COS 330: Great Ideas in Theoretical Computer Science

Fall 2025

Problem Set 7

Module: New Models of Computation

Below is a reminder of key aspects of the PSet:

- The only goal of this PSet is to help you develop your problem-solving skills in preparation for the exams. Your performance on this PSet will not directly contribute to your grade, but will indirectly improve your ability to do well on the exams.
- Because your performance does not directly impact your grade, you may use any resources you like (collaboration, AI, etc.) to help you complete the PSet.
- We strongly suggest taking a serious stab at the PSet alone, to help self-evaluate where you're at. But, we also suggest collaborating with friends, visiting office hours, asking on Ed, and/or using AI tools to help when stuck. Even when able to complete the entire PSet on your own, you may still find any of these methods useful to discuss the PSet afterwards.
- Throughout the PSet, we've included some general tips to help put these into broader context. Exams will not have these, and future PSets may have fewer.
- We strongly suggest treating this like any other PSet, and writing up your solutions as if you were handing them in for a grade. At minimum, we very strongly suggest writing up sufficiently many solutions to discuss with your Coach.

Aligning Expectations

Recall that the 🗎 symbol implies that the following problem is an "exam-style" problem
We highly recommend that you write-up a full solution to this problem.

Problem 1 : Randomized 3-Coloring a Path 🗎

Consider a path graph $P_n = (V, E)$ with n nodes in the distributed model from class. Each node u maintains a state $s(u) \in \{0, 1\}$ (active if 0, stopped if 1) and a color $c(u) \in \{1, 2, 3\}$. Initially all nodes are active.

Time proceeds in rounds $t = 1, 2, \ldots$. In each round, every node u executes the following protocol:

- If $s(u) = 1$: node u keeps its color $c(u)$ and sends $c(u)$ to its neighbors.

- If $s(u) = 0$: node u picks $c(u) \in \{1, 2, 3\}$ uniformly at random (independently of all other random choices), sends $c(u)$ to its neighbors, receives their colors, and stops (sets $s(u) \leftarrow 1$) if and only if $c(u)$ differs from all neighbors' colors in this round.

Once a node stops, it never changes its color. The protocol terminates when all nodes have stopped.

(a) Show that when the protocol terminates, the final colors form a proper 3-coloring of the path.

(b) Show that in each round, every active node stops with probability at least $1/3$.

(c) Let T_u be the stopping time of node u (the first round when $s(u) = 1$). Prove that for every integer $t \geq 0$:

$$\Pr[T_u > t] \;\leq\; \left(\tfrac{2}{3}\right)^t.$$

(Hint: Use part (b) and induction on t.)

(d) Show that after $T := \lceil \frac{2 \ln n}{\ln(3/2)} \rceil = O(\log n)$ rounds, the probability that every node has stopped is at least $1 - 1/n$.

Problem 2 : Linear Search with Predictions

Setup. Consider the real line and suppose there is some hidden target at location $D \in \mathbb{R}$ with $|D| \geq 1$.

The game. Each round i, if you have not yet found the hidden target, you start at the origin 0. You then pick a destination $x_i \in \mathbb{R}$ to walk towards, from 0.

If, along the way, you pass the target at D (because D lies between 0 and your destination), you immediately stop and the game ends. Otherwise, you walk all the way to x_i and then return to 0.

Objective. Your goal is to end the game while walking the minimum total distance.

Strategy. A strategy for the game decides, for each round i, which x_i to choose. The Strategy knows what it chose in all previous rounds when making this decision. Moreover, because round i only takes place if you haven't yet found the target, every strategy is equivalent to picking an infinitely long list $x_1, x_2, \ldots, x_i, \ldots$ which means "if I make it to round i without yet finding the target, it means that $x_1, \ldots, x_{i-1}$ all failed and I would like to go to x_i next."

Offline Strategy. An Offline Strategy knows D when selecting $x_1, \ldots, x_i, \ldots$.

Online Strategy. An Online Strategy knows nothing about D, other than that $D \in \mathbb{R}$, when selecting $x_1, \ldots, x_i, \ldots$.

Competitive Ratio. We denote by $d(S; D)$ the total distance walked by Strategy S when the hidden target is at location D, and further denote $\mathrm{OPT}(D)$ to be the total distance walked by the optimal offline strategy. Finally, we define the competitive ratio of an online strategy S as:

$$\mathrm{CR}(S) \;=\; \max_{|D|\geq 1} \frac{d(S;D)}{\mathrm{OPT}(D)}.$$

(a) Prove that $\mathrm{OPT}(D) = |D|$.

(b) Let S be any strategy such that there exists a round i such that any of the following hold:

- There exists a $j > i$ such that $0 < x_j < x_i$.
- There exists a $j > i$ such that $x_i < x_j < 0$.
- Both x_i and x_{i+1} are positive.
- Both x_i and x_{i+1} are negative.

Prove that there exists another strategy S' such that both of the following hold:

- For all D, $d(S';D) \leq d(S;D)$. That is, S' always performs at least as well as S.
- There exists a D such that $d(S';D) < d(S;D)$.

(c) Consider the following search strategy G, which sets $x_i = (-2)^i$. Prove that

$$\mathrm{CR}(G) \leq 9.$$

Now consider a setting with a <u>prediction</u> that proposes a location $\widehat{D} \in \mathbb{R}$ for the target (which may be incorrect). Define the strategy $G'(\widehat{D})$, which takes a prediction $\widehat{D}$, as follows:

- Let $k := \lceil \log_2 |\widehat{D}| \rceil$ and set a scale parameter $\lambda := 2^k/|\widehat{D}| \geq 1$.
- Set $x'_i := (2^{i-1}/\lambda) \cdot (-1)^{i-1} \cdot (-1)^k \cdot \mathrm{sgn}(\widehat{D})$. That is, $|x'_i| = 2^{i-1}/\lambda$, the steps alternate directions, and $x'_{k+1} = \widehat{D}$.

(d) Show that the strategy $G'(\widehat{D})$ is 3-consistent. In other words, assume the prediction is correct, i.e., $\widehat{D} = D$. Show that for all $|D| \geq 1$:

$$\frac{d(G'(D);D)}{|D|} \;\leq\; 3.$$

(e) Recall that $G'(\widehat{D})$ uses steps with $|x'_i| = 2^{i-1}/\lambda$ where $\lambda = 2^k/|\widehat{D}|$. Prove that for any fixed starting direction and any target distance $|D| \geq 1$:

$$\frac{d(G'(\widehat{D}); |D|)}{|D|} \;=\; \frac{d(G;\, \lambda|D|)}{\lambda|D|}.$$

(Hint: How does scaling all step lengths by $1/\lambda$ affect the round on which the target is found?)

(f) Prove that $G'(\widehat{D})$ has robustness at most 9. In other words, show that:

$$\max_{D,\ \widehat{D}} \frac{d(G'(\widehat{D}); D)}{|D|} \ \leq\ 9.$$

(Hint: Use part (e). Also note that the starting direction choice does not affect the worst-case competitive ratio.)