CS630 Homework 4

Qilin Ye

March 18, 2025

Solution to problem 1. The intuition here is that the stationary distribution must be uniform, as this process is doubly stochastic. To this end, assume $\pi_i = 1/n$ for each state i. We need to verify $\pi P = \pi$: for each i,

$$(\pi P)_i = \sum_{j=1}^{n} \pi_j P_{i,j} = \frac{1}{n} \sum_{j=1}^{n} P_{i,j} = \frac{1}{n}$$

as expected. Finally, under suitable technical assumptions (irreducible, finite MC) the stationary distribution is unique, so our proof is complete.

Solution to problem 2. We consider the coupling (X, Y) modeling the mouse-cat location pair defined on V^2, whose transition is given by the coordinate-wise random walk transition probability. The nontrivial edges are given by E^2, and the trivial ones include walks along one coordinate only, leaving the other coordinate in self-loop (e.g. mouse moves, cat stays), or the double self-loop $(u, v) \to (u, v)$. Under this Markovian coupling, the mouse is caught when starting from some state (u, v), the coupling ends at some (w, w) for $w \in V$.
In fact we will show a stronger result: we show that the expected hitting time from (u, v) to (u, u) is $\mathcal{O}(mn^2)$. In particular we consider any path from (u, v) to (u, u). This can be trivially attained by keeping the mouse fixed at u, while the cat travels from v to u in $\mathcal{O}(n)$ steps. Each $(u, v) \in V^2$ has $\deg(u)\deg(v)$ neighbors, and hence the graph induced by the Markovian coupling has

$$\sum_{u,v \in V} \deg(u, v) = \sum_{u,v \in V} \deg(u)\deg(v) = \left(\sum_{u \in V} \deg(u)\right)\left(\sum_{v \in V} \deg(v)\right) = |E|^2$$

edges. For each edge along the path $(u, v) \to (u, u)$, our result in lecture shows that the expected hitting time is linear w.r.t. the number of edges in the product graph, so $\mathcal{O}(m^2)$. There are $\mathcal{O}(n)$ edges in this path, so the total expected hitting time is $\mathcal{O}(nm^2)$, and we are done.

Solution to problem 3. Let π be the stationary distribution of this MC. Immediately, since every state has probability $1/2$ of going back to state 0,

$$\pi_0 = \frac{\sum_i \pi_i}{2} = \frac{1}{2}.$$

Further, the only way to get to state π_{i+1} is via π_i with probability $1/2$, so $\pi_{i+1} = \pi_i/2$ for each i, which implies $\pi_i = 2^{-(i+1)}$.

Solution to problem 4. Let Z_t denote the walk described by the problem. We define a Markovian coupling (X_t, Y_t) as follows.

- When $X_t \neq Y_t$, in each iteration, increment X_t by 1 w.p. $1/2$, and let it stay put w.p. $1/2$. Likewise and independently for Y_t. Modulo n if necessary.

- After $X_t = Y_t$, in each iteration, increment *both* X_t and Y_t by some c drawn uniformly from $\{0, 1\}$. Modulo n if necessary.

It is easy to verify that (X_t, Y_t) is a Markovian coupling w.r.t. the original lazy walk: in each iteration, $\mathbb{P}(X_{t+1} = X_t + 1) = \mathbb{P}(Y_{t+1} = Y_t + 1) = 1/2$. Let the difference $D_t = X_t - Y_t \pmod n$ be given. Observe that D_t is a bidirectional lazy random walk: $D_{t+1} = D_t + 1$ w.p. $1/4$, $D_{t+1} = D_t - 1$ w.p. $1/2$, and D_t stays put w.p. $1/2$. The coupling is attained when D_t reaches 0 (or n).

We showed in lecture that for adjacent vertices, the hitting time $h_{u,v} = \mathcal{O}(|E|)$, where in our graph $|E| = n$. Our mixing time $\tau_0 = \inf\{t \geqslant 0 : D_t = 0\}$ can be viewed as the hitting time of 0. Since the path from D_0 to 0 is at most length n, we see $\mathbb{E}[\tau_0] = \mathcal{O}(n^2)$.

The rest is a combination of the coupling lemma with Markov inequality. We know $\mathbb{P}(\tau_0 > t) \leqslant \mathbb{E}[\tau_0]/t = \mathcal{O}(n^2)/t$. Naïvely one can set $t = \mathcal{O}(n^2/\epsilon)$ to bound the probability $\mathbb{P}(\tau_0 > t) \leqslant \epsilon$. To attain the additional logarithmic bound I resorted to section 4.5 of [this] book by Levin and Peres. Specifically, setting $t = \mathcal{O}(n^2)$ gives $\mathbb{P}(\tau_0 > t) \leqslant 1/4$, and by their (4.36), $\tau(\epsilon) \leqslant \mathcal{O}(n^2)\log(1/\epsilon)$, as desired.

(I also missed the lecture on mixing times, but I didn't find a relevant identity in the notes, so I had to do some reference reading.)

Solution to problem 5. We imitate the Metropolis-Hastings algorithm mentioned in lecture, where we design a MC over all Δ-colorings of G with the following state transition logic:

Algorithm 1: MCMC improper graph colorings

State space: all Δ-colorings of G
Notations: C = a coloring; $C(v)$ = color of vertex v; $I(C)$ = number of improper edges of C

Current state: coloring C
Choose $v \in V$ uniformly at random
Choose new color $c' \in \{1, \cdots, \Delta\} \backslash \{C(v)\}$ uniformly at random
$C' \leftarrow$ old coloring C except $C(v) = c'$
Calculate $\Delta I = I(C') - I(C)$
State transition:
    with probability $A(C, C') = \min\{1, \lambda^{\Delta I}\}$: move to C'
    with probability $1 - A(C, C')$: stay at C.

To verify the claims, we just need to ensure that given any two colorings C, C', we have $\pi_C P(C, C') = \pi_{C'} P(C', C)$ where π_C is the probability of state C in the stationary distribution and $P(C, C')$ is the transition probability from C to C'. Observe the latter quantity is equal to the probability of proposing $C \to C'$ times that of accepting the proposal, $A(C, C')$. The proposal probability is uniformly equal to $1/(\Delta - 1)$, so the constraints reduce to

$$\pi_C A(C, C') = \pi_{C'} A(C', C) \iff \frac{\lambda^{I(C)}}{\lambda^{I(C')}} = \frac{\pi_C}{\pi_{C'}} = \frac{A(C', C)}{A(C, C')}.$$

But this is precisely how we constructed $A(C, C')$: if $\lambda^{I(C') - I(C)} \leqslant 1$ then $A(C', C)/A(C, C') = 1/\lambda^{I(C') - I(C)} = \lambda^{I(C)}/\lambda^{I(C')}$, and likewise for the other case. The proof is therefore complete!