CS 276: Homework 3

Due Date: Friday September 27th, 2024 at 8:59pm via Gradescope

This problem is based on [DY04, BMR10].

1 A Pseudorandom Function Based on Diffie-Hellman

Let us construct a more efficient variant of the Naor-Reingold PRF.

Definition 1.1 (PRF Construction) *Let $\mathbb{G}$ be a cryptographic group of prime order p. Let $\ell \in \mathbb{N}$ be polynomial in λ. Next, let $s^{*n} = (s_1, \ldots, s_n, h)$ be sampled from $\mathcal{S}^{*n} := \mathbb{Z}_p^n \times \mathbb{G}$, and let $x^{*n} = (x_1, \ldots, x_n)$ be drawn from $\mathcal{X}^{*n} = [\ell]^n$. Finally, define $F^{*n} : \mathcal{S}^{*n} \times \mathcal{X}^{*n} \to \mathbb{G}$ as follows:*

$$F^{*n}(s^{*n}, x^{*n}) = \begin{cases} 1, & \prod_{i\in[n]}(s_i + x_i) = 0 \\ h^{1/\prod_{i\in[n]}(s_i+x_i)}, & \textit{else} \end{cases}$$

This construction is more efficient than Naor-Reingold's PRF. F^{*n} can handle an input x^{*n} of length $n \cdot \lg(\ell)$ bits, whereas the same seed in the Naor-Reingold PRF would handle inputs of length n bits.

Question: Prove that the function F^{*n} given in definition 1.1 is a secure PRF assuming the ℓ-DDH assumption (assumption 1.2).

Assumption 1.2 (ℓ-DDH Assumption) *Let $\mathbb{G}$ be a cryptographic group of prime order p, and let $\ell < p$. Then for any PPT adversary $\mathcal{A}$, the following two hybrids are indistinguishable:*

- $\mathcal{G}_0$*: The challenger samples $(\alpha, g) \xleftarrow{\$} \mathbb{Z}_p \times \mathbb{G}$ and then gives the adversary $(g, g^{\alpha}, g^{\alpha^2}, \ldots, g^{\alpha^\ell}, g^{1/\alpha})$.*
- $\mathcal{G}_1$*: The challenger samples $(\alpha, g, r) \xleftarrow{\$} \mathbb{Z}_p \times \mathbb{G} \times \mathbb{G}$ and then gives the adversary $(g, g^{\alpha}, g^{\alpha^2}, \ldots, g^{\alpha^\ell}, r)$.*

Finally, when $\alpha = 0$, then define $g^{1/\alpha} = 1$.

Note that p must be super-polynomial in λ because otherwise ℓ-DDH does not hold.

Hint: You may wish to use the following strategy. First, let us define a PRF f over a smaller domain $[\ell]$. Let f take a seed $(s, h) \in \mathbb{Z}_p \times \mathbb{G}$ and an input $x \in [\ell]$ and output:

$$f((s,h),x) = \begin{cases} 1, & s + x = 0 \\ h^{1/(s+x)}, & \text{else} \end{cases}$$

First prove that f is a secure PRF when ℓ is polynomial in the security parameter λ.

Second, note that F^{*n} is an n-fold composition of f, where the output of one invocation of f becomes the h-value of the next invocation of f.

$$F^{*n}\big((s_1, \ldots, s_n, h), (x_1, \ldots, x_n)\big) = f((s_n, \ldots f((s_2, f((s_1, h), x_1)), x_2) \ldots), x_n)$$

Then use a similar proof technique to the one used for Naor-Reingold's PRF to prove that the composition of this small-domain PRF f is also a PRF.

Solution

Theorem 1.3 *f is a secure PRF.*

Proof. We will create $\ell+1$ hybrids, and each new hybrid will take a different input $x \in [\ell]$ and switch $f(x)$ from pseudorandom to random. Each successive hybrid is distinguishable from the one before it with only negligible advantage. Since ℓ is polynomial in λ, $\mathcal{H}_0$ and $\mathcal{H}_\ell$ will be distinguishable with only negligible advantage as well.

- $\mathcal{H}_0$ is the PRF security game for f. The challenger samples $(s,h) \xleftarrow{\$} \mathbb{Z}_p \times \mathbb{G}$. Then $\mathcal{A}$ submits a query $x \in [\ell]$, and the challenger responds with

$$F(x) = \begin{cases} 1, & s+x=0 \\ h^{1/(s+x)}, & \text{else} \end{cases}$$

  The adversary may submit many queries. Finally, the adversary outputs a bit b, which is the output of the hybrid.

Then for every $x \in [\ell]$, let $\mathcal{H}_x$ be defined as follows:

- $\mathcal{H}_x$ is the PRF security game for f except inputs $\leq x$ are reprogrammed to random values. The challenger samples $(s,h) \xleftarrow{\$} \mathbb{Z}_p \times \mathbb{G}$ as well as $(r_1,\ldots,r_x) \xleftarrow{\$} \mathbb{G}^x$. Then $\mathcal{A}$ submits a query $x' \in [\ell]$, and the challenger responds with

$$F(x') = \begin{cases} r_{x'}, & x' \leq x \\ 1, & s+x'=0 \\ h^{1/(s+x')}, & \text{else} \end{cases}$$

  The adversary may submit many queries. Finally, the adversary outputs a bit b, which is the output of the hybrid.

Note that in $\mathcal{H}_\ell$, every input receives a uniformly random respondse $r_{x'}$.

Lemma 1.4 *For any $x \in [\ell]$ and any PPT adversary $\mathcal{A}$, $\left|\Pr[\mathcal{H}_{x-1} \to 1] - \Pr[\mathcal{H}_x \to 1]\right| \leq \mathsf{negl}(\lambda)$.*

Proof. Given an adversary $\mathcal{A}_{PRF}$ for which $\left|\Pr[\mathcal{H}_{x-1} \to 1] - \Pr[\mathcal{H}_x \to 1]\right|$ is non-negligible, we can construct an adversary $\mathcal{A}_{DDH}$ that breaks the ℓ-DDH assumption.

Construction of $\mathcal{A}_{DDH}$:

1. Receive $(g, g^\alpha, g^{\alpha^2}, \ldots, g^{\alpha^\ell}, G)$, where $G = g^{1/\alpha}$ or $G = r$ for $(\alpha, g, r) \xleftarrow{\$} \mathbb{Z}_p \times \mathbb{G} \times \mathbb{G}$.
2. For a variable $A \in \mathbb{Z}_p$ and any $x' \in [\ell]$, compute the coefficients of the following polynomials.

$$p(A) = \prod_{x'' \in [\ell] : x'' > x} (A - x + x'') = \sum_{i=0}^{\ell-1} c_i \cdot A^i$$

$$p_{x'}(A) = \frac{p(A)}{A - x + x'} = \sum_{i=0}^{\ell-2} d_{x',i} \cdot A^i$$

3. Let $s = \alpha - x$. s is well-defined, even though $\mathcal{A}_{DDH}$ does not know α and cannot directly compute s. Then compute:
$$h = \prod_{i=0}^{\ell-1} \left(g^{\alpha^i}\right)^{c_i} = g^{p(\alpha)}$$
4. For each $x' < x$, sample $r_{x'} \xleftarrow{\$} \mathbb{G}$, and set $F(x') = r_{x'}$.
5. Set
$$F(x) = G^{c_0} \cdot \prod_{i=1}^{\ell-1} \left(g^{\alpha^{i-1}}\right)^{c_i}$$
6. For each $x' > x$, compute
$$h^{1/(s+x')} = g^{p(\alpha)/(\alpha-x+x')} = g^{p_{x'}(\alpha)} = \prod_{i=0}^{\ell-2} \left(g^{\alpha^i}\right)^{d_{x',i}}$$
and set $F(x') = h^{1/(s+x')}$.
7. Run $\mathcal{A}_{PRF}$. Respond to any queries x' with the value of $F(x')$ that was computed earlier. When $\mathcal{A}_{PRF}$ outputs a bit b, $\mathcal{A}_{DDH}$ outputs b as well.

Analysis: $\mathcal{A}_{DDH}$ correctly simulates $\mathcal{H}_{x-1}$ when $G = g^{1/\alpha}$ and $\mathcal{H}_x$ when $G = r$.

1. The (s, h)-values computed by $\mathcal{A}_{DDH}$ are uniformly random over $\mathbb{Z}_p \times \mathbb{G}$ due to the randomness of α and g.
2. When $G = g^{1/\alpha}$,
$$\begin{aligned} F(x) &= \left(g^{\alpha^{-1}}\right)^{c_0} \cdot \prod_{i=1}^{\ell-1} \left(g^{\alpha^{i-1}}\right)^{c_i} = \prod_{i=0}^{\ell-1} \left(g^{\alpha^{i-1}}\right)^{c_i} \\ &= g^{p(\alpha)/\alpha} = g^{p(\alpha)/(s+x)} \\ &= h^{1/(s+x)} \end{aligned}$$
On the other hand, when $G = r$, then $F(x)$ is uniformly random and independent of $F(x')$ for any $x' \neq x$.
3. Finally, with overwhelming probability, $s + x' \neq 0$ for all $x' \in [\ell]$. This is because $s \xleftarrow{\$} \mathbb{Z}_p$, and p is superpolynomial in λ. So with overwhelming probability, in $\mathcal{H}_x$ or $\mathcal{H}_{x-1}$, the adversary will never query F on an input $x' \in [\ell]$ such that $s + x' = 0$.
4. This shows that $G = g^{1/\alpha}$, $\mathcal{A}_{DDH}$'s messages to $\mathcal{A}_{PRF}$ are statistically close to the messages $\mathcal{A}_{PRF}$ receives in $\mathcal{H}_{x-1}$, and when $G = r$, $\mathcal{A}_{DDH}$'s messages to $\mathcal{A}_{PRF}$ are statistically close to the messages $\mathcal{A}_{PRF}$ receives in $\mathcal{H}_x$.

5. If there exists an $\mathcal{A}_{PRF}$ such that $\big|\Pr[\mathcal{H}_{x-1} \to 1] - \Pr[\mathcal{H}_x \to 1]\big|$ is non-negligible, then $\mathcal{A}_{DDH}$ distinguishes $\mathcal{G}_0$ and $\mathcal{G}_1$ with non-negligible advantage. This would contradict the assumed hardness of ℓ-DDH. Therefore, in fact, for any PPT $\mathcal{A}_{PRF}$, $\big|\Pr[\mathcal{H}_{x-1} \to 1] - \Pr[\mathcal{H}_x \to 1]\big| \leq \mathsf{negl}(\lambda)$.

■

Next, for any PPT $\mathcal{A}_{PRF}$,

$$\big|\Pr[\mathcal{H}_0 \to 1] - \Pr[\mathcal{H}_\ell \to 1]\big| \leq \ell \cdot \mathsf{negl}(\lambda) = \mathsf{negl}'(\lambda)$$

Here, we used the fact that $\ell = \text{poly}(\lambda)$, and $\text{poly}(\lambda) \cdot \mathsf{negl}(\lambda)$ is negligible.

Finally, note that $\mathcal{H}_0$ and $\mathcal{H}_\ell$ are exactly the hybrids that the adversary is asked to distinguish in the PRF security game for f. Therefore, f is a secure PRF. ■

It remains to show that if f is a secure PRF and DDH is hard, then F^{*n} is also a secure PRF. The proof is given in [BMR10], theorem 7. ■

References

[BMR10] Dan Boneh, Hart Montgomery, and Ananth Raghunathan. Algebraic pseudorandom functions with improved efficiency from the augmented cascade. Cryptology ePrint Archive, Paper 2010/442, 2010.

[DY04] Yevgeniy Dodis and Aleksandr Yampolskiy. A verifiable random function with short proofs and keys. Cryptology ePrint Archive, Paper 2004/310, 2004.