COS 445 - PSet 1

Due online Monday, February 10th at 11:59 pm

Instructions:

- Some problems will be marked as *no collaboration* problems. This is to make sure you have experience solving a problem start-to-finish by yourself in preparation for the midterms/final. You cannot collaborate with other students or the Internet for these problems (you may still use the referenced sources and lecture notes). You may ask the course staff clarifying questions, but we will generally not give hints.
- Submit your solution to each problem as a **separate PDF** to codePost. Please make sure you're uploading the correct PDFs to the correct locations![1] If you collaborated with other students, or consulted an outside resource, submit a (very brief) collaboration statement as well. Please anonymize your submission, although there are no repercussions if you forget.
- The cheatsheet gives problem solving tips, and tips for a "good proof" or "partial progress."
- Please reference the course collaboration policy here.

For convenience, we restate some definitions used in this problem set.

[1] We will assign a minor deduction if we need to maneuver around the wrong PDFs. Please also note that depending on if/how you use Overleaf, you may need to recompile your solutions in between downloads to get the right files.

Problem 1: Both Sides Propose (20 points, no collaboration)

Consider the following algorithm, "Both-Proposing Deferred Acceptance:"

- Maintain a tentative matching M, initially empty.
- While there exists an unmatched student:
  - Pick an arbitrary unmatched student, s. s proposes to her favorite university who hasn't yet rejected her. If u prefers s to $t = M(u)$, update $M(s) = u$, $M(u) = s$, and $M(t) = \perp$.
  - Pick an arbitrary unmatched university (if one still exists), u. u proposes to their favorite student who hasn't yet rejected them. If s prefers u to $v = M(s)$, update $M(u) = s$, $M(s) = u$, and $M(v) = \perp$.

Either prove that Both-Proposing Deferred Acceptance always terminates in a stable matching, or provide an example of preferences and order of proposals such that Both-Proposing outputs a matching that is unstable for those preferences.[2]

[2] To clarify: if you provide a proof that Both-Proposing always gives a stable matching, your proof must work no matter which (valid) student/university is selected to propose. If you provide a counterexample, your analysis may select whichever (valid) student/university you like.

Problem 2: Instances with many stable matchings (30 points)

For this problem, you will consider a slight generalization of stable matchings than what we saw in class. Below, the definitions used in class are repeated, with changes ~~crossed out~~ and additions **in bold**.

- Two sides: students (S) and universities (U).
- Every student has a complete preference ordering over the universities. That is, every student s has a list $u_1 \succ_s u_2 \succ_s \ldots \succ_s u_{|U|}$ over the universities they prefer. Also, every university u has a list $s_1 \succ_u s_2 \succ_u \ldots \succ_u s_{|S|}$ over the students they prefer.
- Each student wants to be matched to one university, and each university u has slots for c_u students. ~~For simplicity of notation in this lecture, we'll let $c_u = 1$ for all universities, and $|S| = |U| = n$ (i.e. one student per university).~~ **For this problem, each university has capacity $\mathbf{c_u = 2}$, and therefore $\mathbf{|S| = 2 \cdot |U|}$ (i.e. two students per university)**.

Definition 1 (Blocking Pair)**.** *A student and university form a* blocking pair *for a matching M if:*

- *s strictly prefers u to her match in M ($u \succ_s M(s)$).*
- *u strictly prefers s to* ***one of*** *their match**es** in M ($\exists \mathrm{s}' \in \mathbf{M(u)}$* ***such that*** *$s \succ_u \mathrm{s}'$).*

Prove that, for all $n \geq 1$, there exists a stable matching instance with $4n$ students and $2n$ universities with two slots each, such that there are at least 2^n distinct stable matchings.

That is, for all $n \geq 1$, provide preferences $\succ_{s_1}, \ldots, \succ_{s_{4n}}$ for each of the $4n$ students over the $2n$ universities, and also preferences $\succ_{u_1}, \ldots, \succ_{u_{2n}}$ for each of the $2n$ universities over the $4n$ students, such that there are at least 2^n distinct matchings that are stable for the preferences $\succ_{s_1}, \ldots, \succ_{s_{4n}}, \succ_{u_1}, \ldots, \succ_{u_{2n}}$.

Hint: First try to prove the claim for $n = 1$. See if you can use this solution as a gadget to prove the claim for $n = 2$ and notice a pattern.

Problem 3: Random Preferences (40 points)

Consider an instance with n students, n universities, where each university has capacity one. Each student's preferences are drawn independently and uniformly at random (that is, each student is equally likely to have each of the $n!$ possible orderings). You may not assume anything about the universities' preferences (that is, your proof of the below claim must work *no matter what preferences* the universities have).

Consider an execution of the student-proposing deferred acceptance algorithm, and let M denote the matching output. Let Y_s denote the rank of student s's match under M (that is, one plus the number of universities that student s prefers to their match). Prove that $\mathbb{E}[\sum_s Y_s] = O(n \log n)$.

Hint 1: Prove that deferred acceptance terminates once every university has received a proposal. This hint is useful for getting started.
Hint 2: You may want to make use of the Coupon Collector Problem from PS0! This hint will only be useful once you are deep into the problem, and you should not try to use it as guidance for getting started.
Hint 3: Once you feel like you have the right intuition, you may want to consult the cheatsheet sections on the "principle of deferred decisions" and "coupling arguments" to make your intuition into a formal proof. This hint will only be useful once you are deep into the problem, and you should not try to use it as guidance for getting started.

Extra Credit: Almost Unique Stable Matchings

Recall that extra credit is not directly added to your PSet scores, but will contribute to your participation. Some extra credits are **quite** challenging. We do not suggest attempting the extra credit problems for the sake of your grade, but only to engage deeper with the course material. If you are interested in pursuing an IW/thesis in CS theory, the extra credits will give you a taste of what that might be like.[3]

Consider an instance with n students and n universities where student preferences are uniformly random, and university preferences are arbitrary. However, instead of a full preference ordering over all n universities, each student *truncates* their preferences at the top $c = O(1)$ universities (that is, they prefer to be unmatched rather than partner with a school outside their top c).

Say that a university is *uniquely stable* for this instance if they have the same partner in *all* stable matchings (where "unmatched" counts as a partner). Prove that the expected number of uniquely stable universities is $n - o(n)$.[4]

This is a long problem, and the following hints break down the key steps. If you can clearly state and prove concrete steps (e.g. clearly state claims suggested by some of the hints, and prove them), you will get partial extra credit.

Hint 1: You may want to prove the following fact first. Let M, M' be any two stable matchings. Then every student who is matched in M is also matched in M' (and vice versa). Every university that is matched in M is also matched in M' (and vice versa).

Hint 2: You may also want to prove the following: Let M be output by student-proposing deferred acceptance (i.e. each student stops applying if they are rejected by all of their top c schools), and let $M(u) = s$. Now consider modifying u's preferences by "removing" s and all s' that u likes less than s. That is, u declares that they would rather be unmatched that matched to s or anyone below s. Then any matching M' where $M'(u) \neq M(u)$ is stable for the new preferences if and only if it is stable for the original preferences.

Hint 3: You may next want to prove the following fact using Hints 1 and 2. Let M be output by student-proposing deferred acceptance (where each student only proposes to a university to which they apply), and let $M(u) = s$. Now consider modifying u's preferences by "removing" s and all s' that u likes less than s. That is, u declares that they would rather be unmatched that matched to s or anyone below s. Let M' denote the matching output by student-proposing deferred acceptance with this modified preference (and all others the same). If u is unmatched in M', then u is matched to s in *every* stable matching.

Hint 4: To start wrapping up, you may want to use the fact that Student-Proposing Deferred Acceptance outputs the same matching, independent of the order in which students propose (this is a corollary of a theorem from lecture). In particular, you may want to choose the order in which students propose to make use of the earlier hints.

Hint 5: Finally, you may use the following fact without proof:[5] imagine throwing k balls into n bins uniformly at random without replacement, and then repeating this procedure n times independently (so we pick n uniformly random lists of k distinct bins). Then with probability $1 - e^{-\Omega(n)}$, at least $n \cdot e^{-k}/4$ bins are empty.

[3] Keep in mind, of course, that you will do an IW/thesis across an entire semester/year, and you are doing the extra credit in a week. Whether or not you make progress on the extra credit in a week is not the important part — it's whether or not you enjoy the process of tackling an extremely open-ended problem with little idea of where to get started.

[4] Observe that this means it barely matters which side proposes in this model because almost everyone has the same partner regardless.

[5] This fact is oddly stated, because it is tailored to this problem to remove the need for calculations.