COL866: Special Topics in Algorithms	**Fall 2023**
Assignment 3	
Total points: 100	Deadline: Oct 26 (Thursday)

1. [**20 points**] An instance with indivisible chores is said to have *binary* valuations if for every agent i and for every chore j, it holds that $v_{i,j} \in \{0, -1\}$. Prove that determining whether a given chores instance with binary and additive valuations admits an envy-free allocation is NP-complete.

   Hint: You may reduce from the NP-complete problem `SET SPLITTING`, which is defined as follows: Given a finite set S and a family $\mathcal{F}$ of subsets of S, does there exist a partition of the set S into S_1 and S_2 such that every element of $\mathcal{F}$ is split by this partition (i.e., no element of $\mathcal{F}$ is completely contained in S_1 or in S_2)?

2. [**15 points**] Consider a weighted generalization of the fair division problem with indivisible goods, where each agent i is associated with a weight $w_i > 0$ such that $\sum_i w_i = 1$.

   Let us define weighted generalizations of the fairness notions seen in the lectures as follows: A randomized allocation X is said to be *weighted envy-free* (wEF) if, for any pair of agents i and j, we have $\frac{\mathbb{E}[v_i(X_i)]}{w_i} \geqslant \frac{\mathbb{E}[v_i(X_j)]}{w_j}$. Similarly, a deterministic allocation A is said to satisfy *weighted envy-freeness up to one good* (wEF1) if, for any pair of agents i and j with $A_j \neq \emptyset$, there exists a good $g \in A_j$ such that $\frac{v_i(A_i)}{w_i} \geqslant \frac{v_i(A_j \setminus \{g\})}{w_j}$.

   Prove or disprove: For two agents with additive valuations over indivisible goods, there always exists a randomized allocation that is ex-ante wEF and ex-post wEF1.

3. [**10 points**] For the voting instance given below, compute the election winners under Plurality, Borda, Plurality-with-runoff, STV, Copeland, and Schulze voting rules. Ties are broken according to the lexicographic ordering $a \succ b \succ c \succ d \succ e$.

3 voters:	$c \succ d \succ b \succ a \succ e$
8 voters:	$c \succ e \succ b \succ d \succ a$
18 voters:	$d \succ e \succ c \succ b \succ a$
22 voters:	$e \succ c \succ b \succ d \succ a$
16 voters:	$b \succ d \succ c \succ e \succ a$
33 voters:	$a \succ b \succ c \succ d \succ e$

4. [**15 points**] On the YouTube channel of the online COMSOC video seminar:

   `https://www.youtube.com/channel/UCa_l2EzXiJxzfZKtu2mTkdA/videos`

   pick any one talk of your choice and summarize it in no more than 500 words. (Each video on the channel contains two or more talks; you only need to pick one talk that is sufficiently different from your project topic.) Feel free to use mathematical notation, examples or pictures in your summary if needed. Also provide the bibliographical information of the article(s) that the talk is based on.

5. [**30 points**] The Gibbard-Satterthwaite theorem establishes that no "reasonable" voting rule is strategyproof, meaning there is always a worst-case profile where some voter could misreport and improve. However, this result does not inform us about *how often* such bad profiles arise. In this exercise, we will experimentally analyze the frequency of manipulability of certain voting rules.

   Say there are n voters and m candidates. Design an experiment where the preference of each voter is generated independently and uniformly at random from among the $m!$ possible rankings over the candidates. Consider the following three voting rules: Plurality, Borda, and Copeland. Compare these rules in terms of the *fraction* of manipulable preference profiles (i.e., preference profiles where at least one voter can improve by misreporting). Assume throughout that ties are broken lexicographically.

   Clearly explain the (a) *hypothesis* (e.g., which voting rule did you expect to be the "most frequently manipulable" before you started the experiments and why), (b) *experimental setup* (e.g., what values of n and m did you consider, how many preference profiles did you sample for each setting of n and m and what were the reasons for these choices, which algorithm did you use to determine the manipulability of a profile), (c) *experimental observations*, and (d) *inference*.

6. [**10 points**] A *tournament* graph $G = (V, E)$ is a complete directed graph where, for any pair of vertices $i, j \in V$, either there is a directed edge from i to j or there is one from j to i.

   Given an *arbitrary* tournament graph $G = (V, E)$ on $|V| = m$ vertices, construct an election with m candidates and at most m^2 voters such that, for any pair of candidates i and j, a majority of voters prefer i over j if and only if there is a directed edge from i to j in G.

   In other words, show that *any* preference pattern can be realized as the majority vote of some group of voters.