Algorithmic Game Theory

Assignment 4: due April 16, 2025

Assignment policies:

1. While you may dicuss the problems with others, you must write up the solution by yourself, in your own words.
2. Please write in your submission the people with whom you discussed the problems, as well as any references you used.
3. Late submissions will not be accepted, unless agreed to by me **prior** to the last date for submission.
4. Please write clearly and legibly, and include how you arrived at the solution!

Question 1 [10]: Consider the following single-parameter mechanism design problem: given an undirected graph $G = (V, E)$, bidders correspond to vertices (there is bijection between bidders and vertices). The feasible allocations correspond to matchings in the graph, and bidder i gets (private) value v_i if its corresponding vertex is matched in the allocation chosen. Design a DSIC mechanism that maximizes social welfare and runs in polynomial time.

Question 2 [10]: Consider a single-parameter mechanism design problem with n bidders where any feasible allocation $x \in X$ is the characteristic vector χ_S of a subset $S \subseteq [n]$. That is, each $X \subseteq \{0,1\}^n$. Further, X is *downward-closed*: if $S' \subseteq S$ and $\chi_S \in X$, then $\chi_{S'} \in X$.

For a DSIC mechanism (x, p), let b be the set of bids, and S be the set of winning bidders. Show that for any winning bidder i, i's payment is exactly the difference between (i) the maximum social welfare (assuming truthful bids) that would have been obtained by the other bidders if bidder i were not present in the environment, and (ii) the social welfare obtained by the other bidders in the current allocation.

This gives another way of looking at payments by a bidder: in any allocation, a bidder pays exactly its *externality*, or the loss in social welfare caused by its presence in the environment.

Question 3 [10]: Complete the proof of the following claim from the proof of the Gibbard-Satterthwaite theorem we did in class:

Claim 1. *If a social choice function f is incentive compatible and onto A, and $|A| \geq 3$, then for any preferences $\pi = (\pi_1, \ldots, \pi_n)$ of the agents and $T \subseteq S \subseteq A$, if $f(\pi^S) \in T$ then $f(\pi^T) = f(\pi^S)$.*

Question 4 [10]: For a cost-minimization game with n players, let $s \in \mathcal{S}$ be a pure strategy profile. A subset of players A has a *beneficial deviation* if there exists s_i' for each $i \in A$ so that

$$c_i(s_{-A}, s_A') \leq c_i(s)$$

for each player $i \in A$, with the inequality strict for at least a single player in A. Here $s_{-A} := (s_i)_{i \notin A}$, and $s'_A := (s'_i)_{i \in A}$. A pure strategy profile s is a *strong Nash equilibrium* if no subset has a beneficial deviation.

Prove that in the global connection games we studied earlier, if all players have the same source vertex s and the same destination vertex t, the price of anarchy for strong equilibria is 1.