TD 4: LWE and PRFs (corrected version)

Exercise 1. *Around the DDH assumption*

We recall the definition of the DDH assumption.

Definition 1 (Decisional Diffie-Hellman distribution). *Let $\mathbb{G}$ be a cyclic group of (prime) order p, and let g be a public generator of $\mathbb{G}$. The* decisional Diffie-Hellman *distribution (DDH) is, $D_{\text{DDH}} = (g^a, g^b, g^{ab}) \in \mathbb{G}^3$ with a, b sampled independently and uniformly in $\mathbb{Z}/p\mathbb{Z} =: \mathbb{Z}_p$.*

Definition 2 (Decisional Diffie-Hellman assumption). *The decisional Diffie-Hellman assumption states that there exists no probabilistic polynomial-time distinguisher between D_{DDH} and (g^a, g^b, g^c) with a, b, c sampled independently and uniformly at random in $\mathbb{Z}_p$.*

1. Does the DDH assumption hold in $\mathbb{G} = (\mathbb{Z}_p, +)$ for $p = \mathcal{O}(2^\lambda)$ prime? ☞ No. In this case, the D_{DDH} distribution is $(a \cdot g, b \cdot g, (ab) \cdot g)$. This can be distinguished from (ag, bg, cg) by computing the inverse of g (find a Bézout identity $gu + pv = 1$ in logarithmic time), retrieving a, b and c and checking whether $ab = c$ or not. This is always the case in the DDH, and the case with probability $1/p$ in the uniform case. The advantage of a distinguisher returning the boolean value of $ab = c$ is then $\frac{p-1}{p}$.

2. Same question for $\mathbb{G} = (\mathbb{Z}_p^\star, \times)$ of order $p - 1$, with p an odd prime.

   ☞ No, because $p - 1$ (the order the group) is divisible by 2.

   We know that $x^{\frac{p-1}{2}} = 1$ if $x \in \mathbb{Z}_p$ is a square and -1 otherwise (it is actually the Legendre symbol: $\left(\frac{x}{p}\right)$ and can be efficiently computed). So $\left(\frac{g^a}{p}\right)$ gives us the parity of a, that is $\left(\frac{g^a}{p}\right) = 1$ if a is even and $\left(\frac{g^a}{p}\right) = -1$ if a is odd. Hence, if a is uniformly sampled in $\{0, \cdots, p-1\}$ (meaning that g^a is uniformly sampled in $\mathbb{G}$), then $\left(\frac{g^a}{p}\right)$ is uniformly distributed in $\{-1, 1\}$. But in the case of the DDH distribution, if a or b is even, then ab must be even too (or equivalently, it g^a or g^b is a square, then g^{ab} should be a square too). In the same way, if both a and b are odd, then ab must be odd.

   This enables us to build the following distinguisher $\mathcal{A}$:

   - Return DDH if $\left(\frac{g^{ab}}{p}\right)$ is consistent with $\left(\frac{g^a}{p}\right)$ and $\left(\frac{g^b}{p}\right)$ (i.e. ab is odd and both a and b are odd, or ab is even and a or b is even);
   - Return Unif otherwise.

   Let us now compute the advantage of such a distinguisher.

   $$\begin{aligned} \mathrm{Adv}^{DDH}(\mathcal{A}) &= |\Pr[\mathcal{A} \to DDH \mid DDH] - \Pr[\mathcal{A} \to DDH \mid Unif]| \\ &= |1 - \Pr[\mathcal{A} \to DDH \mid Unif]| \end{aligned}$$

   Our distinguisher returns Unif only is c is odd and either a or b is even of if c is even and both a and b are *odd*. But we have seen that these cases could not appear in the DDH distribution. So we have that $\Pr[\mathcal{A} \to DDH \mid DDH] = 1$.

   It then remains to compute $\Pr[A \to DDH \mid Unif]$. Given a Unif instance (g^a, g^b, g^c), we have seen that $\left(\frac{g^a}{p}\right)$, $\left(\frac{g^b}{p}\right)$ and $\left(\frac{g^c}{p}\right)$ are uniform in $\{-1, 1\}$ because a, b, c are uniform in $\{0, \cdots, q-1\}$. They are also independent because a, b and c are. So all eight possibilities for $(\left(\frac{g^a}{p}\right), \left(\frac{g^b}{p}\right), \left(\frac{g^c}{p}\right))$ have the same probability and we have

   $$\begin{aligned} \Pr[\mathcal{A} \to DDH \mid Unif] &= \Pr[(\left(\tfrac{g^a}{p}\right), \left(\tfrac{g^b}{p}\right), \left(\tfrac{g^c}{p}\right)) = (1,1,1) \text{ or } (1,-1,1) \text{ or } (-1,1,1) \text{ or } (-1,-1,-1)] \\ &= \tfrac{4}{8} = \tfrac{1}{2} \end{aligned}$$

   To conclude, we have $\mathrm{Adv}(\mathcal{A}) = \frac{1}{2}$, which is non-negligible.

   It remains to show that our distinguisher is PPT. This is the case because it just needs to compute $\left(\frac{h}{p}\right) = h^{\frac{p-1}{2}}$ for three different elements h of $\mathbb{G}$. Computing $h^{\frac{p-1}{2}}$ can be done by fast exponentiation, resulting in at most $\log(p)$ multiplications in $\mathbb{Z}_p$. Each such multiplication takes a time polynomial in $\log(p)$, and so our distinguisher $\mathcal{A}$ is indeed polynomial time (in $\log(p)$).

   Remark. The same reasoning can be adapted if the cardinality of the cyclic group $\mathbb{G}$ is $n = km$ for some small k (and any m). In that case, we would have that $(g^a)^m$ is uniformly distributed among $\{1, g^m, g^{2m}, \cdots, g^{(k-1)m}\}$ if a is uniform, and computing $(g^a)^m$ gives us the value of $a \mod k$. We then can check whether $a \mod k$ and $b \mod k$ are coherent with $ab \mod k$. In the DDH case, this will always be coherent whereas in the uniform case, this will be coherent only with probability $\frac{k-1}{k}$. We hence obtain a distinguisher with advantage $\frac{k-1}{k}$ (which is non negligible for any $k \geq 2$) and whose computation time is $\Theta(k \cdot \mathrm{poly}(\log(n)))$. So if k is polynomial in $\log(n)$, this gives us a polynomial time distinguisher with non negligible advantage. This is why, in the next question, we consider of group of prime cardinality.

But this implies knowing a factorisation of n.

Exercise 2. *PRG from LWE*

We recall the Learning with Errors assumption.

Definition 3 (Learning with Errors). *Let $q \in \mathbb{N}$, $B \in \mathbb{N}$, $\mathbf{A} \hookleftarrow U(\mathbb{Z}_q^{m\times n})$. The Learning with Errors (LWE) distribution is defined as follows: $D_{\text{LWE}} = (\mathbf{A}, \mathbf{A}\cdot\mathbf{s} + \mathbf{e} \bmod q)$ for $\mathbf{s} \hookleftarrow U(\mathbb{Z}_q^n)$, $\mathbf{A} \hookleftarrow U(\mathbb{Z}_q^{m\times n})$ and $\mathbf{e} \hookleftarrow U((-B, B]^m)$.*

In this setting, the vector $\mathbf{s}$ is called the secret, and $\mathbf{e}$ the noise.

Remark. If q and B are powers of 2, we are manipulating bits, contrary to the DDH-based PRG from the lecture.

The *LWE assumption* states that, given suitable parameters q, B, m, n, it is computationally hard to distinguish D_{LWE} from the distribution $U(\mathbb{Z}_q^{m\times n} \times \mathbb{Z}_q^m)$.

Let us propose the following pseudo-random generator: $G(\mathbf{A}, \mathbf{s}, \mathbf{e}) = (\mathbf{A}, \mathbf{A}\cdot\mathbf{s} + \mathbf{e} \bmod q)$.

1. By definition, a PRG must have a bigger output size than input size. Give a bound on B that depends on the other parameters if we want G to satisfy this.

   ☞ We want the parameters to satisfy $q^{mn} \cdot q^n B^m \leq q^{nm} \cdot q^m$ i.e. $B^m \leq q^{m-n}$. Then the bound is $B \leq q^{1-n/m}$.

2. Given suitable B, q, n, m such that the LWE assumption and previous bound hold, show that G is a secure pseudo-random generator.

   ☞ Let $\mathcal{A}$ be a PPT adversary that distinguishes with non negligible advantage the output of G from the uniform distribution. Let us use this adversary to solve the LWE problem.

   At the beginning of the game, the reduction $\mathcal{B}$ receives a LWE instance $(\mathbf{A}, \mathbf{b}) \in \mathbb{Z}_q^{m\times n} \times \mathbb{Z}_q^m$ of the LWE problem, the goal is to output `LWE` if it is a LWE instance, and `Unif` if it is uniform.

   The reduction sends $(\mathbf{A}, \mathbf{b})$ to the adversary $\mathcal{A}$ against the PRG. The adversary then returns a bit b' that the reduction returns to its challenger.

   Analysis. $\mathrm{Adv}^{\mathrm{LWE}}(\mathcal{B}) = |\Pr[B \to 1 | b \text{ LWE}] - \Pr[B \to 1 | b \text{ Unif}]| = |\Pr[A \to 1 | b \text{ LWE}] - \Pr[A \to 1 \mid b \text{ Unif}]| = \mathrm{Adv}^{\mathrm{PRG}}(\mathcal{A}) =$ non negl.

Exercise 3. *LWE with small secret*

We once more work in the setting of the LWE assumption. Let q, B, n, m such that the LWE assumption holds. Moreover, we assume that q is prime.

1. **(a)** What is the probability that $\mathbf{A}_1 \in \mathbf{Z}_q^{n\times n}$ is invertible where $\mathbf{A} =: [\mathbf{A}_1^\top | \mathbf{A}_2^\top]^\top$ is uniformly sampled?

   ☞ We have to compute $|GL_n(\mathbb{F}_q)|$, i.e. the number of invertibles matrices with coefficients in $\mathbb{F}_q$. We have $q^n - 1$ choice for the first vector (it can be any vector except the 0 vector), then $q^n - q^1$ for the second vector (anything except a vector collinear to the first one), then $q^n - q^2$ (anything that is not a linear combination of the first two vectors), etc. So we get

$$\Pr_{A_1 \hookleftarrow U(\mathbb{F}_2^{m\times n})}[A_1 \in GL_n(\mathbb{F}_q)] = \frac{1}{q^{n^2}} \prod_{i=0}^{n-1} (q^n - q^i) = \prod_{i=0}^{n-1} (1 - q^{i-n}),$$

   which is always $\geq \prod_{i=0}^{n-1} (1 - 2^{i-n}) \geq 0.288$.

   (b) Assume that $m \geq 2n$. Prove that there exists a subset of n lineraly independent rows of $\mathbf{A} \hookleftarrow U(\mathbb{Z}_q^{m\times n})$ with probability $\geq 1 - 1/2^{\Omega(n)}$ and that we can find them in polynomial time.

☞ If this is not the case, then there exists an hyperplane of $\mathbb{Z}_q^n$ in which each row is sampled. A hyperplane is given by a nonzero vector: there are at most $q^n - 1$ hyperplanes of the space and for a given hyperplane, the probability that each vector falls into it is $q^{(n-1)m}/q^{nm} = 1/q^m$. Then the union bound gives us that the probability is $\geq 1 - \frac{1}{q^{m-n}} \geq 1 - \frac{1}{q^n}$.

To find such rows, the naive greedy algorithm works: select the first row. Then, repeat the following for $i = 2$ to m. If the i-th row is linearly independent from the selected rows, select it.

2. Let us define the distribution $D_B = U((-B, B] \cap \mathbb{Z})$, and $m' = m - n$.

Show that under the $\mathrm{LWE}_{q,m,n,B}$ assumption, the distributions $(\mathbf{A}', \mathbf{A}'\mathbf{s}' + \mathbf{e}') \in \mathbb{Z}_q^{m' \times n} \times \mathbb{Z}_q^{m'}$, with $\mathbf{s}' \hookleftarrow D_B^n$ and $\mathbf{e}' \hookleftarrow D_B^{m'}$, and $(\mathbf{A}', \mathbf{b}')$ with $\mathbf{b}' \leftarrow U(\mathbb{Z}_q^{m'})$ are indistinguishable.

☞ We show how to reduce an instance of the decision problem $\mathrm{LWE}_{q,m,n,B}$ to an instance of this new decision problem. Let $(\mathbf{A}, \mathbf{b}) \in \mathbf{Z}_q^{m \times n} \times \mathbf{Z}_q^m$. With non negligible probability and up to permuting the rows of $\mathbf{A}$ (and $\mathbf{b}$), one can write $\mathbf{A} = \begin{bmatrix} \mathbf{A}_1 \\ \mathbf{A}_2 \end{bmatrix}$, where $\mathbf{A}_1 \in \mathbf{Z}_q^{n \times n}$ is invertible.

Notice that in this case, $\mathbf{A}_2\mathbf{A}_1^{-1} \in \mathbb{Z}_q^{m' \times n}$ is still uniform because $\mathbf{A}_1$ is invertible, and $\mathbf{A}_2$ is uniformly sampled.

Assume that we are given a sample $(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e})$ of the $\mathrm{LWE}_{q,m,n,B}$ distribution. Set $\mathbf{e} =: (\mathbf{e}_1^\top, -\mathbf{e}_2^\top)^\top$ Consider the following:

$$(\mathbf{A}_2\mathbf{A}_1^{-1}, \mathbf{A}_2\mathbf{A}_1^{-1}(\mathbf{A}_1\mathbf{s} + \mathbf{e}_1) - \mathbf{A}_2\mathbf{s} + \mathbf{e}_2) = (\mathbf{A}_2\mathbf{A}_1, \mathbf{A}_2\mathbf{A}_1^{-1}\mathbf{e}_1 + \mathbf{e}_2).$$

This is exactly a sample from the new distribution, with secret $\mathbf{e}_1$ and noise $\mathbf{e}_2$.

Assume now that we are given a sample $(\mathbf{A}, \mathbf{b})$ where $\mathbf{b}$ is uniformly sampled. We write $\mathbf{b} =: (\mathbf{b}_1^\top, \mathbf{b}_2^\top)^\top$. With the previous transformation we get: $\mathbf{A}_2\mathbf{A}_1^{-1}, \mathbf{A}_2\mathbf{A}_1^{-1}\mathbf{b}_1 - \mathbf{b}_2$. Whatever $\mathbf{A}_2\mathbf{A}_1^{-1}\mathbf{b}_1$ is, since it is independent from $\mathbf{b}_2$, we get a uniform sample over $\mathbb{Z}_q^{m' \times n} \times \mathbb{Z}_q^{m'}$.

This means that any distinguisher for the new decision problem is a distinguisher for decision LWE. Under the LWE assumption, any efficient distinguisher has negligible advantage and this concludes the proof.