

CSCI5610: Final Exam


Problem 1 (10%). Problem 3 of Lecture 10 (removing Assumption 1 for union-find).

Problem 2 (20%). Problem 4 of Lecture 15 (dynamic range count).

Problem 3 (20%). Problem 2 of Lecture 17 (high success probability for LSH).

Problem 4 (20%). Let S be a set of n strings $\sigma_1, \sigma_2, \dots, \sigma_n$. Define $m = \sum_{i=1}^n |\sigma_i|$. Given a non-empty string q , a query reports the *largest* $i \in [1, n]$ such that σ_i contains at least one occurrence of q . Design a data structure of $O(m)$ space that can answer any query in $O(|q|)$ time.

Problem 5 (30%). Let P be a set of n points in \mathbb{R}^2 , each of which is associated with a real-valued weight. Given a *two-sided rectangle* $q = (-\infty, x] \times (-\infty, y]$, a query reports the maximum weight of the points in $q \cap P$. Design a data structure of $O(n)$ space that can answer any such query in $O(\log n)$ time.



In the above example, the weight of each point $p \in P$ is indicated next to p . The query with the search rectangle q shown should return 32.