CPSC 455/555 // ECON 425/563, Fall 2011
Solution Set for Exam 2

Question 1:

(a) Let $G = (V, E, t)$ be a trust graph, as the term is defined in Section 27.5, and let v_0 be a distinguished "start node."

See "Definition" 27.2 of the term *sybil strategy*. It order to complete it, we must define what "collapse" means in this context. To "collapse U' into a single node labeled v," do the following for each $x \in V' - U'$: Collapse all edges $i_{xu} = x \to u$ in E' such that $u \in U'$ into a single edge $x \to v$ with trust value $t(x \to v) = \sum_{i_{xu}} t'(i_{xu})$, and collapse all edges $o_{ux} = u \to x$ in E' such that $u \in U'$ into a single edge $v \to x$ with trust value $t(v \to x) = \sum_{o_{ux}} t'(o_{ux})$.

Pathrank$_v(G)$ is the (weighted) distance from v_0 to v. Here, the weight (or length) of a directed edge e is the inverse of the trust value $t(e)$. Note that, with respect to the Pathrank reputation function (unlike the max-flow and PageRank functions), smaller reputational values are more desirable than larger ones.

max-flow$_v(G)$ is the maximum flow from v_0 to v. Here, the capacity of e is simply $t(e)$.

(b) If G is the trust graph in Figure 1, then max-flow$_v(G) = 2$, and max-flow$_w(G) = 2.25$. So w outranks v in G. In one very straightforward sybil strategy (G', U'), the sybil set $U' = \{v, u\}$. The edge set E' of G' is identical to the edge set E of G except that $v \to w$ is not present in E', and the edge $u \to w$ is present. The trust function t' of G' is identical to t in G except that $t(u \to w) = 1$, and of course t' is undefined on $v \to w$, because the latter is not in E'. The trust graph G' is depicted in Figure 2. Now max-flow$_v(G') = 2$, and max-flow$_w(G') = 1.25$. The reputational values of all nodes in V except w are unchanged by this sybil strategy; because v outranks w in G', it has improved its rank.

There is no sybil strategy that v can use to improve its rank with respect to the Pathrank reputation function, because Pathrank is rank-sybilproof. See Theorem 27.9 and HW5.

Question 2:

(a) See Definition 27.1.

(b) The *logarithmic scoring rule* of Sec. 27.4.2 is one example. Full credit will be given for any correct answer.

(c) If one signal is highly likely, *e.g.*, a good meal at a very highly rated restaurant, then a rater who receives a different signal will achieve a higher expected payoff by lying than by reporting truthfully.

(d) This SPNE is very "fragile." If just one player deviates, even by mistake, then everyone will play D forever, and the social efficiency will be lost.

Question 3:

(a) f is computable in this model. If $n < k$, then all agents know that $f(x_1, \ldots, x_n) = 1$, and no bidding or price announcements are needed. If $n \geq k$, then the computation will converge to the correct value of f, because $f(x_1, \ldots, x_n) = 1$ if and only if $\sum_{i=1}^{n} \omega_i x_i \geq 1$, where $\omega_i = \frac{1}{k}$, for $1 \leq i \leq n$.

g is not computable in this model. Suppose that $n = 3$ and that the common prior distribution is uniform. There are just two input vectors on which g is 1, namely $(0, 0, 0)$ and $(1, 1, 1)$. For each i, regardless of his input bit x_i, agent i will bid $\frac{1}{4}$, because only 1 of the 4 equally likely pairs of

inputs of the other agents will cause g to be 1. The price announced after this one round of bidding will be $\frac{1}{4}$, which conveys no information to any agent. Bidding will stop at this point, because the process has converged on a price p^{∞}, but this price is not the value of g on any input vector.
(b) This model does not capture strategic behavior by bidders, *i.e.*, it assumes that all agents bid truthfully in all rounds. It also assumes common knowledge of a prior distribution on the input vector; in practice, it is not clear why this common knowledge would be available to a large number of agents who may be geographically distributed and may not know each other. See the paper "Computation in a Distributed Information Market" for a more in-depth discussion of the strengths and weaknesses of the model.

Question 4:
(a) Cooperate in stage 0. In stage $k+1$, cooperate if the other player cooperated in stage k, and defect if the other player defected in stage k.
(b) See Theorem 2.6 and Example 2.3 in
`www.seas.harvard.edu/courses/cs186/doc/2-game-theory.pdf`, which were covered in Lecture 15, for the proof that this strategy profile is a NE for $\delta = .9$. To see that it is not a SPNE, consider the subgame that starts immediately after a stage in which the players play (D, D). If both players stick with tit-for-tat in that subgame, then they will both play D forever, and the discounted average payoff to player 1 in the subgame will be $\overline{\pi_1} = (1 - .9) \sum_{t=0}^{\infty} (.9)^t \cdot 1 = (.1)(\frac{1}{1-.9}) = 1$. If player 1 deviates to "play C for ever," and player 2 sticks with tit-for-tat, then the players will play (C, D) in the 0^{th} stage of the subgame, after which they will both play C forever. Player 1's discounted average payoff in the subgame would therefore be $\overline{\pi_1} = (1 - .9)(0 + \sum_{t=1}^{\infty} (.9)^t \cdot 3 = (.1)(3)(.9) \sum_{t=0}^{\infty} (.9)^t = (.1)(3)(.9)(\frac{1}{1-.9}) = 2.7$.
(c) Sec. 3.3 of `www.seas.harvard.edu/courses/cs186/doc/3-P2P-file-sharing.pdf`, which was covered in Lecture 16, gives examples of how a peer can benefit by deviating from the reference-client strategy.

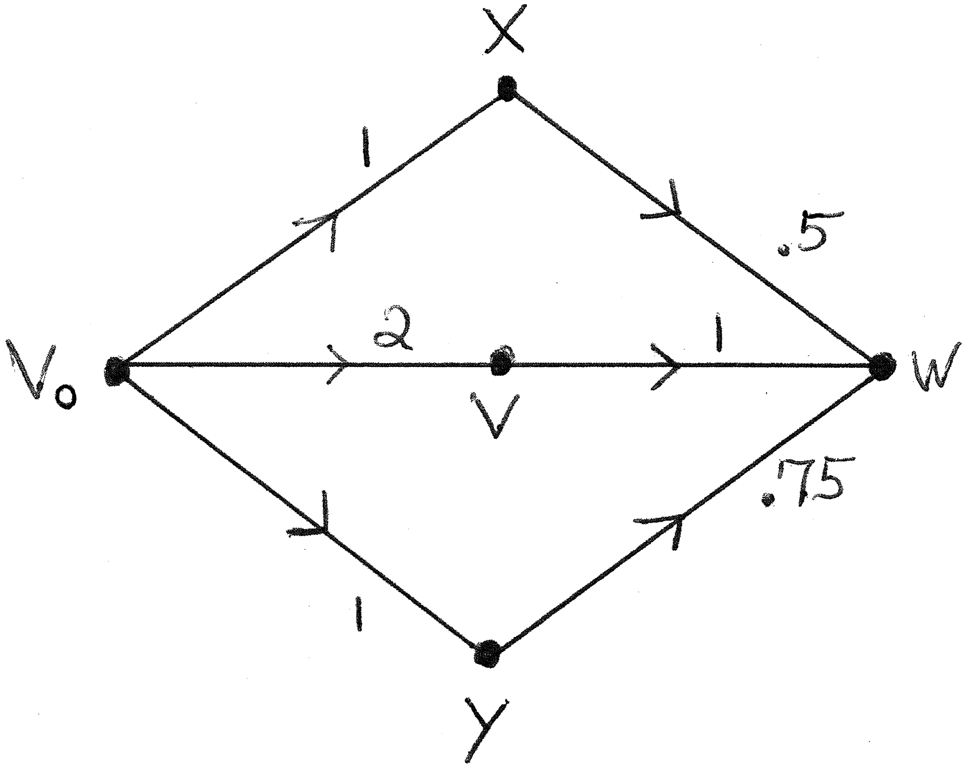

Figure 1

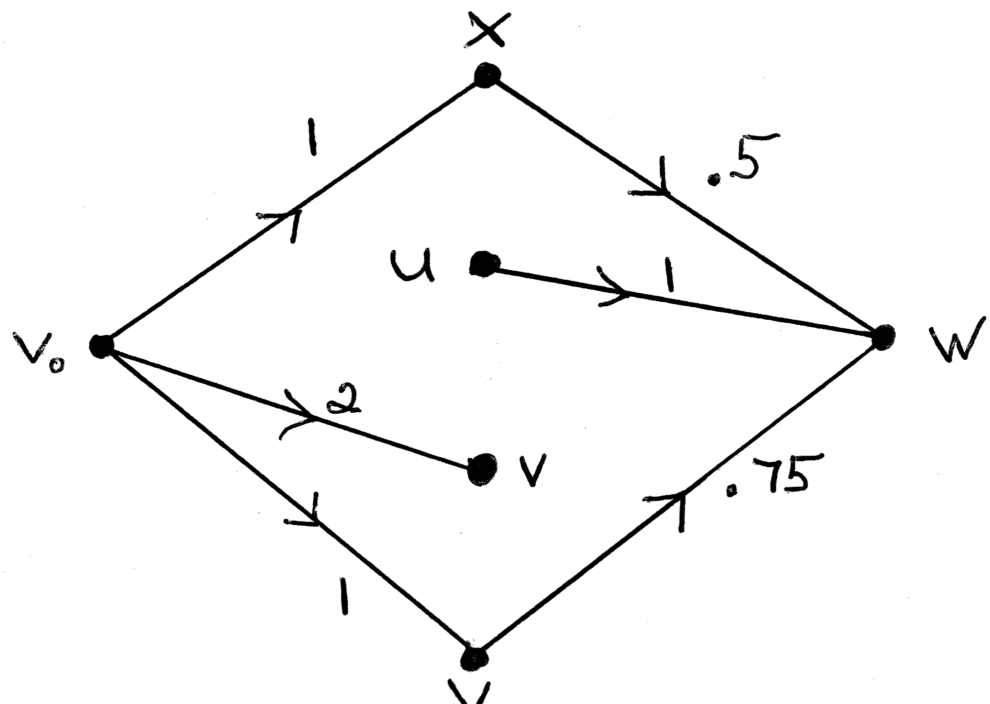

Figure 2