Networks: Fall 2012	Homework
Jon Kleinberg and Eva Tardos	Due 11:15am Friday, October 26, 2012

As noted on the course home page, homework solutions must be submitted by upload to the CMS site, at `https://cms.csuglab.cornell.edu/`. The file you upload must be in PDF format. (It is fine to write the homework in another format such as Word; from Word, you can save the file out as PDF for uploading.)

The CMS site will stop accepting homework uploads after the posted due date. We cannot accept late homework except for University-approved excuses (which include illness, a family emergency, or travel as part of a University sports team or other University activity).

Reading: The questions below are primarily based on the material in Chapters 13, 14 and 15 of the book. Note that we will start covering Chapter 15 on Monday, October 15th.

---

(1) Consider a trading network with intermediaries on Figure 1 in which there are three sellers S_1, S_2 and S_3 (on the left), two buyers B_1 and B_2 (on the right) and two traders (intermediaries) T_1 and T_2. Each seller can trade with only one trader; trader T_1 for seller S_1 and trader T_2 for sellers S_2 and S_3. Buyer B_1 can only trade with trader T_1. Buyer B_2 can only trade with trader T_2. The value for the good of the sellers is 0, buyer B_1 values the good at 1, buyer B_2 has value 2 for the good, as indicated on the figure.

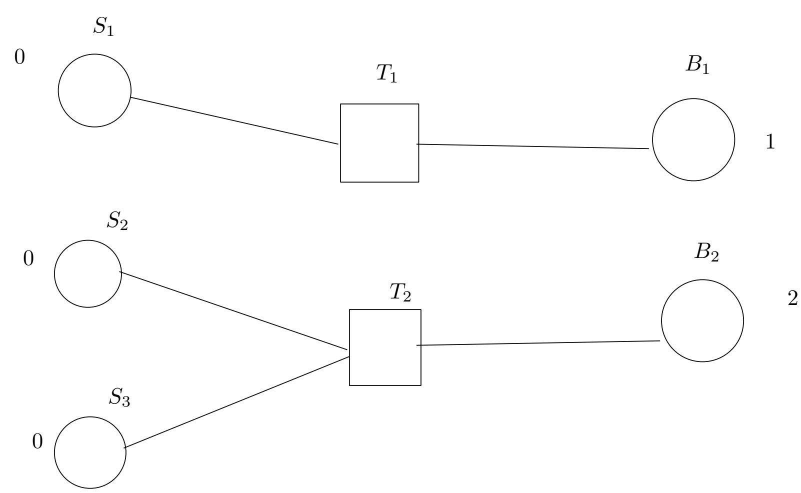

Figure 1: The network buyers, sellers and traders for Question (1).

(a) Describe what all possible Nash equilibria are, including both prices and the flow of goods. Give a brief explanation for your answer. How much profit can each trader make?

(b) Now assume we add an extra connection making it possible for buyer B_1 to trade with either trader. We want to explore how this change effects the possible Nash equilibria. Describe what all possible Nash equilibria are now, including both prices and the flow of goods. How much profit can each trader make? is it the same as in part (a)? Give a brief explanation for your answer to the profits.

(c) Suppose instead of the extra connection in part (b), we add the extra connection making it possible for seller S_2 to trade with either trader. We want to explore how this change effects the possible Nash equilibria. Describe what the possible Nash equilibria are, including both prices and the flow of goods. How much profit can each trader make? is it the same as in part (a)? Give a brief explanation for your answer to the profits.

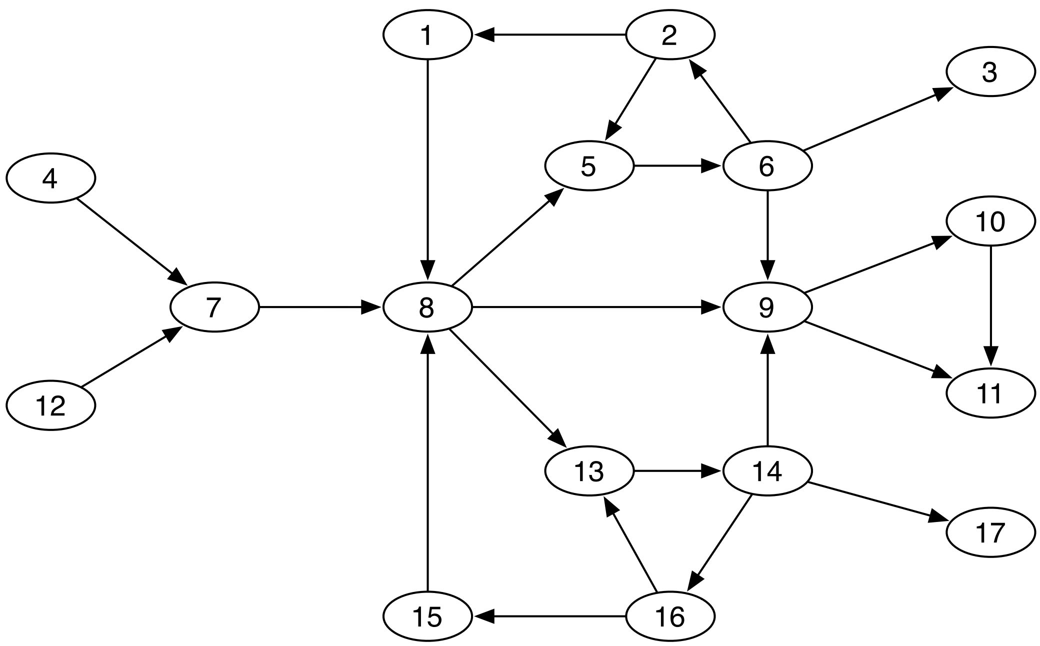

Figure 2: The network of Web pages for Question (2).

(2) Consider the directed graph shown in Figure 2, with nodes representing Web pages and each directed edge representing a link from one Web page to another.

(a) List the nodes in the largest strongly connected component of this graph.

(b) As new links are created or old ones are removed among an existing set of Web pages, such as the one in Figure 2, the set of nodes in the largest strongly connected component can change. Here's an example of how such a change can occur, through the addition of an edge.

Suppose you are allowed to add one link to the graph in Figure 2, going from one node in the figure to another; which link would you add if you wanted to increase the size of the largest strongly connected component by as much as possible? Give an explanation for your answer.

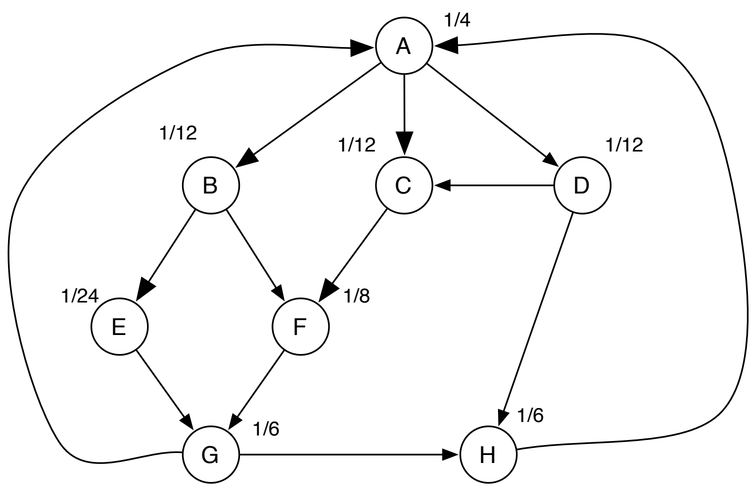

Figure 3: The network of Web pages for Question (3).

(3) Let's consider the limiting PageRank values that result from the Basic PageRank Update Rule (i.e. the version where we don't introduce a scaling factor s). In Chapter 14, these limiting values are described as "exhibiting the following kind of *equilibrium*: if we take the limiting PageRank values and apply one step of the Basic PageRank Update Rule, then the values at every node remain the same. In other words, the limiting PageRank values regenerate themselves exactly when they are updated."

This description gives a way to check whether an assignment of numbers to a set of Web pages forms an equilibrium set of PageRank values: the numbers should add up to 1, and they should remain unchanged when we apply the Basic PageRank Update Rule. (See for example Figure 14.7 in Chapter 14 of the book.)

Try this on the network of Web pages shown in Figure 3. In particular, say whether the indicated set of numbers forms an equilibrium set of PageRank value under the Basic PageRank Update Rule. Also, provide an explanation for your answer: specify either why they form an equilibrium, or how they fail to form an equillibrium.

---

(4) For the network of Web pages shown in Figure 4, determine the equilibrium PageRank values under the Basic PageRank Update Rule. Give an explanation for your answer. (Hint: You can use the approach described in class; let a denote the (unknown) PageRank value of node A, then work out the other PageRank values in terms of a, and then determine a value for a.)

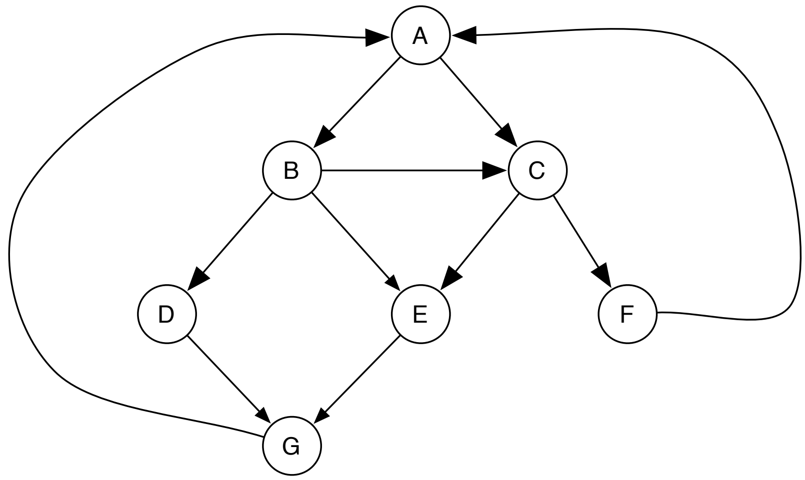

Figure 4: The network of Web pages for Question (4).

(5) In Chapter 14, we discussed the fact that designers of Web content often reason explicitly about how to create pages that will score highly on search engine rankings. In a scaled-down setting, this question explores some reasoning in that style.

(a) Show the values that you get if you run two rounds of computing hub and authority values on the network of Web pages in Figure 5. (That is, the values computed by the k-step hub-authority computation when we choose the number of steps k to be 2.)

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. (We will call the scores obtained after this dividing-down step the *normalized scores*. It's fine to write the normalized scores as fractions rather than decimals.)

(b) Now we come to the issue of creating pages so as to achieve large authority scores, given an existing hyperlink structure.

In particular, suppose you wanted to create a new Web page X, and add it to the network in Figure 5, so that it could achieve a (normalized) authority score that is as large as possible. One thing you might try is to create a second page Y as well, so that Y links to X and thus confers authority on it. In doing this, it's natural to wonder whether it helps or hurts X's authority to have Y link to other nodes as well.

Specifically, suppose you add X and Y to the network in Figure 5. In order to add X and Y to this network, one needs to specify what links they will have. Here are two options; in the first option, Y links only to X, while in the second option, Y links to other strong authorities in addition to X.

- *Option 1:* Add new nodes X and Y to Figure 5; create a single link from Y to X; create no links out of X.

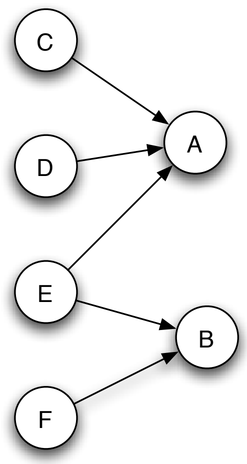

Figure 5: The network of Web pages for Question (5).

- *Option 2:* Add new nodes X and Y to Figure 5; create links from Y to each of A, B, and X; create no links out of X.

For each of these two options, we'd like to know how X fares in terms of its authority score. So, for each option, show the normalized authority values that each of A, B, and X get when you run the 2-step hub-authority computation on the resulting network (as in part (a)). (That is, you should perform the normalization step where you divide each authority value down by the total.)

For which of Options 1 or 2 does page X get a higher authority score (taking normalization into account)? Give a brief explanation in which you provide some intuition for why this option gives X a higher score.

(c) Suppose instead of creating two pages, you create three pages X, Y, and Z, and again try to strategically create links out of them so that X gets ranked as well as possible.

Describe a strategy for adding three nodes X, Y, and Z to the network in Figure 5, with choices of links out of each, so that when you run the 2-step hub-authority computation (as in parts (a) and (b)), and then rank all pages by their authority score, node X shows up in second place.

(Note that there's no way to do this so that X shows up in first place, so second place is the best one can hope for using only three nodes X, Y, and Z.)

---

(6) Suppose a search engine has three ad slots (a, b and c) that it can sell. Slot a has a clickthrough rate of 4, slot b has a clickthrough rate of 3 and slot c has a clickthrough rate of 1. There are three advertisers (x, y and z) who are interested in these slots. Advertiser x values clicks at 7 per click, advertiser y values clicks at 4 per click, and advertiser z values clicks at 2 per click.

(a) Suppose that the search engine which owns the slots does not know the advertiser's values per click. The search engine runs the VCG Procedure to allocate slots. What assignment of slots will occur and what prices will the advertisers pay in total? What are the per-click prices paid by the advertisers for each slot?

(b) Set up this problem of allocating slots to advertisers as a matching market in which the values per click, as well as the clickthrough rates, are known. Run the procedure that we used in Chapter 10 to construct market clearing prices. What are the market clearing prices and what perfect matching between advertisers and slots occurs when you use these market clearing prices? [You do not need to show the steps of the procedure. It's enough to display the final prices and the perfect matching.] What are the per-click prices paid by the advertisers for each slot in this solution?

(c) Consider the following bids submitted by the advertisers for their per-click value in a GSP auction: advertiser x bids 6, advertiser y bids 3, and advertiser z bids 1 (each submitting a bid one smaller than their value). Do these bids form a Nash equilibrium in the GSP auction? Briefly explain your answer.

(d) Can you give a set of Nash equilibrium bids in GSP that will result in the same assignment and same payments as your solution in parts (a) or (b)?