

You are allowed to discuss with others but not allowed to use any references except the course notes and the books “Probability and Computing” and “Randomized Algorithms”. Please list your collaborators for each question. This will not affect your marks. In any case, you must write your own solutions.

There are totally 50 marks, and the full mark is 50. This homework is counted 8% of the course.

1. ***k*-wise independence**

(10 marks) Suppose that we are given m vectors $v_1, \dots, v_m \in \{0, 1\}^l$ such that any k of the vectors are linearly independent modulo 2. Let $v_i = (v_{i,1}, v_{i,2}, \dots, v_{i,l})$ for $1 \leq i \leq m$. Let u be chosen uniformly at random from $\{0, 1\}^l$. Let $X_i = (\sum_{j=1}^l v_{i,j}u_j) \bmod 2$. Show that the X_i are uniform, k -wise independent random bits.

2. **Maximum load**

(10 marks) We have shown that the maximum load when n items are hashed into n bins using a hash function chosen from a 2-universal family of hash functions is at most $\sqrt{2n}$ with probability at least $1/2$. Generalize this argument to k -universal hash functions. That is, find a value such that the probability that the maximum load is larger than that value is at most $1/2$. Then, find the smallest value of k such that the maximum load is at most $3 \ln n / \ln \ln n$ with probability at least $1/2$ when choosing a random hash function from a k -universal family.

3. **Finding approximate median in sublinear time**

(10 marks) We would like to find an approximate median of n distinct numbers in sublinear time. To do so, we sample $m \ll n$ numbers, find the median c of these m numbers, and report c as the approximate median of the n numbers. Let the sorted list of the n numbers be $\{x_1, x_2, \dots, x_n\}$ and so the true median is $x_{n/2}$. The approximate median is said to be a $\pm k$ -approximation if $c \in [x_{\frac{n}{2}-k}, x_{\frac{n}{2}+k}]$. Suppose we want the algorithm to succeed to find a $\pm k$ -approximation with probability at least 0.9999. What is the tradeoff between m and k ? How large should we set m if we want $k \leq \epsilon n$ for some small constant ϵ ? How large should we set m if we want $k \leq n^{1-\eta}$ for $\eta \leq 1/2$?

4. **Distinct elements**

(10 marks) In the distinct element algorithm in L05, we showed that with probability at least $2/3$, we have $(1 - \epsilon)D \leq Y \leq (1 + \epsilon)D$ where Y is the answer that our algorithm returns. Suppose we run k independent copies of the algorithm in parallel and obtain the estimates Y_1, \dots, Y_k and return the median Y' of Y_1, \dots, Y_k as our answer. Prove that $(1 - \epsilon)D \leq Y' \leq (1 + \epsilon)D$ with probability $1 - \delta$ when $k = O(\log \frac{1}{\delta})$.

5. Algebraic matching

(10 marks) In the following two sub-problems, you can assume that the determinant of a matrix where each entry is an element in $\mathbb{F}[x]$ of degree at most d can be computed in $O(dn^\omega)$ field operations, where $\mathbb{F}[x]$ is the set of single variate polynomials with coefficients in a finite field \mathbb{F} and $\omega \approx 2.37$ is the matrix multiplication exponent. Note that the determinant of such a matrix would be a single variate polynomial.

- (a) Given a bipartite graph where each edge is red or blue and a parameter k , determine if there is a perfect matching with exactly k red edges in $O(n^\omega)$ field operations with high probability.
- (b) Given a bipartite graph with a non-negative weight on each edge, determine the weight of a maximum weighted perfect matching in $O(Wn^\omega)$ field operations with high probability, where W is the maximum weight of an edge.