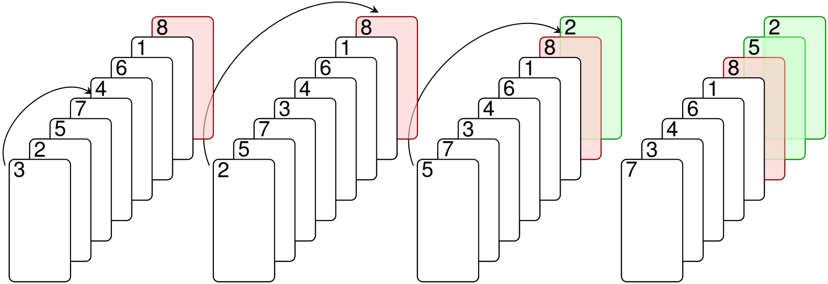

Even if we know which set of cards come after 8, every permutation is equally likely!

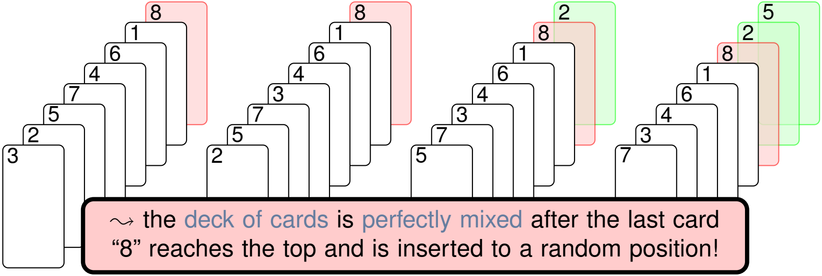

$\rightsquigarrow$ the deck of cards is perfectly mixed after the last card "8" reaches the top and is inserted to a random position!

Analysing the Mixing Time (Intuition)

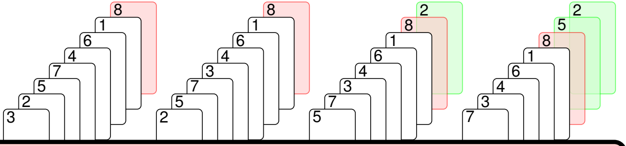

$\rightsquigarrow$ deck of cards is perfectly mixed after the last card "8" reaches the top and is inserted to a random position!

- How long does it take for the last card "n" to become top card?
- At the last position, card "n" moves up with probability $\frac{1}{n}$ at each step
- At the second last position, card "n" moves up with probability $\frac{2}{n}$

$\vdots$

- At the second position, card "n" moves up with probability $\frac{n-1}{n}$
- One final step to randomise card "n" (with probability 1)

This is a **"reversed" coupon collector** process with n cards, which takes $n \log n$ in expectation.

Using the so-called coupling method, one could prove $t_{mix} \leq n \log n$.

Analysis of Riffle-Shuffle

Riffle Shuffle

1. Split a deck of n cards into two piles (thus the size of each portion will be Binomial)
2. Riffle the cards together so that the card drops from the left (or right) pile with probability proportional to the number of remaining cards

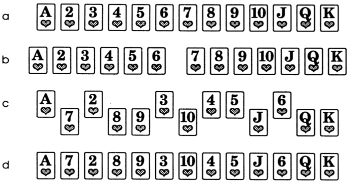

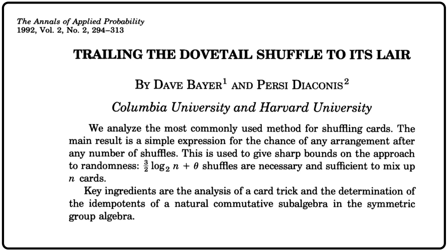

The Annals of Applied Probability
1992, Vol. 2, No. 2, 294–313

TRAILING THE DOVETAIL SHUFFLE TO ITS LAIR

By Dave Bayer[1] and Persi Diaconis[2]

Columbia University and Harvard University

We analyze the most commonly used method for shuffling cards. The main result is a simple expression for the chance of any arrangement after any number of shuffles. This is used to give sharp bounds on the approach to randomness: $\frac{3}{2}\log_2 n + \theta$ shuffles are necessary and sufficient to mix up n cards.

Key ingredients are the analysis of a card trick and the determination of the idempotents of a natural commutative subalgebra in the symmetric group algebra.

t	1	2	3	4	5	6	7	8	9	10
$\|P^t - \pi\|_{tv}$	1.000	1.000	1.000	1.000	0.924	0.614	0.334	0.167	0.085	0.043

Figure: Total Variation Distance for t riffle shuffles of 52 cards.

Outline

The Ehrenfest Markov Chain

Ehrenfest Model

- A simple model for the exchange of molecules between two boxes
- We have d particles labelled $1, 2, \ldots, d$
- At each step a particle is selected uniformly at random and switches to the other box
- If $\Omega = \{0, 1, \ldots, d\}$ denotes the number of particles in the red box, then:

$$P_{x,x-1} = \frac{x}{d} \qquad \text{and} \qquad P_{x,x+1} = \frac{d-x}{d}.$$

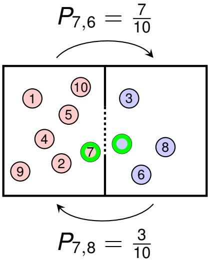

Let us now enlarge the state space by looking at each particle individually!

Random Walk on the Hypercube

- For each particle an indicator variable $\Rightarrow \Omega = \{0, 1\}^d$
- At each step: pick a random coordinate in $[d]$ and flip it

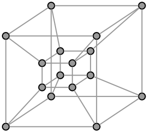

Analysis of the Mixing Time

(Non-Lazy) Random Walk on the Hypercube

- For each particle an indicator variable $\Rightarrow \Omega = \{0,1\}^d$
- At each step: pick a random coordinate in $[d]$ and flip it

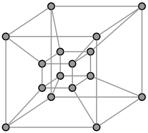

Problem: This Markov Chain is periodic, as the number of ones always switches from odd to even!

Solution: Add self-loops to break periodic behaviour!

Lazy Random Walk (1st Version)

- At each step $t = 0, 1, 2 \ldots$
  - Pick a random coordinate in $[d]$
  - With prob. $1/2$ flip coordinate.

Lazy Random Walk (2nd Version)

- At each step $t = 0, 1, 2 \ldots$
  - Pick a random coordinate in $[d]$
  - Set coordinate to $\{0, 1\}$ uniformly.

These two chains are equivalent!

Example of a Random Walk on a 4-Dimensional Hypercube

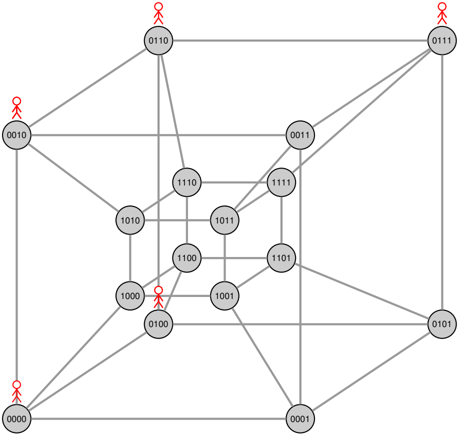

t	Coord.	X_t			
0	2	0	0	0	0
1	3	0	1	0	0
2	3	0	1	0	0
3	4	0	1	1	0
4	2	0	1	1	1
5	4	0	1	1	1
6	2	0	1	1	0
7	4	0	0	1	0
8	3	0	0	1	0
9	1	0	0	1	0
10	done!	0	0	1	0

Once all coordinates have been picked at least once, the state is uniformly at random in $\{0,1\}^d$.

Coupon Collector $\rightsquigarrow$ mixing time should be $O(n \log n)$

We will not formalise this argument here (see exercise question)

Total Variation Distance of Random Walk on Hypercube ($d = 22$)

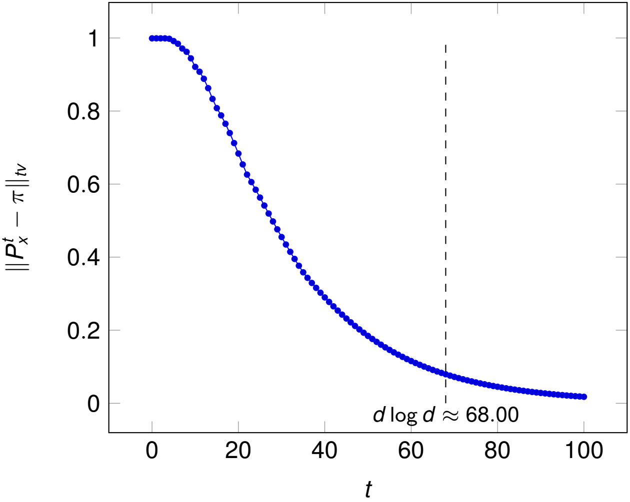

Theoretical Results (by Diaconis, Graham and Morrison)

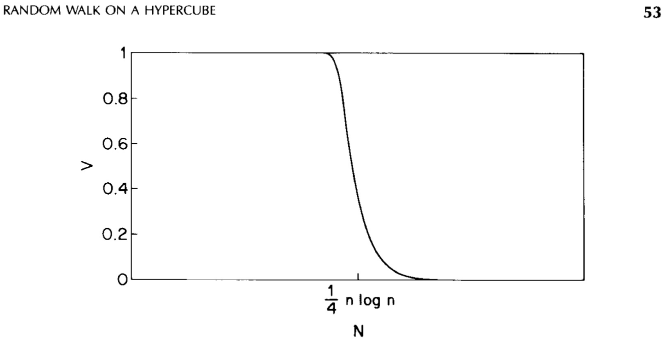

Fig. 1. The variation distance V as a function of N, for $n = 10^{12}$.

Source: "Asymptotic analysis of a random walk on a hypercube with many dimensions", P. Diaconis, R.L. Graham, J.A. Morrison; Random Structures & Algorithms, 1990.

- This is a numerical plot of a theoretical bound, where $d = 10^{12}$
  (Minor Remark: This random walk is with a loop probability of $1/(d+1)$)
- The variation distance exhibits a so-called cut-off phenomena:
  - Distance remains close to its maximum value 1 until step $\frac{1}{4} n \log n - \Theta(n)$
  - Then distance moves close to 0 before step $\frac{1}{4} n \log n + \Theta(n)$

Outline

Recap of Markov Chain Basics

Irreducibility, Periodicity and Convergence

Total Variation Distance and Mixing Times

Application 1: Card Shuffling

Application 2: Ehrenfest Chain and Hypercubes

Application 3: Markov Chain Monte Carlo

A Markov Chain for Sampling Independent Sets (1/2)

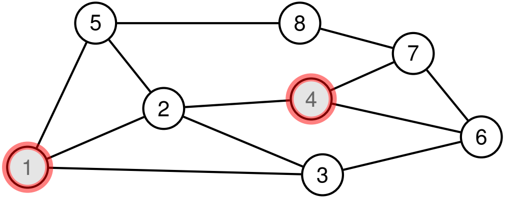

$S = \{1, 4\}$ is an independent set ✓

Independent Set

Given an undirected graph $G = (V, E)$, an independent set is a subset $S \subseteq V$ such that there are no two vertices $u, v \in S$ with $\{u, v\} \in E(G)$.

How can we take a sample from the space of all independent sets?

Naive brute-force would take an insane amount of time (and space)!

We can use a generic Markov Chain Monte Carlo approach to tackle this problem!

A Markov Chain for Sampling Independent Sets (2/2)

INDEPENDENTSETSAMPLER

Let X_0 be an arbitrary independent set in G
For $t = 1, 2, \ldots$:
  Pick a vertex $v \in V(G)$ uniformly at random
  If $v \in X_t$ **then** $X_{t+1} \leftarrow X_t \setminus \{v\}$
  elif $v \notin X_t$ **and** $X_t \cup \{v\}$ is an independent set **then** $X_{t+1} \leftarrow X_t \cup \{v\}$
  else $X_{t+1} \leftarrow X_t$

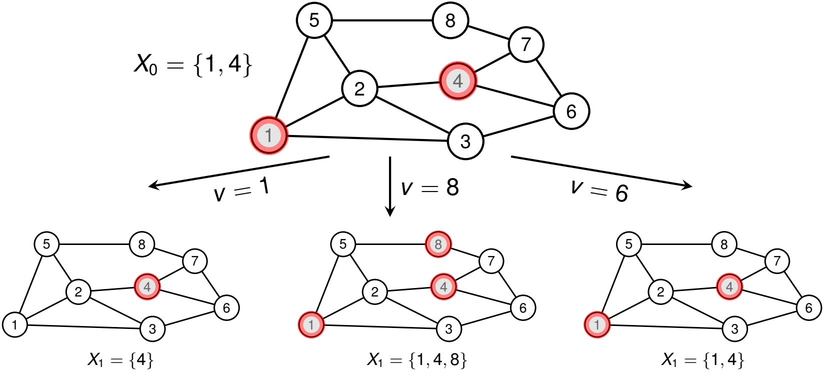

A Markov Chain for Sampling Independent Sets (2/2)

INDEPENDENTSETSAMPLER

Let X_0 be an arbitrary independent set in G
For $t = 1, 2, \ldots$:
    Pick a vertex $v \in V(G)$ uniformly at random
    If $v \in X_t$ **then** $X_{t+1} \leftarrow X_t \setminus \{v\}$
    elif $v \notin X_t$ **and** $X_t \cup \{v\}$ is an independent set **then** $X_{t+1} \leftarrow X_t \cup \{v\}$
    else $X_{t+1} \leftarrow X_t$

Remark

- This is a local definition (no explicit definition of P!)
- This chain is irreducible (every independent set is reachable)
- This chain is aperiodic (Check!)
- The stationary distribution is uniform, since $P_{u,v} = P_{v,u}$ (Check!)

Key Question: What is the mixing time of this Markov Chain?

not covered here, see the textbook of Mitzenmacher & Upfal