

CME 305: Discrete Mathematics and Algorithms

Instructor: Reza Zadeh (rezab@stanford.edu)

HW#1 – Due at the beginning of class Thursday 01/21/16

1. Prove that at least one of G and \overline{G} is connected. Here, \overline{G} is a graph on the vertices of G such that two vertices are adjacent in \overline{G} if and only if they are not adjacent in G .

Solution: Let G be a disconnected graph in which case we can decompose it into k connected components C_1, C_2, \dots, C_k . We want to show that \overline{G} is connected i.e. there is a path between any u and v in \overline{G} . In the case that u and v are in different components we know that there exist an edge (a path of length one) between them in \overline{G} . In the case that u and v are in the same component, say C_i , we can construct a path of two edges between them in \overline{G} as follows. Pick any vertex w from some other component C_j for $j \neq i$ and note that edges $\{u, w\}$ and $\{w, v\}$ are in \overline{G} . Thus u, w, v is a path in \overline{G} and hence \overline{G} is connected.

2. A vertex in G is *central* if its greatest distance from any other vertex is as small as possible. This distance is the *radius* of G .

- (a) Prove that for every graph G

$$\text{rad } G \leq \text{diam } G \leq 2 \text{ rad } G$$

Solution: Since the diameter is the longest shortest path in the graph, and the radius is just a particular shortest path, we have $\text{rad } G \leq \text{diam } G$. Now, since we can always reach any vertex t by going to the center first, then going to t , incurring a cost of at most twice the radius, we have $\text{diam } G \leq 2 \text{ rad } G$.

- (b) Prove that a graph G of radius at most k and maximum degree at most $d \geq 3$ has fewer than $\frac{d}{d-2}(d-1)^k$ vertices.

Solution: Let z be a central vertex in G , and let D_i denote the set of vertices of G at distance i from z . Then $\cup_{i=0}^k D_i$ is all the vertices in the graph. Clearly, $|D_0| = 1$ and $|D_1| \leq d$. For $i \geq 1$ we have $|D_{i+1}| \leq (d-1)|D_i|$, because every vertex in D_{i+1} is a neighbor of a vertex in D_i (why?), and each vertex in D_i has at most $d-1$ neighbors in D_{i+1} (since it has another neighbor in D_{i-1}). Thus $|D_{i+1}| \leq d(d-1)^i$ for all $i < k$ by induction, giving

$$|G| \leq 1 + d \sum_{i=0}^{k-1} (d-1)^i = 1 + \frac{d}{d-2}((d-1)^k - 1) < \frac{d}{d-2}(d-1)^k$$

3. A random permutation π of the set $\{1, 2, \dots, n\}$ can be represented by a directed graph on n vertices with a directed arc (i, π_i) , where π_i is the i 'th entry in the permutation. Observe that the resulting graph is just a collection of distinct cycles.

- (a) What is the expected length of the cycle containing vertex 1?

Solution: Consider the construction of the directed graph where we start at vertex 1. Each step we select an unmarked vertex at random and move to that vertex. We then mark that vertex before repeating the process. Once this construction marks vertex 1 we have a cycle. Let $|C|$ denote the length of this cycle. Then:

$$\begin{aligned}
 E(|C|) &= 1 \times \frac{1}{n} \\
 &+ 2 \times \frac{n-1}{n} \frac{1}{n-1} \\
 &+ \dots \\
 &+ n \times \frac{n-1}{n} \frac{n-2}{n-1} \dots \frac{1}{2} \frac{1}{1} \\
 &= \sum_{i=1}^n i \frac{1}{n} \\
 &= \frac{1}{n} \frac{n(n+1)}{2} \\
 &= \frac{1}{2}(n+1)
 \end{aligned}$$

(b) What is the expected number of cycles?

Solution: Let $f(n)$ denote the expected number of cycles in a graph on n nodes. It is clear that $f(1) = 1$.

Consider $f(n)$ given $f(n-1)$. With probability $1/n$ the new node permutes to itself resulting in an expected number of cycles of $1 + f(n-1)$ and with the remaining probability the new node permutes to a node other than itself, this case then reduces to the $n-1$ case. Hence:

$$f(n) = \frac{1}{n} (1 + f(n-1)) + \frac{n-1}{n} f(n-1) = \frac{1}{n} + f(n-1)$$

It follows recursively that $f(n) = \sum_{i=1}^n 1/i$ which is exactly equal to the n th harmonic number $H(n)$.

4. Let v_1, v_2, \dots, v_n be unit vectors in \mathbb{R}^n . Prove that there exist $\alpha_1, \alpha_2, \dots, \alpha_n \in \{-1, 1\}$ such that

$$\|\alpha_1 v_1 + \alpha_2 v_2 + \dots + \alpha_n v_n\|_2 \leq \sqrt{n}$$

Solution: This can be shown geometrically using similar ideas to Pythagoras' Theorem. Consider the cosine rule:

$$\|\alpha_i v_i + \alpha_j v_j\|_2^2 = \|\alpha_i v_i\|_2^2 + \|\alpha_j v_j\|_2^2 - 2\|\alpha_i v_i\|_2 \|\alpha_j v_j\|_2 \cos \theta$$

Where θ is the angle between $\alpha_i v_i$ and $\alpha_j v_j$.

Fix α_i . Then we can choose α_j such that $\theta \leq \pi/2$. Hence $\cos \theta \in [0, 1]$. It follows that, given α_i we can choose α_j such that:

$$\|\alpha_i v_i + \alpha_j v_j\|_2^2 \leq \|\alpha_i v_i\|_2^2 + \|\alpha_j v_j\|_2^2$$

Applying this result recursively gives:

$$\begin{aligned} \|\alpha_1 v_1 + \alpha_2 v_2 + \dots + \alpha_n v_n\|_2^2 &\leq \|\alpha_1 v_1\|_2^2 + \|\alpha_2 v_2\|_2^2 + \dots + \|\alpha_n v_n\|_2^2 \\ &= \alpha_1^2 \|v_1\|_2^2 + \alpha_2^2 \|v_2\|_2^2 + \dots + \alpha_n^2 \|v_n\|_2^2 \\ &= n \\ \|\alpha_1 v_1 + \alpha_2 v_2 + \dots + \alpha_n v_n\|_2 &\leq \sqrt{n} \end{aligned}$$

5. Consider a graph G on $2n$ vertices where every vertex has degree at least n . Prove that G contains a perfect matching.

Solution: We will prove this in a slightly roundabout way: we first show that G must contain a Hamiltonian path, and then note that a Hamiltonian path contains our desired perfect matching. (A Hamiltonian path is a path which contains every node of the graph.) A direct proof of this is possible, but this proof is shorter and more elegant.

Consider the longest path $P = (v_1, v_2, \dots, v_k)$ in G . All neighbors w of v_1 must be elements of P , otherwise the longer path (w, v_1, \dots, v_k) in G would contradict the definition of P . Similarly, all neighbors of v_k must also be in P . Now since G is simple, we note that all $\geq n$ neighbors of v_1 must be distinct and lie in P , thus we have the bounds $n + 1 \leq k \leq 2n$ on the length of P , where in the lower bound we have also accounted for v_1 itself.

We claim that there exists $j \in \{1, \dots, k - 1\}$ such that v_j and v_{j+1} are neighbors of v_k and v_1 respectively. Suppose for contradiction that this is not the case. Let $S = \{v_i | v_k \sim v_i\}$ be the set of all neighbors of v_k in P . Let $T = \{v_{i-1} | v_1 \sim v_i\}$ be the set of all path vertices immediately preceding the neighbors of v_1 in P . Note that S and T are disjoint by our assumption. Since v_1 and v_k have at least n neighbors in P , we have

$$|S| + |T| \geq n + n = 2n \geq k.$$

But we also know that both S and T are subsets of $\{v_1, \dots, v_{k-1}\}$ so $|S \cup T| = |S| + |T| \leq k - 1$, a contradiction. Thus there exists j such that $v_1 \sim v_{j+1}$ and $v_k \sim v_j$.

Then we have the cycle $C = (v_1, v_2, \dots, v_j, v_k, v_{k-1}, \dots, v_{j+1}, v_1)$ in G which contains each vertex of P exactly once. Now we claim that $k = 2n$; in that case C contains our desired Hamiltonian path of G . To see this consider a vertex $x \notin C$. Since G is simple and $|C| = k \geq n + 1$, one of the n neighbors of x , call it y must lie in C . But then cutting either one of the edges in C incident to y and including the edge $\{x, y\}$ would result in a path longer than P , contradicting our original longest path assumption. Thus, our cycle C must have length $2n$ — it must contain a Hamiltonian path.

Finally, we prove that the Hamiltonian path found above contains a perfect matching. Let the Hamiltonian path P be $(v_1, v_2, \dots, v_{2n})$. Choose the edges $(v_1, v_2), (v_3, v_4), \dots, (v_{2i-1}, v_{2i}), \dots, (v_{2n-1}, v_{2n})$. These edges are all in the Hamiltonian path, and every node in the path is present in exactly one of these edges. As the path contains every vertex in the graph, each node of the graph is the endpoint of exactly one of the edges. Thus, this is a perfect matching in the graph, as desired.

6. Let $G = (V, E)$ be a graph and $w : E \rightarrow R^+$ be an assignment of nonnegative weights to its edges. For $u, v \in V$ let $f(u, v)$ denote the weight of a minimum $u - v$ cut in G .

(a) Let $u, v, w \in V$, and suppose $f(u, v) \leq f(u, w) \leq f(v, w)$. Show that $f(u, v) = f(u, w)$, i.e., the two smaller numbers are equal.

Solution: Let $c = \min(f(u, w), f(w, v))$. Consider the two ends of the smallest path between u and v . We can route c units of flow from u to w and then from w to v . This means $f(u, v) \geq c = \min(f(u, w), f(w, v)) = f(u, w)$, which is only possible if $f(u, v) = f(u, w)$.

(b) Show that among the $\binom{n}{2}$ values $f(u, v)$, for all pairs $u, v \in V$, there are at most $n - 1$ distinct values.

Solution: We prove this by induction on the number of nodes. The result is clearly true for a graph with 3 nodes from part a. Assume the result for all graphs G' of size n , and consider a graph G with $n + 1$ nodes. There will be a largest edge, pick one of its two vertices, call it v . Order the edges incident upon v in decreasing order: f_1, f_2, \dots, f_n . So f_1 is the largest edge in G . Note that the f_i are sides of triangles of all whom have one edge in the smaller graph G' , where there are only $n - 2$ distinct edges by induction hypothesis. We argue that other than f_1 , all the other f_i are equal to some edge in G' , thus the number of distinct edges in G can only one larger, with the contribution coming from f_1 . This is true because of the decreasing ordering on the f_i 's and the triangle property from part a, enforcing each f_2, \dots, f_n be equal to some edge in G' . Thus the addition of v can only add one new distinct edge weight: f_1 , making for at most $n - 1$ distinct weights.

7. Let T be a spanning tree of a graph G with an edge cost function c . We say that T has the *cycle property* if for any edge $e' \notin T$, $c(e') \geq c(e)$ for all e in the cycle generated by adding e' to T . Also, T has the *cut property* if for any edge $e \in T$, $c(e) \leq c(e')$ for all e' in the cut defined by e . Show that the following three statements are equivalent:

- (a) T has the cycle property.
- (b) T has the cut property.
- (c) T is a minimum cost spanning tree.

Remark 1: Note that removing $e \in T$ creates two trees with vertex sets V_1 and V_2 . A *cut* defined by $e \in T$ is the set of edges of G with one endpoint in V_1 and the other in V_2 (with the exception of e itself).

Solution: In order to show that (a), (b), and (c) are equivalent, it is enough to show: (a) \Leftrightarrow (c), and (b) \Leftrightarrow (a).

(c) \Rightarrow (a): By contradiction suppose T does not have the cycle property: there exists $e' \notin T$ such that $T \cup \{e'\}$ has a cycle C in which there exists $e \in T$ and $e \in C$ where $c(e) > c(e')$. Let tree T' be the tree obtained by adding e' to T and removing e ; T' is a tree with cost strictly less than cost of T which is contradicting with T being an MST.

(a) \Rightarrow (c): By contradiction suppose T is not an MST: let e' be the first edge that was picked by Kruskal's algorithm but does not belong to T . Adding e' to T would create cycle C . Since T has cycle property, $c(e) \leq c(e')$, $e \in C$. Therefore, all $e \in C$, $e \neq e'$ have been visited by the Kruskal's algorithm. We have two cases:

case 1: All $e \in C \setminus \{e'\}$ were picked by the algorithm. In this case the algorithm would not pick e' because it creates a cycle with the existing edges.

case 2: There exists $e^* \in C$, $e^* \neq e'$ such that it was not picked by the algorithm. The reason for not picking an edge is that it would create a cycle with the existing edges. However, since e' was the first edge picked by the algorithm that does to belong to T , this would mean that T has a cycle, which is a contradiction.

(a) \Rightarrow (b): By contradiction suppose T does not have the cut property: there exists $e \in T$ such that there exists edge $e' = (v_1, v_2)$ such that $v_1 \in V_1$ and $v_2 \in V_2$ (V_1 and V_2 are the set of vertices of the two connected components after removing e , see Remark 1.), and $c(e') < c(e)$. Since T is connected graph there exist path P_T between v_1 and v_2 such that all the edges of P_T belong to T . Adding e' to P_T creates a cycle in which there exist $e \in T$ where $c(e) > c(e')$ which is contradicting with T having the cycle property.

(b) \Rightarrow (a): By contradiction suppose T does not have the cycle property: there exists edge e' such that when adding it to T and creating cycle C , there exists $e \in C$, where $c(e) > c(e')$. In T , if we remove e , we have two connected components with vertex sets V_1 and V_2 . Let v_1, v_2 be the endpoints of e , where $v_1 \in V_1$ and $v_2 \in V_2$. Since $e \in C$ there exists another path between v_1 and v_2 therefore at least one edge from C belongs to cut (V_1, V_2) ; since all the edges of C belong to T except e' and $T \cap \text{cut}(V_1, V_2) = \emptyset$, e' should belong to $\text{cut}(V_1, V_2)$. However, $c(e') < c(e)$, which contradicts with T having the cut property.

8. Prove that there is an absolute constant $c > 0$ with the following property. Let A be an $n \times n$ matrix with pairwise distinct entries. Then there is a permutation of the rows of A so that no column in the permuted matrix contains an increasing subsequence of length $c\sqrt{n}$.

Solution: Let $A = (a_{ij})_{n \times n}$ be a matrix with pairwise distinct entries. Let $\pi \in G_n$ be a random permutation of $[n]$. For ever $c\sqrt{n}$ -tuple $S = (s_1, s_2, \dots, s_{c\sqrt{n}})$ of elements of $[n]$, let X_S be the event that the sequence $(a_{1,s_1}, a_{1,s_2}, \dots, a_{1,s_{c\sqrt{n}}})$ is increasing. Since all entries of A are distinct, there is exactly one tuple S for which the sequence formed is increasing. Then $P[X_S] = 1/(c\sqrt{n})!$. Let Y_S be the indicator random variable for

event X_S , and let

$$Y = \sum_{S \subseteq [n], |S|=c\sqrt{n}} Y_S$$

be the number of increasing sub-sequences of length $c\sqrt{n}$ in the first column. By linearity of expectations:

$$\begin{aligned} E[Y] &= \sum_{S \subseteq [n], |S|=c\sqrt{n}} E[Y_S] \\ &= \sum_{S \subseteq [n], |S|=c\sqrt{n}} P(X_S) \\ &= \sum_{S \subseteq [n], |S|=c\sqrt{n}} \frac{1}{(c\sqrt{n})!} \\ &= \binom{n}{c\sqrt{n}} \frac{1}{(c\sqrt{n})!} \end{aligned}$$

Again by linearity of expectation, the expected number of increasing subsequences of length $c\sqrt{n}$ in the columns of A is upper bounded by $n \frac{\binom{n}{c\sqrt{n}}}{(c\sqrt{n})!}$. This is less than

$$n \frac{\left(\frac{ne}{c\sqrt{n}}\right)^{c\sqrt{n}}}{\left(\frac{c\sqrt{n}}{e}\right)^{c\sqrt{n}}} = n \left(\frac{e^2}{c^2}\right)^{c\sqrt{n}}$$

by Stirling's approximation.

For any $c > e$, and any sufficiently large value of n , such an upper bound is less than 1, implying that $E[Y] < 1$. Hence there is a permutation that has less than $E[Y] < 1$ increasing subsequences of length $c\sqrt{n}$ in any column. I.e.: it must have zero increasing subsequences of size $c\sqrt{n}$.

9. At lunchtime it is crucial for people to get to the food trucks as quickly as possible. The building is represented by a graph $G = (V, E)$, where each room, landing, or other location is represented by a vertex and each corridor or stairway is represented by an edge. Each corridor has an associated capacity c , meaning that at most c people can pass through the corridor at once. Traversing a corridor from one end to the other takes one timestep and people can decide to stay in a room for the entire timestep. Suppose all people are initially in a single room s , and that the building has a single exit t . Give a polynomial time algorithm to find the fastest way to get everyone out of the building.

Solution: We solve this problem by modifying this graph, and then applying the Ford-Fulkerson algorithm. First, fix a "deadline" time k , such that we need to empty the building by time k , under the conditions above. We will develop a polynomial time algorithm for this problem. Note that once we have this algorithm we can solve our problem by increasing k one by one until we find the smallest value such that everyone can move out. Let the total number of people equal P .

We modify our graph as follows. For each room $r \in V$, make $k+1$ copies of it r_0, r_1, \dots, r_k . In our modified graph, these will correspond to the room r at each individual time step. For each pair of rooms r and w with a hallway of capacity c_{rw} between them, create the directed edges with weight c_{rw} each $r_i \rightarrow w_{i+1}$ and $w_i \rightarrow r_{i+1}$, for each time $i = 0, 1, 2, \dots, k-1$. These edges correspond to the act of moving from room r to w (or from w to r) in some timestep. For each room r , add the directed edges with weight P $r_i \rightarrow r_{i+1}$ for each time $i = 0, 1, 2, \dots, k-1$. Finally, add a single room t , and connect t_k to t with an edge of weight P . These edges correspond to the act of staying in a room for a single timestep. On this modified graph, run the Ford-Fulkerson algorithm with source w_0 and target t . We first show that this algorithm will give us an evacuation plan if one exists.

First, we prove that if a plan existed which evacuated the building in k timesteps, the scheme above will return a flow of size P or more. This is fairly obvious: if we route units of flow as we route people in such a scheme, and then send all P units of flow from t_k to t , we will clearly obtain a flow of size P . As Ford-Fulkerson finds the maximum flow of the graph, it must find a flow of greater size than this, and so Ford-Fulkerson finds a flow of size greater than or equal to P in our graph, as desired.

Now, we prove that if Ford-Fulkerson finds a flow of size P or greater, we can convert this to a valid evacuation plan to get everyone out of the building in k timesteps. Note that the flow in this case must have value exactly P , since the cut defined by t has size P . We convert the obtained flow as follows. For every timestep $i = 1 \dots k$, look for all pairs of rooms r and w such that people are moving from r_{i-1} to w_i and from w_{i-1} to r_i . The problem with this is that although no more than c_{rw} people are crossing in either direction, the total number of people utilizing the hallway between r and w may be greater than c_{rw} . To remedy this, we do the following: for all rooms r and w and all times k , if a people are "swapping" from r to w , make those people stay in their original room for that timestep and let the rest cross the hallway as normal. We see that this operation does not change the number of people in every room after the i^{th} timestep. It is easy to see that after every change has been made, we are left with another nonnegative integer-weighted flow, and the number of people in each room at each timestep remains the same. Clearly, then, the flow obtained after all of the swaps still has size P . We can then convert this into an evacuation plan by routing each person as a distinct unit of flow in the algorithm.

Thus, the algorithm will create a scheme to empty the building in k time steps if such a scheme exists. Ford-Fulkerson's runtime on a graph with m edges and maximum flow f is $O(mf)$. If we assume our original graph has m edges, our modified graph has km edges. Since the maximum flow of this graph is at most P , the runtime of Ford-Fulkerson on the modified graph is $O(kmP)$. In the optional cleanup stage, we look at each of the m pairs of connected edges at each of the k timesteps and check if there is a swap between them. Clearly, the entire phase takes $O(mk)$ time. Thus, this algorithm runs in polynomial time $O(kmP)$. Via our original observation, increasing k one by one in this algorithm will give us a polynomial time algorithm for the original problem, as desired.