Randomness and Computation 2022
Tutorial 6 (week 9)
Solution

1. **Sampling from Independent Sets of a graph**. A independent set (IS) of a graph $G = (V, E)$ is a set of vertices belonging to V for which no edge $e \in E$ connects two vertices of the set. We want to construct a Markov chain that allows to sample uniformly at random from the ensemble of independent sets of G, where G is assumed to be composed of a single connected component. Let's defined the following Markov chain X_t over the space Ω_{IS} of independent set (Section 12.6 of Mitzenmacher and Upfal). Let's X_0 be a trivial IS, for example the empty set. Then at each step we:

   - Select an edge $e = (u, v) \in E$ from the graph uniformly at random.
   - We proceed as follows:
     - (M1): with probability $1/3$ set $X_{t+1} = X_t - \{u, v\}$
     - (M2): with probability $1/3$ let $Y = (X_t - \{u\}) \cup \{v\}$. If Y is an IS, then $X_{t+1} = Y$; otherwise $X_{t+1} = X_t$ (M2).
     - (M3): with probability $1/3$ let $Y = (X_t - \{v\}) \cup \{u\}$. If Y is an IS, then $X_{t+1} = Y$; otherwise $X_{t+1} = X_t$ (M3).

   (a) Explain the effect of the three action that have $1/3$ probability in the MC.

   (b) Show that the MC is irreversible and aperiodic.

   (c) Use detail balance to show that it converges to the uniform distribution over Ω_{IS}.

   Solution

   (a) Lets enumerate the effect of each action:

   - M1: If the two vertices u and v of the selected edge do not belong to the current IS X_y, i.e, $\{u, c\} \notin X_t$ then $X_{t+1} = X_t$. The onlt other option is that one of the two belongs to X_t, as by definition two vertices connected by an edge can not belong to an IS. Therefore, if $\{u\} \in X_t$ we have $X_{t+1} = X_t - \{u\}$ and similarly for v.
   - M2: If $\{u, v\} \notin X_t$ we have $X_{t+1} = X_t + v$. If $u \in X_t$ and $v \notin X_t$ we get instead $X_{t+1} = X_t + v - u$. Otherwise, If $u \notin X_t$ and $v \in V_t$, we get $X_{t+1} = X_t$.
   - M3: If $\{u, v\} \notin X_t$ we have $X_{t+1} = X_t + u$. If $v \in X_t$ and $u \notin X_t$ we get instead $X_{t+1} = X_t + u - v$. Otherwise, If $v \notin X_t$ and $u \in V_t$, we get $X_{t+1} = X_t$.

   (b) It is always possible two connect two IS via a walk over the space Ω_{IS} using movement between neighboring IS that are different only in one exact vertex. This is possible because removing one vertex from an IS always generates another IS. So at worst there is a path from IS1 to IS2 passing via the empty set. The Markov chain below will always have a non-zero probability to produce the required steps, therefore there is always a non-zero probability of transitioning from IS1 to IS2. The MC is aperiodic as there is always a non-zero probability of remaining in the same IS that we started.

(c) Let assume that we have two independent set X and Y. Because the detailed balance condition $\pi_i P_{i,j} = \pi_j P_{j,i}$ concerns IS that are strictly different, we are interested in two scenarios: (Case I) where $|X| = |Y|$ but they differ in two vertices, which can only be reached via moves $M2$ and $M3$, and $|X| = |Y| \pm 1$.

- Case I: without lost of generality we can consider $X = Y - \nu + u$. In this scenario the transition from X to Y has non-zero probability and involves selecting the edge (u, ν) followed by selecting move $M2$, giving a probability of transition $P_{x,y} = 1/(3|E|)$. The transition from Y to X involves also selecting the edge (u, ν) but now the move $M3$, which has the same probability $1/(3|E|)$.
- Case II: Let's assume that $X = Y + u$, where the scenario $Y = X + u$ follows a similar argument. The transition from X to Y need the selection of one edge $e = (u, \nu)$ for which ν does not belong to X, followed by the move $M3$. Let's define $P(u) = \sum_{(u,\nu):\nu \notin X} P(e = (u, \nu))$, the transition probability reads $P_{x,y} = 1/3P(u)$. Similarly, the transition from Y to X can only happen selection an edge $e = (u, \nu)$ for which ν does not belong to X followed by move $M1$, which has also probability $1/3P(u)$.
  let's define the probability of selecting that edge being P_e that contains the vertex u

2. ***An s-t connectivity algorithm using $O(\log n)$ space.** Let a graph $G = (V, E)$ and two vertices s and t in G. Let $|V| = n$ and $|E| = m$. One can easily find a connection between s and t in linear time using standard breath-first search or depth-first search, however, requiring $\Omega(n)$ space. Consider the following algorithm:

- Start a random walk from s
- If the walk reaches t within $4n^3$ steps, return that there is a path.
- Otherwise return no path exists.

(a) Discuss why the algorithm needs only $O(\log n)$ space.

(b) Using the relation between hitting time and the stationary distribution, i.e., $h_{i,i} = 1/\pi(i)$, proof that $h_{x,y}$, i.e., the expected time to reach x from y, satisfies $h_{x,y} < 2m$.

(c) The cover of G is the maximum over all $\nu \in V$ of the expected time to visit all nodes starting from ν. Show that the cover time is bounded by $4nm$.

(d) Show that the algorithm only errs by returning there is no path when there is one with probability smaller than 1/2.

(e) Discuss how to reduce that probability to a very small constant δ.

Solution:

(a) The algorithm only needs to keep track of the current position and the number of step taken in the random walk, both can be done in time $O(\log n)$, where the number of steps is bounded by $4n^3$.

(b) Let $N(u)$ be the set of neighbors of vertex u in G and $h_{w,u}$ the expected time to reach u from w, We can write:

$$\frac{2|E|}{d(u)} = \frac{1}{\pi(u)} = h_{u,u} = \frac{1}{d(u)} \sum_{w \in N(u)} (1 + h_{w,u}), \tag{1}$$

where in the last equality we used the fact the the expectation of returning to u should be the expectation of the time to go to u from its neighbors w averaged over the probability of going to w from u as first step of the walk. Therefore,

$$2|E| = \sum_{w \in N(u)} (1 + h_{w,u}), \tag{2}$$

which easily implies $h_{w,u} < 2|E|$.

(c) Choose a spanning tree of G; that is, any subset of the edges that gives an acyclic graph connecting all vertices of G. There exist a cyclic tour on this spanning tree in which every edge is traversed once in each direction (ex: sequences of vertices passed through when doing depth-first search). Let $v_0, v_1, ..., v_{2|V|-2} = v_0$ be the sequence of vertices in the tour. The expected time to go over all this vertices in the tour is an upper-bound to the cover time. Hence, the cover time is bounded above by

$$\sum_{i=0}^{2|V|-3} h_{v_i, v_{i+1}} < (2|V| - 2)(2|E|) < 4|V||E|, \tag{3}$$

where we used the result in the previous sub-question in the first inequality.

(d) Remark that $4|V||E| = 4nm \leq 2n^3$ (as $m \leq n^2/2$). Using Markov inequality, it is easy to show that the probability of non finding the path on time $> 2h_{s,t} = 4n^3$ is smaller than 1/2. The choice of bounding the algorithm to $4n^3$ was done to exactly guarantee this probability or error $< 1/2$.

(e) Repeating the algorithm M times and only deciding to output "no path" when the M runs agree on a "no path"solution, we can decrease the failure probability from 1/2 to $1/2^M$, which allow us to decrease the failure probability to δ by choosing $M = \log(1/\delta)$.

Raul Garcia-Patron Sanchez