Complex Networks
Solutions 5

---

1. There is an easy way and a hard way to do this.

   The easy way is to note that the modified voter model is identical to the original voter model, but with contacts between nodes happening according to a Poisson process of rate λp. This is because the contacts that result in a node copying its neighbour are obtained by a Bernoulli(p) thinning of the Poisson process of all contacts. (The Poisson process of contacts between nodes in the same state should be thinned by the same parameter p for our claim to be justified. This isn't stated in the problem, but as such contacts don't have any effect, we may as well assume they are similarly thinned.)

   Hence, the probability of reaching a particular absorbing state, and the time to absorption, are the same as in the original voter model, but with λ replaced by λp. The probability of reaching n was k/n, which doesn't depend on λ, and so stays the same. The time to hit an absorbing state was $\frac{n}{\lambda}h(\frac{k}{n})$, which therefore changes to $\frac{n}{\lambda p}h(\frac{k}{n})$; here $h(\cdot)$ is the binary entropy function.

   The hard way to solve this problem is to repeat the analysis of the voter model from lectures, modifying the jump rates appropriately.

2. The contact rates are $q_{ij} = 1$ whenever $(i,j) \in E$, and zero otherwise (for $j \neq i$). Hence, the Q matrix is

$$Q = \begin{pmatrix} -3 & 1 & 1 & 1 \\ 1 & -2 & 1 & 0 \\ 1 & 1 & -3 & 1 \\ 1 & 0 & 1 & -2 \end{pmatrix}.$$

   As each of the columns sum to zero, it is easy to see that $\pi = \begin{pmatrix} 1 & 1 & 1 & 1 \end{pmatrix}/4$ solves the global balance equations $\pi Q = 0$. Alternatively, the local balance equations are $\pi_i \cdot 1 = \pi_j \cdot 1$ whenever $(i,j) \in E$, i.e., $\pi_i = \pi_j$ if there is an edge between i and j. Consequently, $\pi_i = \pi_j$ if there is a path between i and j. As the graph is connected, $\pi_i = \pi_j$ for all i and j. In other words, the uniform distribution solves the local balance equations, and hence is invariant. Moreover, π is the unique invariant distribution as it is the only probability vector satisfying either the global or local balance equations.

   Let $\mathbf{X}(t) \in \{0,1\}^V$ denote the vector of states of all nodes at time t, which evolves as a Markov process. By the results in the lecture notes, $M(t) = \pi \cdot \mathbf{X}(t)$ is a martingale, i.e.,

$$\mathbb{E}[M(t+\tau)|\mathbf{X}(s), s \leq t] = M(t) \text{ for all } \tau > 0 \text{ and all } t.$$

   Note that we need to condition on the past of $\mathbf{X}(t)$; it may not be enough just to condition on $M(s), s \leq t$.

   (More precisely, the definition of a martingale has to state what we are conditioning on - and the martingale has to be measurable with respect to that - but we are going to be less formal and avoid measure-theoretic terminology. However, you do need to have the right intuition about what to condition on in order to get the martingale property.)

Using the expression for π, this says that $M(t) = \sum_{i=1}^{4} X_i(t)/4$ is a martingale. The random time T at which consensus is reached is a stopping time, and the martingale is bounded (by 1), so we can use the optional stopping theorem to conclude that $\mathbb{E}M(T) = M(0)$. But $M(0) = 2/4$ as we start with two of the nodes in state 1, and so

$$\mathbb{E}M(T) = 1 \cdot \mathbb{P}(\mathbf{X}(T) = \mathbf{1}) + 0 \cdot \mathbb{P}(\mathbf{X}(T) = \mathbf{0}) = M(0) = \frac{1}{2}.$$

Hence, $\mathbb{P}(\mathbf{X}(T) = \mathbf{1}) = 1/2$ is the probability of reaching consensus on the all-1 state.

3. (a) Suppose that at time t the hub and $k-1$ leaves are in state 1, for some $k \in \{1, 2, \ldots, n-1\}$. Then, either one of the $n-k$ leaves in state 0 flips to state 1, which happens at aggregate rate $n-k$ (the total rate at which one of these $n-k$ leaves becomes active, since it then necessarily copies the hub), or the hub flips from state 1 to state 0, which happens at rate $(n-k)/(n-1)$ (the probability that, when the hub becomes active, it chooses one of the leaves in state 0 to copy). Hence, in this case,

$$\mathbb{E}[M(t+dt) - M(t) | (\mathbf{X}(u), u \le t)] = 1 \cdot (n-k)dt - (n-1) \cdot \frac{n-k}{n-1} dt = 0,$$

neglecting $o(dt)$ terms.

The analysis is identical if the hub is in state 0, since the two states are perfectly symmetrical. Hence, we have shown that the expectation of $M(t)$ remains constant so long as all the nodes aren't in the same state. But once all nodes are in the same state (0 or 1), they remain in that state for ever. Hence, $M(t)$ remains constant, deterministically, from then on; trivially, so does its expectation. Thus, we have shown that $M(t)$ is a martingale.

(b) For the given initial state, we have $M(0) = (n-1) + (k-1) = n + k - 2$. Let T be the random time to hit one of the two absorbing states, the all-0 or all-1 states. Then T is a stopping time. Moreover, $M(T) = (n-1) + (n-1)$ on the event that the all-1 state is reached at time T, while $M(T) = 0$ on the event that the all-0 state is reached. Hence,

$$\mathbb{E}[M(T)] = 2(n-1) \cdot \mathbb{P}_k(\mathbf{X}(T) = \mathbf{1}) + 0 \cdot \mathbb{P}_k(\mathbf{X}(T) = \mathbf{0}).$$

By the Optional Stopping Theorem, $\mathbb{E}[M(T)] = \mathbb{E}[M(0)] = M(0)$, i.e.,

$$2(n-1) \cdot \mathbb{P}_k(\mathbf{X}(T) = \mathbf{1}) = n + k - 2,$$

which implies that

$$\mathbb{P}_k(\mathbf{X}(T) = \mathbf{1}) = \frac{n+k-2}{2(n-1)}.$$

4. (a) Each leaf copies the hub at rate 1. The hub copies each leaf at rate $1/(n-1)$, adding up to rate 1 for activity. Hence, the random walk has rate 1 of moving from a leaf to the hub, and rate $1/(n-1)$ of moving from the hub to any given leaf.

(b) Let X_t denote the distance between the two random walks at time t. From the answer to the last part, we see that X_t is a Markov process with transition rates $q_{21} = 2$, $q_{12} = (n-2)/(n-1)$ and $q_{10} = 1 + 1/(n-1)$. All other off-diagonal transition rates are zero. Note that for X_t to go from 2 to 1 (the event whose rate is denoted q_{21}, we must have the two random walks at different leaves at time t, and one of them moving to the hub during the interval $(t, t+dt)$. As either of them could move at rate 1, the overall rate is 2. Likewise, X_t goes from 1 to 2 if one walk is at the hub, one at the leaf, and the one at the hub moves to another leaf; the rate for this is the rate that the walk at hub moves to any one of $n-2$ leaves (but not the one occupied by the other random walk). Finally, for X_t to go from 1 to 0, either the random walk at the leaf should move to the hub, or the one at the hub should move to the specific leaf occupied by the other walk.

(c) Let α_x denote the expected time for X_t to hit state 0 starting in state x. We will obtain simultaneous equations for the α_x and solve them. Clearly, $\alpha_0 = 0$. To compute α_2, we note that X_t spends a random Exp(2) time, with mean $1/2$, in state 2 before moving to state 1, after which it needs expected time α_1 to get to state 0. Hence,

$$\alpha_2 = \frac{1}{2} + \alpha_1. \tag{1}$$

To compute α_1, we note that the total jump rate out of this state is $q_{12} + q_{10} = 2$. So again, the mean time spent in this state is $1/2$. Then, with probability $(n-2)/2(n-1)$, the next state is 2, from which it takes expected time α_2 to get to zero, while with probability $n/2(n-1)$, the next state is 0 and it takes $\alpha_0 = 0$ time to get to state 0 from there. Hence,

$$\alpha_1 = \frac{1}{2} + \frac{n-2}{2(n-1)}\alpha_2 + \frac{n}{2(n-1)}\alpha_0 = \frac{1}{2} + \frac{n-2}{2(n-1)}\alpha_2. \tag{2}$$

Solving (1) and (2) simultaneously, we get $\alpha_2 = 2(n-1)/n = 2 - (2/n)$ and $\alpha_1 = (3/2) - (2/n)$.

(d) We saw in the last part that the mean time for two random walks to meet is $2 - (2/n)$ if they start at different leaves, and $(3/2) - (2/n)$ if one starts at the hub, and the other at a leaf. Following the hint, this gives an upper bound of $2(n-1) - 2(n-1)/n$ $= 2(n-1)^2/n$ for all walks to coalesce (if you add up the worst case bounds). Or you could argue that it is enough if the walk started at each leaf coalesces with the one started at the hub, which leads to the bound $3(n-1)/2 - 2(n-1)/n$. Either of these is an acceptable answer as are the approximations $2n$ and $3n/2$, neglecting terms of smaller order in n.

Remark. The bounding technique used in this homework problem was for illustration. It yields a very loose bound. A more careful analysis, using Chernoff bounds, yields an upper bound on the time to consensus which is logarithmic in n rather than linear in n.

5. (a) Suppose Y_t lies between 1 and $n-1$. Note that each particle moves at rate 1, and moves left or right with equal probability. Hence, $Y_{t+dt} - Y_t = \pm 1$ with equal probability $1 \cdot dt$ each. Hence,

$$\mathbb{E}[Y_{t+dt}^2 - Y_t^2 | Y_t] = \Big((Y_t+1)^2 - Y_t^2\Big)dt + \Big((Y_t-1)^2 - Y_t^2\Big)dt = 2dt,$$

and so $\mathbb{E}[Y_{t+dt}^2 - 2(t+dt) - Y_t^2 + 2t] = 0$, which establishes that $Y_t^2 - 2t$ is a martingale.

(b) Note that T is a stopping time as we can determine whether $T \le t$ by observing Y_t only on $0 \le s \le t$. However, the martingale $Y_t^2 - 2t$ is not bounded; it is bounded above by n^2, but not bounded below. The statement of the optional stopping theorem that we saw in lectures is not strong enough for our purposes in this problem, but it turns out that the OST holds under weaker conditions than we stated in lectures, and is applicable.

Using the optional stopping theorem, we get $\mathbb{E}[Y_T^2 - 2T] = Y_0^2 - 0 = |i-j|^2$. But Y_T is either 0 or n by definition of T, so we also have

$$\begin{aligned}\mathbb{E}[Y_T^2 - 2T] &= \mathbb{E}[Y_T^2] - 2\mathbb{E}[Y_T] = n^2\mathbb{P}(Y_T = n) + 0 \cdot \mathbb{P}(Y_T = 0) - 2\mathbb{E}[T] \\ &= n^2\mathbb{P}(Y_T = n) - 2\mathbb{E}[T].\end{aligned}$$

Equating this to $|i-j|^2$, and using the given expression, $\mathbb{P}(Y_T = n) = |i-j|/n$, we get $\mathbb{E}[T] = |i-j|(n-|i-j|)/2$.

(c) For $i = 2, 3, \ldots, n$, let T_i denote the time until particle i has merger with particle 1 (i and 1 could already be component parts of merged particles at this time). Then, the time at which all particles have merged and there is just one partice left is $\max_{i=2}^n T_i$. Clearly, the T_i are not independent. Nevertheless, using the hint in the question and linearity of expectation, we have,

$$\mathbb{E}\big[\max_{i=2}^n T_i\big] \le \mathbb{E}\big[\sum_{i=2}^n T_i\big] = \sum_{i=2}^n \mathbb{E}[T_i].$$

Now, from the answer to the last part, $E[T_i] = k(n-k)/2$ if the initial distance between particles 1 and i is k. Now, both particle $k+1$ and particle $n+1-k$ are at distance k from particle 1, and the largest possible distance is $n/2$. Hence, we can rewrite the above as

$$\mathbb{E}\big[\max_{i=2}^n T_i\big] \le 2\sum_{k=1}^{n/2} \frac{k(n-k)}{2} \approx \int_0^{n/2} x(n-x)dx = \frac{n^3}{8} - \frac{n^3}{24} = \frac{n^3}{12}.$$