Homework Assignment 3: Some partial solutions.

CSC 200Y: Social and Economic Networks

Due: February 10, 2016

Be sure to include your name, student number and tutorial room with your assignment. If your handwriting is *possibly* illegible, be sure to hand in your assignment in some typed form. Make sure that your assignement is *stapled* together.

1. Consider a stable marriage problem with four women Xina, Yael, Zoe and Winnie, and four men, Elvis, Farid, Gabe, and Hugh. Using initials instead of names, the women's preference for the men are:

   $X : F \succ H \succ E \succ G \quad Y : E \succ F \succ G \succ H \quad Z : E \succ F \succ G \succ H \quad W : G \succ E \succ F \succ H$

   And the men's preferences for the women are:

   $E : X \succ W \succ Y \succ Z \quad F : Y \succ Z \succ X \succ W \quad G : Y \succ Z \succ W \succ X \quad H : Z \succ W \succ Y \succ X$

   (a) Suppose we run the female-proposing deferred acceptance (FPDA) algorithm. Show how each iteration will proceed by: (i) clearly labelling each iteration; (ii) stating exactly which proposals will be made at that iteration; (iii) and stating exactly which engagements will be in place at the end of that iteration (once relevant proposals are accepted or rejected). Indicate clearly which iteration is the final one and what stable marriages result from FPDA.

   Solution:
   Round 1

Proposals	Engagements
Y to E	E:Y
Z to E	F:X
X to F	G:W
W to G	

   Round 2

Proposal	Engagements
	E:Y
Z to F	F:Z
	G:W

Round 3

Proposals	Engagements
	E:Y
	F:Z
X to H	G:W
	H:X

6b.

(b) We say a man *lies* in FPDA when he rejects a proposal from a woman even though he prefers the proposer to his current fiancee. Identify some man who can *lie* by falsely rejecting a proposal, thereby changing the outcome of FPDA so that he ends up married to a more preferred partner that he did in part (a). State when he should lie in FPDA and what stable marriages will result. (Assume all other men continue to accept and reject proposals truthfully.)

Man F should reject woman Zs proposal in Round 2. If so, he ends up married to Y, whom he prefers most. Once F rejects Z, Z must move down her list to G in round 3. From that point, the following will occur, leading to man F ending up with woman Y.
Round 3: Z proposes to G, G accepts and jilts W.
Round 4: W proposes to E, E accepts and jilts Y.
Round 5: Y proposes to F, F accepts and jilts X.
This is why F benefits from his lie!
Round 6: X proposes to H, H accepts. We are done.

(c) Now consider the *male-proposing* deferred acceptance (MPDA) algorithm on the same problem (i.e., same set of preferences). As in part (a), clearly indicate what happens at each iteration of MPDA and what stable marriages result when MPDA terminates.

2. Consider a matching market in which there are four condos, C1, C2, C3, and C4, for sale in a certain neighborhood. There are four buyers Dan, Etta, Fahiem, and Gabor wishing to purchase these condos. The buyers have the following valuations:

	C1	C2	C3	C4
Dan	2	7	6	4
Etta	7	9	5	4
Fahiem	3	2	5	4
Gabor	4	7	2	1

(a) Suppose we run the "auction" mechanism described in Ch.10 to determine market clearing prices, but *without using the price reduction step (i.e., we do not insist that the smallest price is* 0*).* Recall that at each round of the auction, if there is a constricted set of buyers in the preferred seller graph, then we trigger price increases for one or more condos. We want you to run this auction using the following rule: *if there is more than one possible constricted set of buyers at any round of the auction, choose a **minimal** set of constricted buyers to determine which condo prices increase.* For each round of the auction, show each of the following: i. prices for each condo; ii. the preferred seller graph; iii. if the prices are not market clearing at that round, the *minimal* constricted set of buyers you selected and the condos whose prices will increase; and iv. if the prices are market clearing at that round, a perfect matching that assigns a condo to each buyer.

Solution
******* Round 1
all prices 0; D,G constricted by links to C2; C2 price raised to 1
******* Round 2
E,G constricted by links to C2; C2 price raised to 2
******* Round 3
D,F constricted by links to C2,C3; C3 price raised to 1
******* Round 4
With prices (0,2,1,0) the PSG is market clearing

(b) Repeat part (a), but this time, if there is more than one possible constricted set of buyers at any round of the auction, choose a **maximal** set of constricted buyers to determine which condo prices increase. Be sure to show all details of every round as in part (a). How do the market clearing prices reached compare to those in part (a)?

(c) Starting from the market clearing prices computed in part (a), suppose the sellers each decided to increase their prices by the same amount i. What is the largest price increase i they could impose without changing the preferred seller graph?

Solution: As long as everyone maintains a non negative value for their most desired condos, we can increase the prices. Thus we can increase the prices by 4 but no further.

(d) Suppose another very attractive condo C5 comes onto the market. All four buyers value this condo more than any of the other four condos. The owner of this condo recognizing its attractiveness prices it higher than any of the buyer valuations. (i) Briefly explain why the addition of this condo (at this price) won't change the market clearing prices of the other condos. (ii)

In real-world real estate markets, the entrance of a new condo into the market at a high price will often cause prices for other condos in the neighborhood to increase in value. Using notions described in Chapter 9, explain what features of the market clearing model in Chapter 10 prevent this from happening, and what features of real-world real estate might cause this to occur.

(e) Suppose one of the original four sellers takes their condo off the market. Which seller, if she took her condo off the market before the auction in part (a) begins, would cause the greatest loss in social welfare among the buyers when the auction ends? Justify your response by explaining how the auction would proceed, and with what result it ends, when only the remaining three sellers participate.

3. Consider a collection of directed graphs $G_n = (V_n, E_n)$ where we think of these graphs being generated over time with V_n strictly increasing in size. For example, this could be the collection of web graphs that are generated say each month. Suppose we say that a set of nodes $S \subseteq V_n$ is a *giant* strongly connected component if $|S| \geq |V_n|/c$ for some fixed constant c. Suppose now that edges in G_n are generated randomly and independently with each directed edge having probability $p > 0$. (Note: This is not realistic for the web graph.) Given an informal argument to show that with high probability there cannot be two disjoint giant strongly connected components. I am not looking for an exact probability bound but rather a high level argument that would be the basis for establishing such a probabilistic bound.

   Solution: Suppose there were two giant strongly connected components C_1 and C_2 and just to be specific lets say that each has size at least $|V_n|/100$. The claim is that it is highly unlikely that these two components would not become one component given the random edge generation process that has been stated. In order for these giant components to be merged into one component we component all we need is that

   - there is an edge from some node $v_1 \in C_1$ to some $v_2 \in C_2$ and
   - there is an edge from some node $u_2 \in C_2$ to some node $u_1 \in C_1$.

   But the probability of this not happening is at most $2[1 - p^{(|V_1| \cdot |V_2|)}]$

4. Consider a toy small example of a directed web graph $G = (V, E)$ (this is Fig. 13.8 in the text):

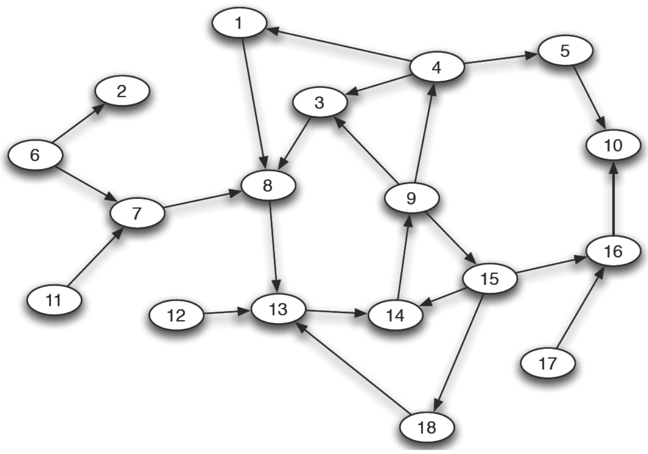

(a) Suppose we assign all 18 nodes equal initial page rank values of $\frac{1}{18}$. Now suppose we apply the *unscaled* Page Rank algorithm to this graph. After one iteration (or round) of the algorithm, will any nodes have a page rank value of zero? If so, which ones and why? If not, briefly explain why not.

Solution:

Nodes 6, 11, 12, and 17 will give up their initial page rank after the first iteration (since they all have outgoing edges). Since their in-degree is 0, they will receive no page rank from any other node; hence they will have page rank zero.

(b) Answer the same question as in part (a) for both two and three iterations of unscaled page rank.

Solution:
After two iterations, the nodes above will still have page rank zero. In addition, node 7 have page rank zero, since all of its incoming edges are from nodes that lost their page rank after the first iteration. No other nodes have page rank zero after two iterations (all other nodes have an in-path of length 2). No further nodes will lose their page rank after the third iteration.

(c) When the algorithm converges (i.e., reaches equilibrium), which nodes in the graph will have non-zero page rank values? What will their equilibrium page rank values be? Briefly justify your answer.

Solution:
A node with out-degree 0 never loses any page rank, hence nodes 2 and 10 will have non-zero page rank. But every other node has a directed path leading to either node 2 or 10, so all page rank gradually leaks away into these nodes, so they are the only nodes with positive page rank. (Formally, these are the only equlibrium values since this is the only solution to the equilibrium

equations.) Node 2 obtains and keeps half of the initial page rank node 6, plus it maintains its original page rank; so at equilibrium, its page rank is $\frac{3}{36}$. All other nodes will gradually lose their page rank to node 10 which is a sink; thus it obtains the remaining page rank of $1 - \frac{3}{36} = \frac{33}{36}$.

(d) Suppose we repeat part (a) but consider a variation of the graph in which a single node can delete a single one of their outgoing edges. Which nodes can increase their (equilibrium) page rank values by deleting a single outgoing edge? Briefly justify your answer. *(Do not consider simultaneous deletion of multiple edges, just one single edge.)*

(e) What is the minimum number of edges that would need to be added to the graph so that every node has non-zero page rank? Briefly explain what edges would be needed and why.

(f) Suppose we use *scaled* Page Rank on the original graph in part (a) with scaling factor $s = 0.9$. Which nodes will have non-zero (equilibrium) page rank values? Qualitatively, briefly describe which node will now have the highest page rank value.

5. The following is a toy web graph with hubs C, D, E and F, and authorities A, B, and G. (This graph and question is a slight modification of Question 3 in Sec. 14.7 of the text.)

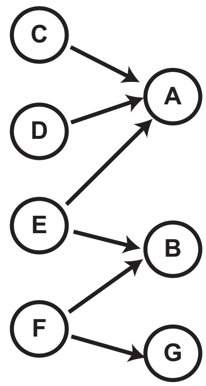

   (a) Show the (hubs and authorities) values obtained by running two rounds of the hubs and authorities algorithm on this network. Show the values both before and after the final normalization step, in which we divide each authority score by the sum of all authority scores, and divide each hub score by the sum of all hub scores. You should express the normalized scores as fractions rather than decimals.

   (b) Suppose you wish to add a new authority web site X to the network and want it to have as high an *authority* score as possible. You also have the ability to create another new hub web site Y which we can use as a hub to elevate the score of X. Consider the following three options:

      - You add a link from Y to X (and add no other new links).
      - You add a link from Y to A and X (and no other links).
      - You add a link from Y to A, B, G and X (and no other new links).

      For each option, show the normalized authority values that each of A, B, G and X obtain when you run 2-steps of hub-authority computation on the resulting network (as you did in part (a)). Which option gives X the highest (normalized) score?

6. Suppose a search engine has three slots $\{a, b, c\}$ to sell and there are three advertisers $\{x, y, z\}$. We will assume the basic advertising valuation model in which: each slot i has clickthough rate of r_i; each advertiser j has a value v_j per click; and hence the value of advertiser j for slot i is $v_{ij} = r_i \cdot v_j$. Suppose the clickthrough rates and advertiser per-click values are as follows: $r_a = 11, r_b = 5, r_c = 4$ and $v_x = 9, v_y = 7, v_z = 4$. We will also assume that there is a *reserve price* of 1 per click for any slot.

   (a) How will the slots be allocated and what prices will be charged to each advertiser if the search engine uses the *VCG mechanism*? Explain your answer.

   (b) Now let's consider the *generalized second-price (GSP)* mechanism. Assume that all advertisers bid truthfully (i.e., state their true per-click values). How will the slots be allocated and what prices will be charged to each advertiser if the search engine uses the GSP mechanism? Explain your answer.

   (c) As in part (b), assume slots are allocated using the GSP mechanism. For *each* of the three advertisers, explain what their *best response* would be if the other two advertisers bid truthfully. In other words, describe a per-click bid (it may not be unique) that would give them the highest payoff if the others bid truthfully. Which advertisers (if any) benefit by not bidding truthfully? Please justify your answer.

   (d) Unlike the VCG mechanism, GSP may not induce truthful bidding. However, consider the following claim: no advertiser can improve his utility by bidding *greater* his true per-click value (and may possibly decrease his utility). Give an informal argument that justifies this claim.