Homework 3

Due date: Tuesday, April 11, 2023 at 11:59PM (Gradescope)

Question 1 (40 points)

Let $N = 2^n$ and define the "*Powers of Two Graph*" (a folklore construction of a depth-robust graph) $G_n = (V, E)$ with nodes $V = [N]$ and the edge $E = \{(i-2^j, i) : i \le N \text{ and } i-2^j \ge 1\}$.

Part A. We say that a node $v \le N$ is α-forward good with respect to a set S of deleted nodes, if for all $r > 0$ the interval $[v, v+r-1]$ contains at most $\alpha \times r$ nodes in S. Suppose that v is α-forward good and let $U(j)$ denote the number of nodes in the interval $[v, v+2^j-1]$ that are not reachable from v in $G_n - S$. Similarly, let s_j denote the number of deleted nodes in the interval $[v, v+2^j-1]$. Use induction to prove that $U(j) \le \sum_{i=0}^{j} s_i 2^{j-i} \le j2^j\alpha$.

> *Answer:*
>
> ...

Part B. Suppose that $2^j < r < 2^{j+1}$. Show that at least $r - 2(j+1)r\alpha$ nodes in $[v, v+r-1]$ are reachable from v in $G_n - S$.

> *Answer:*
>
> ...

Part C. We say that a node $w \le N$ is α-backward good with respect to a set S of deleted nodes if for all $0 < r \le w$ the interval $[w-r+1, w]$ contains at most $\alpha \times r$ nodes in S. Show that if node w is α-backward good and node $v < w$ is α-forward good with respect to S with $\alpha = 0.01/n$ then there is a directed path connecting v to w in $G_n - S$.

> *Answer:*
>
> ...

Part D. Show that G_n is (e, d)-depth robust with $e = \Omega(N/n)$ and $d = \Omega(N)$ and lower bound the cumulative pebbling cost $\mathsf{CC}(G_n)$.

> *Answer:*
>
> ...

Part E. Assume that G is (e, d)-depth-robust with $e > d$. Suppose that we delete $|S| = e/2$ nodes from G. Show that the graph $G - S$ contains at least $e/(2d)$ node disjoint paths of length d. (**Hint:** To get started, let $S_0 = S$ and let $S_1 = S_0 \cup P$ where P is a directed path in $G - S_0$ containing exactly d nodes.)

Answer:
...

Part F. We say that a directed graph $G = (V = [N], E)$ is (e, d, f)-fractionally depth-robust if for any subset $|S| \le e$ of at most e nodes there is a subset $T \subseteq [N] \setminus S$ of $|T| \ge f$ nodes such that for every node $v \in T$ the graph $G - S$ contains a directed path of length d ending at node v. Supposing that $N = 2^n$ and G is $(\Omega(N), \Omega(N/n))$-depth robust show that G is (e, d, f)-fractionally depth-robust with $e = \Omega(N)$, $d = \Omega(N/n)$ and $f = \Omega(N)$. (**Hint:** You should used what you proved in part E to get started.)

Answer:
...

Part G. Suppose that G is (e, d, f)-fractionally depth-robust and consider the pebbling challenge game used in the analysis of Proofs of Space. In particular, suppose that Alice can place $e' < e$ pebbles on the graph G and then a challenger asks Alice to place pebbles on randomly selected nodes $v_1, \ldots, v_k$. Alice can place pebbles in parallel, but is not finished until she has placed pebbles on *all* of the challenge nodes $v_1, \ldots, v_k$. Upper bound the probability that Alice can complete the challenge within $d' < d$ steps.

Answer:
...

Resource and Collaborator Statement:
...

Question 2 (30 points)

Recall that a point function $f_{\alpha,\beta}(x)$ outputs β if $x = \alpha$ and $f_{\alpha,\beta}(x) = 0$ otherwise. Consider the following construction of a distributed point function. The setup algorithm picks a random Puncturable PRF key $K \in \{0,1\}^\lambda$ and sets $K_0 = \mathtt{iO}(1^\lambda, C_0)$ to Alice and $K_1 = \mathtt{iO}(1^\lambda, C_{1,\alpha,\beta})$ to Bob where functionality of the circuits C_0 and C_1 are described as follows $C_0(x) \doteq F_K(x)$ and $C_{1,\alpha,\beta}(x) = F_k(x)$ if $x \neq \alpha$; otherwise if $x = \alpha$ we have $C_{1,\alpha,\beta}(x) = F_K(x) \oplus \beta$. Consider the following security game: The attacker fixes (α_0, β_0), (α_1, β_1) and a role $i \in \{0,1\}$ (indicating whether the attacker plays the role of Alice/Bob) and sends these values to the challenger. The challenger picks a random coin b, sets $(\alpha, \beta) = (\alpha_b, \beta_b)$ and then generates $K_0 = \mathtt{iO}(C_0)$ and $K_1 = \mathtt{iO}(C_{1,\alpha,\beta})$ and sends K_i back the the attacker. Finally, the attacker outputs a guess b'. The attacker wins if $b' = b$ and we use $\mathtt{WIN}_{\mathcal{A}}(\lambda)$ to denote the event that the attacker $\mathcal{A}$ wins when using securing parameter λ. The advantage of an attacker $\mathcal{A}$ over random guessing is denoted $\mathtt{ADV}_{\mathcal{A}}(\lambda) = \Pr[\mathtt{WIN}_{\mathcal{A}}(\lambda)] - \frac{1}{2}$. We say that the DPF is secure if all PPT attackers $\mathcal{A}$ there exists a negligible function $\mu(\lambda)$ upper bounding $\mathtt{ADV}_{\mathcal{A}}(\lambda)$.

Part A. (5 points) Explain how Alice and Bob can locally generate their shares of $f_{\alpha,\beta}(x)$ given any input x.

Answer:

...

Part B. (25 points) Prove that DPF construction is secure according to the above distribution. You may assume that the PPRF and iO constructions are both secure.

Answer:

...

Resource and Collaborator Statement:

...

Question 3 (30 points)

Alice wants to design a delegated signature scheme. In particular, the delegated signature scheme should implement four PPT algorithms ($\mathsf{KeyGen}, \mathsf{DelegateKey}, \mathsf{Sign}, \mathsf{Verify}$). $\mathsf{KeyGen}(1^\lambda)$ takes as input a security parameter (λ) and outputs a secret-public key pair (sk, pk) and $\mathsf{DelegateKey}(sk, x)$ takes as input a prefix x and the secret key sk and outputs a key sk_x which can be used to sign any message of the form $m = x||y$. $\mathsf{Sign}(sk, m)$ outputs a signature σ such that $\mathsf{Verify}(pk, \sigma, m) = 1$. If $m = x||y$ then $\mathsf{Sign}(sk_x, m)$ outputs a signature σ such that $\mathsf{Verify}(pk, \sigma, m) = 1$. However, if x is not a prefix of m then $\mathsf{Sign}(sk_x, m) = \perp$.

Selective security game: In the selective security game, we fix a target message m^* and then the challenger $\mathcal{C}$ generates (sk, pk) and sends pk to the attacker $\mathcal{A}$. The attacker may make $q = \texttt{poly}(\lambda)$ queries to $\mathsf{DelegateKey}(sk, .)$ but may not submit a query x_i which is a prefix of m^*. The game ends when the attacker outputs an attempted forgery for m^*. The scheme is secure, if for all PPT attackers there is a negligible function upper bounding the probability that the attacker wins.

Part A. Use indistinguishability obfuscation to design a secure delegated signature scheme according to the above game.

Answer:

...

Part B. Prove that your construction is secure according to the above definition of selective security.

Answer:

...

Resource and Collaborator Statement:

...