ETH Zurich
Institute of Theoretical Computer Science
Dr. Karl Bringmann, Dr. Johannes Lengler, and Dr. Tsur Luria
Frank Mousset, Felix Weissenberger

FS 2015
Solution GHW 2

Randomized Algorithms and Probabilistic Methods: Advanced Topics

Solution to Exercise 1

Let X be a set of size n and let Ω be the set of all subsets of X of size k, where k is a fixed constant. We consider the symmetric (lazy) random walk on the graph G with vertex set Ω where AB is an edge if and only if the size of the symmetric difference between A and B is exactly two. We will use $\triangle$ for the symmetric difference of sets, i.e., $A \triangle B = (A \setminus B) \cup (B \setminus A)$.

To bound the mixing time t_{mix}, we give a lower bound for the edge expansion

$$h(G) = \min_{0<|S|<|\Omega|/2} \frac{e(S, \overline{S})}{|S|}.$$

For this, we use the canonical paths method. Thus, we should start by defining paths between any two elements $A, B \in \Omega$. It turns out that the exact definition of the paths is not really all that important; the only important thing is that on the path from A to B, we never leave the set $A \cup B$. Clearly, there are many ways of definiting paths that satisfy this condition. Assume now that we fix such a system of paths.

Consider any edge $A'B' \in E(G)$ and count the number of pairs (A, B) for which the corresponding A-B-path uses the edge $A'B'$. We first note that by the defining property of our paths, the union $A' \cup B'$ must be a subset of $A \cup B$. Since $|A' \triangle B'| = 2$ and $|A'| = |B'| = k$, we have

$$|A' \cup B'| = k + 1.$$

On the other hand, clearly $|A \cup B| \leq |A| + |B| = 2k$. Then we can count the pairs AB that use the edge $A'B'$ as follows. First, we colour the elements of $A' \cup B'$ red, blue, or green, depending on whether the element is in $A \setminus B$, $B \setminus A$ or $B \cap A$. Note that there are 3^{k+1} such colourings. Fix any such colouring and let r, b, g be the numbers of red, blue and green elements, respectively. Then the number of choices for A that are consistent with the colouring of $A' \cup B'$ is at most n^{k-r-g} (one needs to choose the $k - r - g$ elements that are not in $A \cap (A' \cup B')$), and, similarly, the number of consistent choices for B is at most n^{k-b-g}. All in all, for each colouring, the number of coices for (A, B) that are consistent with the colouring is at most $n^{2k-r-b-2g}$. However, since $r + b + g = |A' \cup B'| = k + 1$, we have $n^{2k-r-b-2g} = n^{k-1-g} \leq n^{k-1}$. Therefore, there are at most $3^{k+1} n^{k-1}$ pairs (A, B) that use the edge $A'B'$.

From this, we obtan a lower bound on $h(G)$ as usual. If $|S| \leq |\Omega|/2 = \binom{n}{k}/2$, then

$$e(S, \overline{S}) \cdot 3^{k+1} n^{k-1} \geq |S||\overline{S}| \geq |S| \cdot \binom{n}{k}/2,$$

so we have

$$h(G) \geq \frac{\binom{n}{k}}{2 \cdot 3^{k+1} n^{k-1}} = \Omega(n).$$

Using the result from the lecture, and as G is d-regular with $d \leq n$, we obtain

$$t_{\text{mix}} \leq \frac{4d^2 \log(2\sqrt{|\Omega|})}{h(G)^2} = O(\log(\sqrt{|\Omega|})) = O(\log n).$$

Solution to Exercise 2

Let X and Y be two configurations. We construct a path from X to Y.

We think of our board as having k rows and n columns. Instead of moving directly from X to Y, we move first to nearby configurations X' and Y' with a single king in each row. We build a path in three stages:

- Stage 1: Using only vertical moves, move from X to a configuration X' such that each row contains a single king.
- Stages 2 to $k+1$: Using only horizontal moves, move the king in the first row from its column in X' to its column in Y'. Then move the second king from its column in X' to its column in Y'. Continue in this fashion until all of the kings have been moved to their column in Y'.
- Stage $k+2$: Using only vertical moves, move from Y' to Y.

We will assume for now that that Stages 1 and $k+2$ can be accomplished, i.e., that it is possible to move from X to X' and from Y' to Y using only vertical moves.

Given this, we can bound the maximal flow on an edge as follows. If the edge corresponds to a vertical move, then the paths only use it in the first or last stage. In the first stage, the number of X's that use it is at most k^k, since each king must be in the same column in X as it is in the endpoints of the edge, but it can be in one of k rows. Y can be anything, so the number of paths is at most $k^k \cdot |\Omega|$. The last stage is the same, so the load of the horizontal moves is at most $2k^k \cdot |\Omega|$.

For a horizontal move e, assume that we are moving a king in the i-th row. Then the first $i-1$ kings have already been moved, and the last $k-i$ kings haven't been moved yet.

Consider a pair (X', Y') that uses e. How many possibilities are there for X'? Well, we know the locations of the last $k-i$ kings, and the first i kings could be anywhere (so long as there is a single one in each row), so there are n^i possibilities. Similarly, there are n^{k-i+1} possiblities for Y', so the number of pairs (X', Y') using e is at most $n^{k+1} < n \cdot \binom{nk}{k} \leq n \cdot |\Omega|$. Given X' and Y', there are again at most k^k possibilities for X and the same for Y, so the total number of paths using e is at most $n^{k+1} \cdot k^{2k}$.

Therefore, by Theorem 2.7, $h(G) \geq 1/(2k^{2k} \cdot n) \in \Omega(n^{-1})$ since k is a constant. Finally, by the hint from the exercise

$$t_{mix} \in \mathcal{O}(n^2 \log n).$$

It remains to prove that Stages 1 and $k+2$ are possible. We will focus on Stage 1; Stage $k+2$ is the same. We can prove this by induction on k.

For $k=1$ there is nothing to prove. In general, we first move the lowest king down to the bottom row (unless it is already there). On the other hand, if there are multiple kings on the bottom row, we make a series of moves to move all but one of them upwards, so that there is only one king on the bottom row. That is, if there is a "pile" of kings of size r on the bottom r places in some column, we move the top king in the pile up one square, and then the next one and so on until all of the pile has moved up one square.

Now the bottom row has a single king, and the top $k-1$ rows have $k-1$ kings, so we can use the induction hypothesis.

Solution to Exercise 3

(a) The Markov chain is connected and it has an odd and an even cycle, so it is ergodic.

(b) This works like in the hypercube. Contrary to what is stated in the exercise sheet, we need to assume that the random walk is lazy. The paths are obtained by going through the coordinates from left to right and fixing one coordinate at a time. Consider an edge

$$(x_1, \ldots, x_n) \mapsto (x_1, \ldots, x_{i-1}, y_i, x_{i+1}, \ldots, x_n),$$

where $y_i = x_i \pm 1$. This edge is used only by pairs $(s,t) \in \Omega^2$ such that for all $j < i$, we have $t_j = x_j$, and for all $j > i$, we have $s_j = x_j$. Therefore, the number of pairs using the edge is at most $k^{i+1} \cdot k^{n-i+1} = k^{n+2}$. From this, we obtain a lower bound for the Cheeger constant:

$$h(G) \geq |\Omega|/(2k^{n+2}) = 1/(2k^2).$$

By the result from the lecture (for lazy random walks on d-regular graphs),

$$t_{\text{mix}} \leq \frac{4d^2 \log(2\sqrt{|\Omega|})}{h(G)^2} = O(n^3).$$

(c) The eigenvalues of C_k (the cycle on k vertices) are $2\cos(2r\pi/k)$, $0 < r < k$. The easiest way to see this is to write the adjacency matrix A_1 of C_k as the sum $A_1 = A + A^T$, where

$$A = \begin{bmatrix} 0 & 1 & 0 & \cdots & 0 & 0 \\ 0 & 0 & 1 & \cdots & 0 & 0 \\ & & & \ddots & & \\ 0 & 0 & 0 & \cdots & 0 & 1 \\ 1 & 0 & 0 & \cdots & 0 & 0 \end{bmatrix}.$$

The characteristic polynomial of A is

$$p_A(\lambda) = \det(A - \lambda I) = (-\lambda)^k - (-1)^k,$$

which is zero if and only if $\lambda^k = 1$. Thus, the eigenvalues of A are of the form ω^r, where $0 \le r < k$ and where $\omega = e^{2\pi i/k}$. Since multiplication with A corresponds to a cyclic left-shift of the input vector, the eigenvector associated with ω^r must be $(1, \omega^r, \omega^{2r}, \ldots, \omega^{kr})$. Similarly, for A^T, the eigenvalues are also ω^r, with associated eigenvector $(1, \omega^{-r}, \omega^{-2r}, \ldots, \omega^{-kr})$. Therefore, the eigenvectors of $A_1 = A + A^T$ are exactly the vectors

$$(1, \omega^r, \omega^{2r}, \ldots, \omega^{kr})$$

with associated eigenvalues $\omega^r + \omega^{-r} = 2\cos(2r\pi/k)$.

Now note that the adjacency matrix of C_k^n (the graph on which our random walk takes place) has the recursive definition

$$A_n = I_k \otimes A_{n-1} + A_1 \otimes I_{k^{n-1}},$$

where $\otimes$ denotes the Kronecker product of matrices (note that A_n has dimensions $k^n \times k^n$). Let v be an eigenvector of A_{n-1} with eigenvalue λ and let u be an eigenvector of A_1 with eigenvalue μ. Then $u \otimes v$ is an eigenvector of A_n with eigenvalue $\lambda + \mu$ (check this). Moreover, these are all eigenvalues, since if $v_1, \ldots, v_{k^{n-1}}$ is an eigenbasis for A_{n-1} and $u_1, \ldots, u_k$ is an eigenbasis for A_1, then the products $u_i \otimes v_j$ are linearly independent: assuming

$$\sum \lambda_{ij}(u_i \otimes v_j) = 0,$$

we have in particular $\sum \lambda_{ij} u_{ik} v_j = 0$ for all k, which, by independence of the v_j implies $\sum \lambda_{ij} u_{ik} = 0$ for all k. But then $\sum \lambda_{ij} u_i = 0$, which by independence of the u_i, implies $\sum \lambda_{ij} = 0$.

By induction, we thus see that the eigenvalues of A_n are all sums of the form

$$2\cos(2r_1\pi/k) + 2\cos(2r_2\pi/k) + \cdots + 2\cos(2r_n\pi/k),$$

where $0 < r_i < k$.

The largest eigenvalue is $2n$, the second largest is $2(n-1) + 2\cos(2\pi/k)$, and the smallest is $2n\cos(2(k-1)\pi/k)$. One sees that of the latter two, the one with largest absolute value is the first, provided n is large enough, i.e.,

$$2n - \max\{|\lambda_2|, |\lambda_{k^n}|\} = 2 - 2\cos(2\pi/k).$$

Every vertex has degree $2n$, so using the result from the lecture, it follows that

$$t_{\mathrm{mix}}(1/4) \le \frac{2n\log(\sqrt{2|\Omega|})}{2 - 2\cos(2\pi/k)} = \frac{2\log(\sqrt{2})n + n^2 \log k}{2 - 2\cos(2\pi/k)}.$$