TUTORIAL 2

1 Hashing with SIS (⋆⋆)

The objective of this exercise is to study a construction of a collision resistant hash function based on SIS.

Let F be a family of functions from a set X to a set Y (which we will call "hash functions", but really they are just functions) and let D_F be a distribution over this set of functions.

Definition: The advantage of a probabilistic polynomial time (p.p.t.) algorithm $\mathcal{A}$ against the collision resistance of the family of hash functions (F, D_F) is defined as

$$\mathrm{Adv}_F(\mathcal{A}) := \Pr_{f \leftarrow D_F}\Big(\mathcal{A}(f) = (x, x') \in X^2 \text{ with } f(x) = f(x') \text{ and } x \neq x'\Big),$$

where the probability is taken over the random choice of f and the internal randomness of $\mathcal{A}$.

Recall also the SIS problem, which is as follows.

Definition: Let q, m, n be integers with $m \geq n$ and $B > 0$ be some bound. The advantage of a p.p.t. adversary $\mathcal{A}$ against the $\mathrm{SIS}_{q,n,m,B}$ problem is defined as

$$\mathrm{Adv}_{\mathrm{SIS}}(\mathcal{A}) := \Pr_{A \leftarrow \mathcal{U}(\mathbb{Z}_q^{m \times n})}\Big(\mathcal{A}(A) = x \in \mathbb{Z}^m \text{ with } x^T \cdot A = 0 \bmod q \text{ and } 0 < \|x\| \leq B\Big),$$

where the probability is over the random choice of A and the internal randomness of $\mathcal{A}$.

We will consider the following family F of functions, from $\{0,1\}^m$ to $\mathbb{Z}_q^n$. The functions of F are indexed by a matrix $A \in \mathbb{Z}_q^{m \times n}$ and are defined as

$$\begin{aligned} f_A : \{0,1\}^m &\to \mathbb{Z}_q^n \\ x &\mapsto x^T \cdot A \end{aligned}$$

The distribution D_F over F is obtained by sampling $A \in \mathbb{Z}_q^{m \times n}$ uniformly at random and outputting f_A.

1. Assume that $B \geq \sqrt{m}$. Show that if there exists an adversary $\mathcal{A}$ against the collision resistance of (F, D_F) with advantage $\varepsilon > 0$, then there exists an adversary $\mathcal{B}$ against the $\mathrm{SIS}_{q,n,m,B}$ problem with advantage $\geq \varepsilon$. This proves that (F, D_F) is a family of collision resistant functions, provided that the SIS problem is hard.

   A: Let us assume that there is an adversary $\mathcal{A}$ as in the question and construct an adversary $\mathcal{B}$ against SIS. The algorithm $\mathcal{B}$ gets as input some uniformly random matrix $A \in \mathbb{Z}_q^{m \times n}$. It sends to $\mathcal{A}$ the function f_A. The adversary $\mathcal{A}$ outputs a pair $(x, x') \in \{0,1\}^m \times \{0,1\}^m$ and $\mathcal{B}$ finally outputs the element $z = x - x'$.

   Observe first that the view of $\mathcal{A}$ is exactly the same as in the true collision-resistant game. Hence, the probability that $\mathcal{A}$ outputs $x, x' \in \{0,1\}^m$ with $x \neq x'$ and $f_A(x) = f_A(x')$ is $\mathrm{Adv}_F(\mathcal{A}) = \varepsilon$.

   The second observation is that when $\mathcal{A}$ succeeds in finding a collision, then $\mathcal{B}$ succeeds in computing a solution to SIS. Indeed, since $x^T \cdot A = f_A(x) = f_A(x') = (x')^T \cdot A$ (all equalities are modulo q), we have $z^T \cdot A = 0 \bmod q$. Moreover, since $x \neq x'$, then $z \neq 0$. Finally, since x and x' have coefficients in $\{0,1\}$, then $z = x - x'$ has coefficients in $\{-1, 0, 1\}$. Hence, we have $\|z\| \leq \sqrt{m} \leq B$, where the last inequality comes from the assumption in the question. We conclude that z is a solution to SIS with parameters q, m, n and B, and the success probability of $\mathcal{B}$ is at least the same as the one of $\mathcal{A}$, i.e., ε.

2 QR-factorization (⋆⋆)

The objective of this exercise is to define the QR factorization of a matrix and prove useful properties of this decomposition, which will be used in exercise 3.

In this exercise, we admit the following result:
Lemma: There exists a polynomial time algorithm that takes as input any matrix $B \in \mathrm{GL}_n(\mathbb{R})$, and outputs two matrices $Q, R \in \mathrm{GL}_n(\mathbb{R})$ such that

- $B = Q \cdot R$;
- Q is orthonormal, i.e., $Q^{-1} = Q^T$;
- R is upper triangular and has non negative diagonal coefficients.

The pair (Q, R) is called a *QR-factorization* of the matrix B. We will see below that it is unique. In the rest of this exercise sheet, it might be useful to remember that an orthonormal matrix Q has the following properties:

- all the rows and columns of the matrix Q have euclidean norm 1;
- the rows (resp. columns) of Q are orthogonal;
- for any vector v it holds that $\|Qv\| = \|v\|$.

1. Let $B \in \mathrm{GL}_n(\mathbb{R})$. Show that the QR-factorization of B is unique (i.e., show that if $B = QR = Q'R'$ with Q, Q' orthonormal and R, R' upper triangular with positive diagonal coefficients, then $Q = Q'$ and $R = R'$) (⋆⋆)

   A: Let Q, Q', and R, R' be as in the question and such that $QR = Q'R'$. Rewriting the equality, we have $\tilde{Q} = \tilde{R}$, where $\tilde{Q} = (Q')^{-1} \cdot Q$ and $\tilde{R} = R' \cdot R^{-1}$.

   Observe that the set of orthonormal matrices is stable by inversion and multiplication. Hence $\tilde{Q}$ is orthonormal. Similarly, the set of upper triangular matrices with positive diagonal coefficients is stable by inversion and multiplication, hence $\tilde{R}$ is upper triangular with positive diagonal coefficients.

   We will show that the intersection of the set of orthonormal matrices with the set of upper triangular matrices with positive diagonal coefficients only contains I_n, which will prove the equality $Q = Q'$ and $R = R'$.

   Let $\tilde{Q} = \tilde{R}$ be a matrix which is both orthonormal and upper-triangular with positive diagonal coefficients. Then $\tilde{R}^T = \tilde{Q}^T = \tilde{Q}^{-1} = \tilde{R}^{-1}$, where we used the fact that the transpose of an orthonormal matrix is its inverse. But since $\tilde{R}$ is upper triangular, we know that its inverse is also upper-triangular and its transpose is lower-triangular. Since both are equal, the matrix must be diagonal.

   Let us now prove that the diagonal coefficients are all equal to 1. This comes from the fact that the euclidean norm of every column of an orthonormal matrix is 1. Since this norm is equal to the absolute value of the diagonal coefficient (which is the only non-zero coefficient in each column), this coefficient must be ± 1. Using the fact that $\tilde{R}$ has positive diagonal coefficients, we conclude that they must be all 1.

   We say that a basis B of a lattice is *size-reduced* if its QR-factorization (Q, R) satisfies the following property: for all $j \geq i$, $|r_{i,j}| \leq r_{i,i}$ (remember that $r_{i,i} > 0$). In other words, the diagonal coefficients of R are the largest coefficients of their rows (in absolute value).

2. Let $B \in \mathrm{GL}_n(\mathbb{R})$ and (Q, R) be its QR-factorization. Show that there exists an efficiently computable unimodular matrix U such that $B \cdot U$ is size-reduced and has QR-factorization (Q, R') with $r'_{i,i} = r_{i,i}$ for all i. (⋆⋆)
   (You do not have to describe the algorithm very properly, getting the idea is sufficient.)

A: This transformation, which consist in reducing the non-diagonal coefficients modulo the diagonal coefficients is a very common operation performed on lattices bases (for instance in the LLL algorithm). It is usually called size-reduction. It allows in particular to avoid the explosion of the size of the coefficients during the execution of multiple algorithms.

This transformation is done on the columns of R by operations like $C_j \leftarrow C_j + \lfloor r_{i,j}/r_{i,i} \rceil C_i$ for all $j \geq i$. This reduces the non-diagonal coefficients modulo the diagonal coefficients, hence it ensures that all the coefficients on a row are smaller (in absolute value) than the diagonal coefficient $r_{i,i}$. These operations are unimodular since they can be inverted by performing only integer operations, and they preserve the diagonal coefficients, as desired. (One needs to perform these operations in an appropriate order, otherwise reduced coefficients might be increased again afterwards, but this is doable).

In the rest of this exercise sheet, we call `size_reduce` the polynomial time algorithm that takes as input a matrix B and returns a sized-reduced matrix $B' := B \cdot U$ as in the above question, i.e., with $r_{i,i} = r'_{i,i}$ and $\mathcal{L}(B') = \mathcal{L}(B)$.

3. Let $B \in \mathrm{GL}_n(\mathbb{R})$ and (Q, R) be its QR-factorization. Let b_j be the column vectors of B. Show that $\max_j r_{j,j} \leq \max_j \|b_j\|$. If B is size-reduced, show that we also have the inequality $\max_j \|b_j\| \leq \sqrt{n} \cdot \max_j r_{j,j}$ (in other words, the size of the diagonal coefficients of R are a relatively good approximation of the size of the vectors of B when B is size-reduced). $(\star\star)$
(Hint 1: observe that $b_j = Q \cdot r_j$ with r_j the j-th column of R)
(Hint 2: remember the property that $\|Qv\| = \|v\|$ for any vector v)

A: Let us first show that $\max_j r_{j,j} \leq \max_j \|b_j\|$. We will actually show the stronger property $r_{j,j} \leq \|b_j\|$ for all j's. Fix some column index j. Since $B = Q \cdot R$, then $b_j = Q \cdot r_j$, where r_j is the j-th column of R. Moreover, since Q is orthonormal, then $\|b_j\| = \|r_j\|$. Finally, note that $\|r_j\| \leq |r_{j,j}| = r_{j,j}$ (since the diagonal coefficients are positive), which concludes the proof of the first inequality.

For the second inequality, we use again the fact that $\|b_j\| = \|r_j\|$. A closer look at r_j shows that $\|r_j\| \leq \sqrt{j} \cdot \max_{i \leq j} |r_{i,j}| \leq \sqrt{n} \cdot \max_{i \leq j} r_{i,i}$ (in the last inequality we used the fact that the basis is size-reduced). From this, we conclude that $\|b_j\| = \|r_j\| \leq \sqrt{n} \cdot \max_i r_{i,i}$ as desired.

3 Computing a short basis from a short generating set $(\star\star)$

The objective of this exercise is to show that given an arbitrary basis B of a lattice $\mathcal{L}$ and a set of n linearly independent (short) vectors S in $\mathcal{L}$, then one can create a new basis $\tilde{B}$ of $\mathcal{L}$ with vectors of length not much larger than the ones of S. In other words, finding short linearly independent vectors in $\mathcal{L}$ is sufficient to obtain a short basis of $\mathcal{L}$.
This exercise uses results from exercise 2.

1. Let B be a basis of a lattice $\mathcal{L}$ and $S \in \mathrm{GL}_n(\mathbb{R})$ be a set of n linearly independent vectors in $\mathcal{L}$. Make sure you remember why there exists an integer matrix X such that $S = B \cdot X$. Is X unimodular?

A: Every column vector of S belongs to $\mathcal{L}$, hence is an integer linear combination of the columns of B. Hence $S = B \cdot X$ with X integer. The matrix X is unimodular (i.e., has an integral inverse) if and only if S is a basis of $\mathcal{L}$ (which might not be the case here).

2. Let Y be the HNF basis of the lattice $\mathcal{L}(X^T)$ and let U be the unimodular matrix such that $X^T = Y \cdot U$. Verify that $B' = B \cdot U^T$ is a basis of $\mathcal{L}$ and that $S = B' \cdot Y^T$.

A: Since U is unimodular, then so is U^T (it has integer coefficients and determinant ± 1). Hence, B' is indeed a basis of $\mathcal{L}$. Moreover, since $S = B \cdot X$ and $X = U^T \cdot Y^T$, then we indeed have $S = (B \cdot U^T) \cdot Y^T$ as desired.

3. Let $S = Q_S \cdot R_S$ be the QR factorization of the matrix S and $B' = Q_B \cdot R_B$ be the one of B'. Show that $Q_S = Q_B$ and that $R_S = R_B \cdot Y^T$.
(Hint: use the unicity of the QR-factorization that you proved in exercise 2)

A: From the equality $S = B' \cdot Y^T$, we have $Q_S R_S = Q_B \cdot (R_B \cdot Y^T)$. Note that Y is lower triangular with positive diagonal coefficients (since it is an HNF basis), hence Y^T is upper triangular with positive diagonal coefficients, and so is $(R_B \cdot Y^T)$. We conclude by using the unicity of the QR decomposition which we proved in question 1.

Let $\tilde{B} = \texttt{size\_reduce}(B')$. Our objective is to show that $\tilde{B}$ is a basis of $\mathcal{L}(B)$ which has vectors almost as short as the ones of S. (You can check from the way we defined it that $\tilde{B}$ can be computed in polynomial time from B and S).

4. Let $(\tilde{Q}, \tilde{R})$ be the QR-factorization of $\tilde{B}$. Show that $\max_j \tilde{r}_{j,j} \leq \max_j \|s_j\|$.
(Hint 1: remember from question 2 in exercise 2 that $\tilde{r}_{j,j} = (R_B)_{j,j}$ when we use the size-reduction algorithm)
(Hint 2: observe that the triangular matrix Y is integral and has positive diagonal coefficients, hence its diagonal coefficients are ≥ 1.)

A: Since $\tilde{r}_{j,j} = (R_B)_{j,j}$, it suffices to prove that $\max_j (R_B)_{j,j} \leq \max_j \|s_j\|$.

We have seen in the previous question that $R_S = R_B \cdot Y^T$. Since all those matrices are upper triangular, then the diagonal coefficients satisfy $(R_S)_{j,j} = (R_B)_{j,j} \cdot (Y^T)_{j,j}$ for all j's. But Y^T is an integer matrix, hence its diagonal coefficients are ≥ 1. And we conclude that $(R_B)_{j,j} \leq (R_S)_{j,j}$ (recall that all those diagonal coefficients are positive).

Finally, we use question 3 to conclude that $(R_S)_{j,j} \leq \max_j \|s_j\|$.

5. Conclude that $\tilde{B}$ is a new basis of $\mathcal{L}$ with columns vectors $\tilde{b}_j$ satisfying $\max_j \|\tilde{b}_j\| \leq \sqrt{n} \cdot \max_j \|s_j\|$. In other words, the vectors of $\tilde{B}$ are almost as short as the linearly independent vectors from S.
(Hint: this question consists mainly in combining what you have seen in this exercise and in exercise 2.)

A: From the definition of $\tilde{B}$ and B', one can check that $\mathcal{L}(\tilde{B}) = \mathcal{L}(B') = \mathcal{L}(B)$. Let us now show that $\max_j \|\tilde{b}_j\| \leq \sqrt{n} \cdot \max_j \|s_j\|$. Using question 3 from exercise 2 and the fact that $\tilde{B}$ is size reduced, we see that $\max_j \|\tilde{b}_j\| \leq \sqrt{n} \cdot \max_j \tilde{r}_{j,j}$. From there, we conclude using the previous question.

4 Ideal lattices (⋆⋆)

Let R be the ring $\mathbb{Z}[X]/(X^d+1)$ where d is a power-of-two (so that X^d+1 is irreducible, and $K = \mathbb{Q}[X]/(X^d+1)$ is a field). An ideal in R is a subset I of R such that for all $x, y \in I$, the sum $x + y$ is also in I, and for any $x \in I$ and $\alpha \in R$, the product $x \cdot \alpha$ is in I.

1. Recall that the coefficient embedding

$$\Sigma : K \to \mathbb{Q}^d$$
$$a = \sum_{i=0}^{d-1} a_i X^i \mapsto (a_0, \cdots, a_{d-1})$$

maps elements of K to vectors in $\mathbb{Q}^d$ (and elements of R to vectors in $\mathbb{Z}^d$). Show that if $a \in K$ is non-zero, then the d vectors $\Sigma(a \cdot X^i)$ for $i = 0$ to $d-1$ are linearly independent. (⋆⋆)
(Hint 1: assume you have a $\mathbb{Q}$-linear relation $\sum_{i=0}^{d-1} y_i \cdot \Sigma(a \cdot X^i) = 0$ with the y_i's in $\mathbb{Q}$ and not all zero and try to obtain a contradiction.)
(Hint 2: Σ is a $\mathbb{Q}$-morphism and is a bijection between K and $\mathbb{Q}^d$. Also, K is a field so all non-zero elements are invertible.)

A: Assume by contradiction that the vectors $v_i = \Sigma(a \cdot X^i)$ are not linearly independent. Since $\mathbb{Q}$ is a field containing the v_i's, then there must exist a relation involving the v_i's with coefficients in $\mathbb{Q}$, i.e., there exist $y_0, \cdots, y_{d-1} \in \mathbb{Q}$ not all zero such that $\sum_i y_i \cdot v_i = 0$.

Note that Σ is an additive isomorphism between $K = \mathbb{Q}[X]/(X^d+1)$ and $\mathbb{Q}^d$. Hence, applying Σ^{-1} to the previous equality yields $\sum_i y_i \cdot a \cdot X^i = 0$, i.e., $a \cdot \left(\sum_i y_i \cdot X^i\right) = 0$ (here the operations are performed in $K = \mathbb{Q}[X]/(X^d+1)$, i.e., modulo X^d+1). Let us write $y = \sum_i y_i \cdot X^i \in K$. Since K is a field and $a \cdot y = 0$, then either $a = 0$ or $y = 0$. We assumed that a was non-zero, hence y must be zero. But again, because Σ is an isomorphism, this implies that the y_i's are all 0, which is a contradiction. This shows that the vectors v_i's are indeed linearly independent.

Remember that during the lecture, we have seen that a principal ideal is an ideal of rank d once embedded into $\mathbb{Q}^d$ via the canonical embedding. The objective of the next question is to show that this is true for all ideals (not only the principal ideals).

2. Show that for any non-zero ideal I, the set $\Sigma(I)$ is a lattice of rank d in $\mathbb{R}^d$. ($\star\star$)

A: We use the equivalent definition of a lattice from tutorial 1. First, observe that $\Sigma(I)$ is indeed stable by addition and subtraction (since I is and Σ is an additive morphism). Then, we see that $\Sigma(I)$ is discrete since it is included in $\mathbb{Z}^d$. Finally, let us exhibit d linearly independent vectors in $\Sigma(I)$. Since I is non-zero, it must contain a non-zero element $a \in I$. Moreover, since I is an ideal and $X^i \in R$ for all $i \geq 0$, then the elements $a \cdot X^i$ are in I, i.e., the vectors $\Sigma(a \cdot X^i)$ are in $\Sigma(I)$. We have seen in the previous question that for $i = 0$ to $d-1$, those vectors are linearly independent, which concludes the proof.

3. Let I be an ideal of R and $s \in I$ be a non-zero element of I. Show that one can efficiently construct d elements s_i (for $1 \leq i \leq d$) in I such that the vectors $\Sigma(s_i)$ are linearly independent and have euclidean norm $\|\Sigma(s_i)\| = \|\Sigma(s)\|$. ($\star\star$)

A: Let us again take $s_i = s \cdot X^{i-1}$ for $i = 1$ to d. Those elements are in I since I is an ideal. Moreover, by definition of R, one can see that if $s = \sum_{j=0}^{d-1} x_j X^j$, then

$$s_{i+1} = s \cdot X^i = \sum_{j=0}^{d-1} x_j \cdot X^{i+j} = \sum_{k=i}^{d-1} x_{k-i} X^k - \sum_{k=0}^{i-1} x_{k+d-i} X^k,$$

(here, we use the fact that $X^\ell = -X^{\ell-d}$ in R for $d \leq \ell < 2d$). From this, one can see that $\Sigma(s_i)$ is obtain by permuting the coefficients of $\Sigma(s)$, and multiplying some of them by -1. This does not change the euclidean norm, i.e., we have $\|\Sigma(s_i)\| = \|\Sigma(s)\|$ for all i's.

4. Conclude that in an ideal lattice $\Sigma(I)$, finding one short vector $v \in \Sigma(I)$ is sufficient to construct a short basis B of $\Sigma(I)$ where all vectors b_i of B have euclidean norm at most $\sqrt{d} \cdot \|v\|$.
(Hint: you may want to use the result of question 5 from exercise 3)

A: This is done by combining the previous question with exercise 3. Using the previous question, we construct d linearly independent vectors in $\Sigma(I)$ with the same euclidean norm as v. Then, using exercise 3, we use this set of short linearly independent vectors to create a short basis of I, with a loss of a factor $\sqrt{d}$ on the size of the vectors.

Note: in this exercise, we used special properties of the ring R. In more generality, from one short vector $v \in \Sigma(I)$, one can construct a short basis with vectors of norm at most $\gamma_K \cdot \|v\|$ for some γ_K depending on the number fields K. For most number fields K used in cryptography, this quantity γ_K is small (and so the intuition that "one short vector in an ideal is sufficient to have a short basis" is true).