

15-451 Algorithms, Spring 2017

Recitation #5 Worksheet

Zero-Sum Games

Solving for minimax optimality: Here is a version of the shooter-goalie game. Shooter is the row player, Goalie is the column player. The matrix entry is the payoff to the shooter (row player) if the shooter and goalie each choose the corresponding action (think of it as the chance the ball goes in). So, the row player wants to maximize the expected payoff and the column player wants to minimize it.

	<i>L</i>	<i>R</i>
<i>L</i>	1/2	1
<i>R</i>	1	1/4

What are the minimax optimal strategies for the two players, and what is the minimax value of the game (to the shooter)?

What about the following game:

	<i>L</i>	<i>R</i>
<i>L</i>	5/7	1
<i>R</i>	8/7	1/7

Lower bounds for randomized algorithms: Go through the analysis from the lecture notes for how to give a lower bound for *randomized* comparison-based sorting algorithms. This is in the lecture notes (but we went quickly over it in lecture saying that we would go through it in recitation). Here are the bullet points:

- First of all, in a zero-sum game, if there is a randomized strategy for the column player that guarantees the row-player pays at least V in expectation no matter what row it plays, this implies the row-player pays at least V in expectation even if it uses a randomized strategy.
- This implies that if there is a probability distribution over inputs such that any *deterministic* algorithm has expected cost at least V , then any *randomized* algorithm also has expected cost at least V . (This is called “Yao’s principle”).
- In the case of comparison-based sorting, we can choose the uniform distribution over all $n!$ permutations of $\{1, \dots, n\}$ as the randomized strategy for the adversary / column-player.
- Any deterministic sorting algorithm can be viewed as a binary tree, with each internal node having a comparison, and with $n!$ leaves, one for each ordering of the input elements (since we know the algorithm cannot do the same thing on two different input orderings if it wants to be correct, since the input elements are all distinct).
- The number of leaves at depth $\leq \lg(n!) - 10$ is at most $1 + 2 + 4 + 8 + \dots + n!/1024 < n!/512$. This means that over 99% of the leaves are at depth $\geq \lg(n!) - 10$, which implies the average depth is $\Omega(n \log n)$.
- Since the average depth is $\Omega(n \log n)$, this means the algorithm pays $\Omega(n \log n)$ in expectation to sort a random ordering of $\{1, \dots, n\}$.
- By Yao’s principle, this then implies that even a *randomized* comparison-based sorting algorithm cannot beat the $\Omega(n \log n)$ lower bound.