EXERCISES

The exercises 1 to 6 in this sheet are the most interesting ones and should be prioritized. Exercises 7 to 10 are more advanced, and are targeted for students who already have some knowledge about lattices, and want to learn more (exercise 10 is available only in the online version: `https://apelletm.pages.math.cnrs.fr/page-perso/documents/enseignement/CIMPA_school_Pondicherry/exercises.pdf`).

Exercises 1 to 4, and advanced exercises 7, 8, 10 can be done after the first lecture. Exercise 5 can be done after the second lecture (if we go fast enough), exercise 6 can be done after the third lecture, and exercise 9 can be done after the fourth lecture.

1 Lattice bases ($\star$)

The objective of this exercise is to prove a bunch of properties regarding bases of lattices. Throughout this exercise, the matrix B (or the matrices B_1, B_2) are invertible matrices in $\mathrm{GL}_n(\mathbb{R})$ for some dimension $n > 0$. Recall that we write $\mathcal{L}(B)$ for the lattice spanned by the columns of the matrix B.

1. Let $B_1, B_2 \in \mathrm{GL}_n(\mathbb{R})$. Show that $\mathcal{L}(B_1) = \mathcal{L}(B_2)$ if and only if $B_1 = B_2 \cdot U$ for some $U \in \mathbb{Z}^{n\times n}$ such that $\det(U) = \pm 1$. Such a matrix U is called unimodular. It is an invertible integer matrix whose inverse is also an integer matrix.

   A: Assume first that $\mathcal{L}(B_1) = \mathcal{L}(B_2)$. Then, every column of B_1 belongs to $\mathcal{L}(B_1) = \mathcal{L}(B_2)$. Hence, by definition of the lattice $\mathcal{L}(B_2)$ (integer linear combinations of the columns of B_2), we know that there exists an integer square matrix U_1 such that $B_1 = B_2 \cdot U_1$. Since B_1 and B_2 are both invertible, then U_1 is also invertible (over $\mathbb{R}$). Our objective is to show that U_1 is invertible over $\mathbb{Z}$ (i.e., it's inverse is also an integer matrix). By a similar argument, we know that there exist an invertible (over $\mathbb{R}$) integer matrix U_2 such that $B_2 = B_1 \cdot U_2$.

   Combining both equations, we obtain $B_1 = B_2 \cdot U_1 = B_1 \cdot U_2 \cdot U_1$. Since B_1 is invertible, we can simplify this into $I_n = U_2 \cdot U_1$. Since U_1 and U_2 are invertible over $\mathbb{R}$, their inverse is unique and we conclude that $U_1^{-1} = U_2$ is an integer matrix as desired.

   To conclude, observe that since U_1 and U_2 are integer matrices, then their determinant is also an integer. But we have $1 = \det(I_n) = \det(U_1 \cdot U_2) = \det(U_1) \cdot \det(U_2)$. Hence, the only possibility for $\det(U_1)$ is 1 or -1 (these are the only invertible elements in $\mathbb{Z}$).

   In the other direction, assume that $B_1 = B_2 \cdot U$ with U integer and $\det(U) = \pm 1$. Then, U is invertible over $\mathbb{R}$ and its inverse matrix U^{-1} has integer coefficients (recall that $U^{-1} = 1/\det(U) \cdot \mathrm{adj}(U)$ where the adjugate matrix $\mathrm{adj}(U)$ is integral since U is).

   Since U is integral, then by definition every column of $B_1 = B_2 \cdot U$ is in the lattice spanned by B_2. Hence we have $\mathcal{L}(B_1) \subseteq \mathcal{L}(B_2)$. Since U^{-1} is also integral, then every column of $B_2 = B_1 \cdot U^{-1}$ is in the lattice spanned by B_1, and we conclude that $\mathcal{L}(B_2) \subseteq \mathcal{L}(B_1)$.

2. Let B_1 and B_2 be two bases of the same lattice $\mathcal{L}$. Prove that $|\det(B_1)| = |\det(B_2)|$.
   This shows that the quantity $|\det(B)|$ does not depend on the choice of the basis B of $\mathcal{L}$, but only on the lattice $\mathcal{L}$. It is usually called the volume or the determinant of the lattice $\mathcal{L}$, and written $\mathrm{vol}(\mathcal{L})$ or $\det(\mathcal{L})$.

   A: We have seen in the previous questions that if $\mathcal{L}(B_1) = \mathcal{L}(B_2)$, then $B_1 = B_2 \cdot U$ for some matrix U with $\det(U) = \pm 1$. Taking the absolute value of the determinant of this equation proves that $|\det(B_1)| = |\det(B_2)|$.

3. Let $\mathcal{L}_1$ and $\mathcal{L}_2$ be two lattices of rank n. Show that if $\mathcal{L}_1 \subseteq \mathcal{L}_2$, then $\det(\mathcal{L}_1) = k \cdot \det(\mathcal{L}_2)$ for some integer $k > 0$. This integer k is called the index of $\mathcal{L}_1$ inside $\mathcal{L}_2$ and is written $[\mathcal{L}_2 : \mathcal{L}_1]$.

A: Let B_1 be a basis of $\mathcal{L}_1$ and B_2 be a basis of $\mathcal{L}_2$. Since $\mathcal{L}_1 \subseteq \mathcal{L}_2$, then every column of B_1 is in $\mathcal{L}(B_2)$, i.e., there is an integer matrix X such that $B_1 = B_2 \cdot X$. Taking the determinant, we have $\det(B_1) = \det(B_2) \cdot \det(X)$. Hence, $k = |\det(X)|$ and k is indeed an integer since X has integer coefficients (and k is non-zero since B_1 and B_2 are both invertible).

The determinant of a lattice is an important quantity, mostly useful in cryptography thanks to Minkowski's first theorem. This theorem states that in any lattice $\mathcal{L}$ of dimension n, there exists a non-zero vector $v \in \mathcal{L}$ such that $\|v\| \leq \sqrt{n} \cdot \det(\mathcal{L})^{1/n}$.

4. Show that the upper bound in Minkowski's first theorem can be quite loose for some lattices: construct a lattice with $\det(\mathcal{L}) = 1$ and which contains a non-zero vector v whose euclidean norm is arbitrarily close to 0.

A: Take $\varepsilon > 0$ and define $\mathcal{L}$ to be the lattice with basis $b_1 = (\varepsilon, 0)^T$ and $b_2 = (0, \varepsilon^{-1})^T$. Then $\det(\mathcal{L}) = 1$ but $\mathcal{L}$ contains the vector b_1 whose norm can be arbitrarily close to 0.

The objective of the next questions is to observe that when dealing with lattices, a maximal set of independent vectors is not always a basis, and a minimal set of generating vectors is also not always a basis (which differs from what we are used to in vector spaces).

5. Exhibit a family of n linearly independent vectors in $\mathbb{Z}^n$ which do not form a $\mathbb{Z}$-basis of $\mathbb{Z}^n$.

A: One example is the family $b_i = (0, \ldots, 0, 2, 0, \ldots, 0)$ with a 2 in i-th position, for $i = 1$ to n. Those vectors are linearly independent but they generate the lattice $(2\mathbb{Z})^n$, which is included strictly in $\mathbb{Z}^n$. Note that one cannot add a vector to this family of vectors and still have independent vectors (because independence is defined over $\mathbb{R}$, where things work as expected: the maximal size of an independent set of vectors in $\mathbb{R}^n$ is n).

6. Exhibit a family of $n + 1$ vectors generating $\mathbb{Z}^n$ such that it is not possible to remove any vector from this set to obtain a $\mathbb{Z}$-basis of $\mathbb{Z}^n$.

A: Take $b_0 = (2, 0, \ldots, 0)$, $b_1 = (3, 0, \ldots, 0)$ and $b_i = (0, \ldots, 0, 1, 0, \ldots, 0)$ with a i at the i-th position for $i = 2$ to n. Then $(b_i)_{0 \leq i \leq n}$ generates $\mathbb{Z}^n$. This is because 2 and 3 are coprime, hence one can find an integer linear combination of b_1 and b_2 with a 1 in its first coordinate (just take $b_1 - b_0 = (1, 0, \ldots, 0)$).

However, one can check that removing b_0 or b_1 from the list of generator does not generate $\mathbb{Z}^n$ anymore: the first coordinate will always be a multiple of 2 or 3. Similarly, we cannot remove one of the b_i for $i \geq 2$ since the i-th coordinate would always be 0.

7. Compute a basis for the lattice generated by $c_1 = (2\pi, 4)^T$, $c_2 = (0, 3)^T$ and $c_3 = (4\pi, 4)^T$. Same question for $c_1 = (1, 0)^T$, $c_2 = (1, 1)^T$ and $c_3 = (1, \pi)^T$. (⋆⋆)
(Hint: the question might be lying to you. In this case, show what is wrong in the question. :)).

A: A basis for the first lattice is given by $b_1 = (2\pi, 0)^T$ and $b_2 = (0, 1)^T$. A way to check that this is indeed a basis of the lattice generated by c_1, c_2 and c_3 is to check that each of the b_i is in the $\mathbb{Z}$-span of the c_i-s (note: $b_2 = 2c_1 - c_3 - c_2$ and $b_1 = c_1 - 4b_2$) and that reciprocally each of the c_i is in the $\mathbb{Z}$-span of the b_i's. This shows that the b_i an the c_i generates the same lattice. Then observe that the b_i are 2 linearly independent vectors in $\mathbb{R}^2$ hence they form a basis of their lattice.

For the second example, it turns out that the $\mathbb{Z}$-span of c_1, c_2 and c_3 is not a lattice. A way to see this is that a lattice must be discrete (see the alternative definition in Section **??**). But the $\mathbb{Z}$-span of c_1, c_2 and c_3 is not discrete. Indeed, we have $(0, 1)^T$ and $(0, \pi)^T$ in the $\mathbb{Z}$-span. Since π is not a rational number, we can create a vector $(0, \varepsilon)^T$ with ε as small as we want by taking integer linear combinations of those two vectors. This shows that the $\mathbb{Z}$-span of the c_i contains an accumulation point at 0, and so it is not a lattice.

2 HNF basis (⋆⋆)

In this exercise, we will see how to compute a special basis of a lattice $\mathcal{L}$, called the HNF basis of $\mathcal{L}$. The main advantage of this basis is that it can be computed in polynomial time from any *basis of $\mathcal{L}$, hence, it is a "worst possible" basis: revealing this basis does not leak more information on $\mathcal{L}$ than what any other basis would leak.*

The algorithm to compute the HNF basis is very similar to the way one would use Gaussian elimination to compute the echelon form of matrices over a field. The main difference is that since we are only allowed to perform integer linear combinations over the vectors of our basis, we cannot multiply by the inverse of a coefficient, in order to annihilate the other coefficients on the same row.

1. Let's review Gaussian elimination a little. Run Gaussian elimination (over $\mathbb{R}$) on the columns of the matrix $M = \begin{pmatrix} 2 & 3 \\ 3 & 4 \end{pmatrix}$ in order to obtain a triangular matrix of the form $\begin{pmatrix} * & 0 \\ * & * \end{pmatrix}$. (Here, running Gaussian elimination on the columns means that you are only allowed to perform operations on the columns of the matrix. Said differently, you can only multiply M by invertible matrices on the right).

   A: In order to obtain a 0 on the top-right part of the matrix, we perform the operation $C_2 \leftarrow C_2 - 3/2 \cdot C_1$ (where C_1 and C_2 are the columns of the matrix M). This corresponds to multiplication on the right by the matrix $\begin{pmatrix} 1 & -3/2 \\ 0 & 1 \end{pmatrix}$. We then obtain the matrix $\begin{pmatrix} 2 & 0 \\ 3 & -1/2 \end{pmatrix}$, which has the desired shape.

2. In the previous question, the operations we performed on the columns were not integer. We now want to focus on integer operations on the columns of M. Show that there exists an integer matrix U with determinant 1 such that $M \cdot U = \begin{pmatrix} 1 & * \\ * & * \end{pmatrix}$.

   A: The matrix $U = \begin{pmatrix} -1 & 0 \\ 1 & -1 \end{pmatrix}$ has integral coefficient, determinant 1 and satisfies $M \cdot U = \begin{pmatrix} 2 & 3 \\ 3 & 4 \end{pmatrix} \cdot \begin{pmatrix} -1 & 0 \\ 1 & -1 \end{pmatrix} = \begin{pmatrix} 1 & -3 \\ 1 & -4 \end{pmatrix}$ as desired.

3. More generally, show that for any matrix $M = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$ there is a unimodular matrix U such that $M \cdot U = \begin{pmatrix} \gcd(a,b) & * \\ * & * \end{pmatrix}$. (⋆⋆)

   A: We know by Bézout's identity that there exists u, v integers such that $au + bv = \gcd(a, b)$. Moreover, this equality also shows that such u and v must be coprime, since $\gcd(a, b)$ already divides a and b. Hence, applying Bézout's identity once more to u and v, we have x and y such that $ux + vy = 1$. Take the matrix $U = \begin{pmatrix} u & -y \\ v & x \end{pmatrix}$. The first column of this matrix is constructed such that the top-left coefficient of $M \cdot U$ is equal to $au + bv = \gcd(a, b)$. The second column of the matrix is added so that the matrix U has determinant 1 (so that it is invertible over $\mathbb{Z}$). This is ensured by the second Bézout's identity, which shows that $\det(U) = ux + yv = 1$, i.e., U is unimodular.

4. Using the previous question, show that for any matrix $M = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$ there is a unimodular matrix U such that $M \cdot U = \begin{pmatrix} \gcd(a,b) & 0 \\ * & * \end{pmatrix}$.

A: Once we have applied the unimodular matrix U from the previous question, we obtain a basis of the form $\begin{pmatrix} \gcd(a,b) & z \\ * & * \end{pmatrix}$. Moreover, we know that z must be a multiple of $\gcd(a,b)$, since it is an integer linear combination of a and b (all top coefficient of vectors in $\mathcal{L}(M)$ must be integer linear combinations of a and b). Hence, from now on, we can use regular Gaussian elimination and perform $C_2 \leftarrow C_2 - z/\gcd(a,b)C_1$ to annihilate the top-right coefficient (where C_1 and C_2 are the columns of the matrix $M \cdot U$). This operation is obtained by multiplying on the right by the matrix $U' = \begin{pmatrix} 1 & -z/\gcd(a,b) \\ 0 & 1 \end{pmatrix}$ which is integer and unimodular as desired.

5. Compute a matrix U as in the previous question for $M = \begin{pmatrix} 9 & 2 \\ 3 & 1 \end{pmatrix}$.

   A: $U = \begin{pmatrix} -1 & -2 \\ 5 & 9 \end{pmatrix}$ which gives $M \cdot U = \begin{pmatrix} 1 & 0 \\ 2 & 3 \end{pmatrix}$

6. Let $M_1 = \begin{pmatrix} 2 & 1 & 0 \\ 8 & 1 & 4 \\ 0 & 1 & 7 \end{pmatrix}$. Generalize the algorithm from the previous questions to compute a matrix M_2 such that $M_2 = M_1 \cdot U$ for some unimodular matrix U and M_2 is of the form $M_2 = \begin{pmatrix} * & 0 & 0 \\ * & * & 0 \\ * & * & * \end{pmatrix}$.

   A: $U = \begin{pmatrix} 0 & -1 & -2 \\ 1 & 2 & 4 \\ 0 & 2 & 3 \end{pmatrix}$ and $M_2 = M_1 \cdot U = \begin{pmatrix} 1 & 0 & 0 \\ 1 & 2 & 0 \\ 1 & 16 & 25 \end{pmatrix}$

7. Let $\mathcal{L}$ be a lattice of dimension n. Show that there is a unique basis B of $\mathcal{L}$ such that $b_{i,j} = 0$ when $j > i$, $b_{i,i} > 0$ and $0 \leq b_{i,j} < b_{i,i}$ for $j < i$. This is the basis which is called the Hermite normal form (HNF) basis of $\mathcal{L}$. (⋆⋆)

   A: First, observe that the algorithm that we described in the previous question provides an algorithmic proof that such a basis exists (the condition that $b_{i,j} \in [0, b_{i,i})$ for $j < i$ is ensured by reducing the non-diagonal coefficients modulo the diagonal coefficients, from top to bottom).

   Let us now prove that such a basis is unique. Assume for a contradiction that there exists two such bases B and C, with columns b_j and c_j. Let j_0 be maximal such that $b_{j_0} \neq c_{j_0}$. Since B and C span the same lattice, then b_{j_0} is an integer linear combination of the vectors $(c_j)_{1 \leq j \leq n}$. Moreover, because of the special shape of C and since the top coefficients of b_{j_0} are 0, then is must be that b_{j_0} is a combination of the columns c_j for $j \geq j_0$. This implies that the diagonal coefficient b_{j_0,j_0} is an integer multiple of c_{j_0,j_0}. But a similar argument shows that c_{j_0,j_0} is an integer multiple of b_{j_0,j_0}, hence we conclude that $|b_{j_0,j_0}| = |c_{j_0,j_0}|$. Since both are positive by assumption, we conclude that $b_{j_0,j_0} = c_{j_0,j_0}$.

   From this, we know that $b_{j_0} = c_{j_0} + \sum_{j>j_0} a_j \cdot c_j$ for some integers a_j's. However, we know that $c_j = b_j$ for any $j > j_0$ by choice of j_0, which means that the diagonal coefficients $b_{j,j}$ and $c_{j,j}$ are the same for $j > j_0$. We also know that the bottom coefficients of both b_{j_0} and c_{j_0} are reduced modulo those diagonal coefficients $c_{j,j} = b_{j,j}$. Hence, a recursive argument shows that a_j must be equal to 0 for all $j > j_0$, and we conclude that $b_{j_0} = c_{j_0}$, which is a contradiction.

3 LWE and SIS lattices (⋆⋆)

Let q ,$m \geq r > 0$ be integers and $A \in \mathbb{Z}_q^{m \times r}$. Recall that the SIS lattice associated to A is defined by $\Lambda^\perp(A) := \{x \in \mathbb{Z}^m \,|\, x^T \cdot A = 0 \bmod q\}$. Recall similarly that the LWE lattice associated to A is $\Lambda(A) := \{x \in \mathbb{Z}^m \,|\, \exists s \in \mathbb{Z}^n$ s.t. $As = x \bmod q\}$.

1. Show that $\Lambda(A)$ is generated by the columns of A and the m vectors $q \cdot e_i$ (with $1 \leq i \leq m$), where e_i is the vector with a 1 at the i-th position and 0's everywhere else.

   A: First, one can check that $\Lambda(A)$ indeed contains the column vectors of A (take $s = (0, \ldots, 0, 1, 0, \ldots, 0)$ in the definition of $\Lambda(A)$) and the m vectors $q \cdot e_i$ (take $s = 0$).

   Let us then show the reverse inclusion. Let $x \in \Lambda(A)$. By definition, there must exist a vector $s \in \mathbb{Z}^r$ and $z \in \mathbb{Z}^m$ such that $x = A \cdot s + q \cdot z$. This shows that x is an integer linear combination of the columns of A and the $q \cdot e_i$ vectors. Hence, those vectors indeed generate the lattice $\Lambda(A)$.

2. Assume that q is prime. Using the previous question, exhibit a set of generating vectors for the lattice $\Lambda^\perp(A)$. (Hint: you might want to show that $\Lambda^\perp(A) = \Lambda(B)$ for some well chosen matrix B).

   A: Let $B \in \mathbb{Z}^{m \times k}$ be a basis (in columns) of the left kernel of A modulo q, i.e., $B^T \cdot A = 0 \bmod q$ (here, we use the fact that q is prime so that $\mathbb{Z}_q$ is a field and the kernel of A is a vector space). We know that $k \geq m - r$, but it could be bigger if the rank of A modulo q is $< r$. We have that $x \in \Lambda^\perp(A)$ if and only if $x^T \cdot A = 0 \bmod q$, which is equivalent to x belongs to the span of the columns of B modulo q, i.e., $x \in \Lambda(B)$. Using the previous question, we conclude that the column vectors of B together with the $q \cdot e_i$ vectors form a generating set of $\Lambda^\perp(A)$.

3. Assume again that q is prime. Assume also that the rank of A modulo q is r (i.e., the r column vectors of A are linearly independent modulo q). Show that up to permuting the rows of A (i.e., permuting the coefficients of the vectors in $\Lambda(A)$), there exists a basis of $\Lambda(A)$ of the form $\begin{pmatrix} I_r & 0_{n \times (m-r)} \\ A' & q \cdot I_{m-r} \end{pmatrix}$, for some integer matrix $A' \in \mathbb{Z}^{(m-r) \times r}$. $(\star\star)$

   Similarly, show that up to permuting the rows of A, there exists a basis of $\Lambda^\perp(A)$ of the form $\begin{pmatrix} I_{m-r} & 0_{(m-r) \times r} \\ B' & q \cdot I_r \end{pmatrix}$, for some integer matrix $B' \in \mathbb{Z}^{r \times (m-r)}$.

   A: First, observe that by definition of $\Lambda(A)$, the lattice only depends on the span over $\mathbb{Z}_q$ of the columns of A, and not the actual choice of the basis A. Also, since the rank of the columns of A is r, then up to permuting the rows of A, we can assume that $A = \begin{pmatrix} A_1 \\ A_2 \end{pmatrix}$ with $A_1 \in \mathbb{Z}^{r \times r}$ invertible modulo q.

   Hence, we have $\Lambda(A) = \Lambda(A \cdot A_1^{-1})$, where $A \cdot A_1^{-1} = \begin{pmatrix} I_r \\ A' \end{pmatrix}$, with $A' = A_2 \cdot A_1^{-1}$.

   By a previous question, we know that the columns of $\tilde{A} := \begin{pmatrix} I_r \\ A' \end{pmatrix}$ together with the $q \cdot e_i$ vectors generate the lattice $\Lambda(A)$.

   Observe now that because of the special shape of $\tilde{A}$, the first r vectors $q \cdot e_i$ are already in the span of the columns of $\tilde{A}$ and the other $q \cdot e_j$ vectors for $j > r$ (for $i \leq r$, the vector $q \cdot e_i$ can be obtained by multiplying the i-th column of $\tilde{A}$ by q and annihilating the bottom $m - r$ coordinates using the $q \cdot e_j$ with $j > r$ since those coordinates will be integer multiples of q).

   Hence, the r column vectors of A together with the $(m - r)$ vectors qe_j for $j > r$ generate the lattice $\Lambda(A)$. Since those are exactly m vectors, they form a basis of the lattice, with the desired shape.

   Regarding $\Lambda^\perp(A)$, we have already seen in a previous question that this lattice is equal to $\Lambda(B)$ where B forms a basis of the kernel of A. Since A has rank r modulo q, then we know that B has dimension $m \times (m - r)$ and rank $m - r$ modulo q. Applying what we have done above to the matrix B solves the second part of the question.

4. Assuming that q is prime and that A has rank n modulo q, show that the SIS lattice $\Lambda^\perp(A)$ contains a non-zero vector of norm $\leq \sqrt{m} \cdot q^{r/m}$ and that the LWE lattice $\Lambda(A)$ contains a non-zero vector of norm $\leq \sqrt{m} \cdot q^{1-r/m}$.

A: From the previous question, we know that $\det(\Lambda(A)) = q^{m-r}$ and $\det(\Lambda^\perp(A)) \leq q^r$ (permuting the coefficients of the vectors does not change the volume of the lattices). The shortness of the vectors then follows from Minkowski's first theorem.

4 Solving the closest vector problem ($\star$)

Babai's round-off algorithm solves the approximate closest vector problem as follows. Given as input a basis $(b_i)_{1\leq i\leq n}$ of the lattice $\mathcal{L}$ (of dimension n) and a target t, the algorithm writes $t = \sum_{i=1}^n t_i b_i$ with $t_i \in \mathbb{R}$ and output the vector $s = \sum_i \lceil t_i \rfloor b_i$.

1. Show that Babai's round-off algorithm finds a point $s \in \mathcal{L}$ such that $\|t - s\| \leq 1/2 \cdot n \cdot \max_i \|b_i\|$.

A: Since the $\lceil t_i \rfloor$'s are integers, then s belongs indeed to the lattice $\mathcal{L}$.

Let us now compute the distance to t. For $x \in \mathbb{R}$, we write $\{x\} = x - \lceil x \rfloor$ the fractional part of x. It belongs to $[-1/2, 1/2]$.

$$\begin{aligned}\|s - t\| &= \|\sum_i \{t_i\} \cdot b_i\| \\ &\leq \sum_i |\{t_i\}| \cdot \|b_i\| \\ &\leq 1/2 \cdot n \cdot \max_i \|b_i\|.\end{aligned}$$

5 LWE is a BDD problem ($\star\star$)

In this exercise, we will fix an LWE instance (A, b), with a prime modulus q, with $A \in (\mathbb{Z}/q\mathbb{Z})^{m\times n}$ (where $m \geq n > 0$ are integers), and with $b = A \cdot s + e \bmod q$ for some secret $s \in (\mathbb{Z}/q\mathbb{Z})^n$ and $e \in \mathbb{Z}^m$ satisfying $\|e\| \leq B$ (for some bound $B > 0$). We will show that, under some conditions on the parameters B, q, m and n, it is possible to recover the secret s by solving a BDD instance in the lattice $L := \Lambda(A)$ from Exercise 3 (hence the name "LWE lattice").[1]

1. Let $t \in \mathbb{Z}^m$ be any lift of b in $\mathbb{Z}^m$ (i.e., $t \bmod q = b$). Show that there exists a vector $v_0 \in L$ such that $\|v_0 - t\|_2 \leq B$. Show also that if one recovers v_0, then one can recover the secret s of the LWE instance.

A: Let $v_0 = t - e \in \mathbb{Z}^m$. We have that $v_0 \equiv b - e \equiv As \bmod q$, hence, $v_0 \in \Lambda(A)$ by definition of $\Lambda(A)$. Moreover, $\|v_0 - t\|_2 = \|e\|_2 \leq B$ by assumption on e. If one recovers $v_0 = As \bmod q$, then using Gauss pivoting (in the field $\mathbb{Z}/q\mathbb{Z}$), one can recover s from v_0 and A in polynomial time, and hence, solve the LWE instance. Note that if the rank of A is $< n$, s may not be unique, in which case the algorithm will recover one possible s. This is not really an issue because any such s would be a solution to the LWE instance, and so in practice it would be sufficient to recover one such s to break cryptosystems. Moreover, we will see with the next questions that s is unique with overwhelming probability over the random choice of A.

To prove that s can be recovered by solving a BDD instance in L, is then "only" remains to prove that $B \leq 1/\gamma \cdot \lambda_1(L)$ for some $\gamma > 2$ (this is actually the hardest part of the proof): this will prove that t is a γ-BDD instance in L, with closest vector v_0, and so solving BDD in L with target t will recover v_0 and hence we can obtain the secret s.

In Exercise 3, we have shown that $\lambda_1(L) \leq \sqrt{m} \cdot q^{1-n/m}$ (when A has rank n). The next questions will be devoted to the proof that this bound is close to optimal (with overwhelming probability over the random choice of A). Recall that we asked that q is prime.

[1] More precisely, this will be possible with overwhelming probability over the random choice of A.

2. Let $\rho \in (0, q)$ be some bound, and write $\mathcal{B}_\rho$ the m-dimensional euclidean ball centered in 0 and of radius ρ. Show that

$$\Pr_A\big(\lambda_1(\Lambda(A)) \leq \rho\big) \leq \sum_{\substack{y \in \mathcal{B}_\rho \cap \mathbb{Z}^m \\ x \in (\mathbb{Z}/q\mathbb{Z})^n \setminus \{0\}}} \Pr_A\big(Ax = y \bmod q\big),$$

where A is sampled uniformly at random in $(\mathbb{Z}/q\mathbb{Z})^{m \times n}$.
(Hint: you may want to use a union bound)

A: By definition of $\Lambda(A)$, we know that $\lambda_1(\Lambda(A)) \leq \rho$ if and only if there exists $y \in \mathbb{Z}^m$ with $y \neq 0$ and $\|y\| \leq \rho$ and $x \in (\mathbb{Z}/q\mathbb{Z})^n$ such that $y = Ax \bmod q$. Since $\rho < q$ and $0 < \|y\| \leq \rho$, it must be that $y \neq 0 \bmod q$. Hence, it must also be that $x \in (\mathbb{Z}/q\mathbb{Z})^n \setminus \{0\}$. This implies that

$$\begin{aligned}\Pr_A\big(\lambda_1(\Lambda(A)) \leq \rho\big) &= \Pr_A\Big(\exists y \in (\mathcal{B}_\rho \cap \mathbb{Z}^m) \setminus \{0\},\, \exists x \in (\mathbb{Z}/q\mathbb{Z})^n \setminus \{0\} \,|\, y = Ax \bmod q\Big) \\ &\leq \sum_{\substack{y \in (\mathcal{B}_\rho \cap \mathbb{Z}^m) \setminus \{0\} \\ x \in (\mathbb{Z}/q\mathbb{Z})^n \setminus \{0\}}} \Pr_A\big(y = Ax \bmod q\big),\end{aligned}$$

where the last inequality follows from the union bound.

3. Show that for any $x \in (\mathbb{Z}/q\mathbb{Z})^n \setminus \{0\}$ and $y \in \mathbb{Z}^m$ fixed, it holds that $\Pr_A\big(Ax = y \bmod q\big) = q^{-m}$.
(Hint: this is where you use that q is prime)

A: Let us write a_i the rows of the matrix A and y_i the coefficients of y. Since A is sampled uniformly at random, its rows are also uniform in $(\mathbb{Z}/q\mathbb{Z})^n$ and independent. Hence, it holds that

$$\Pr_A\big(Ax = y \bmod q\big) = \prod_{i=1}^{m} \Pr_A\big(\langle a_i, x\rangle = y_i \bmod q\big).$$

Since $x \neq 0$, there must be some index $\ell \in [1, n]$ such that $x_\ell \neq 0$. Since q is prime, then x_ℓ must be invertible modulo q. Hence, if we write $a_{i,j}$ the coefficients of the vector a_i, we see that

$$\langle a_i, x\rangle = y_i \iff a_{i,\ell} x_\ell = y_i - \sum_{j \neq \ell} a_{i,j} x_j \iff a_{i,\ell} = x_\ell^{-1} \cdot (y_i - \sum_{j \neq \ell} a_{i,j} x_j).$$

But $a_{i,\ell}$ is uniform in $\mathbb{Z}/q\mathbb{Z}$ and independent from the other $a_{i,j}$'s, hence we conclude that $\Pr_A\big(a_{i,\ell} = x_\ell^{-1}(y_i - \sum_{j \neq \ell} a_{i,j} x_j)\big) = 1/q$. Taking the product over all i's leads to $\Pr_A\big(Ax = y \bmod q\big) = (1/q)^m$ as desired.

4. Conclude that

$$\Pr_A\big(\lambda_1(\Lambda(A)) \leq \rho\big) \leq |\mathcal{B}_\rho \cap \mathbb{Z}^m| \cdot q^{n-m}.$$

A: Combining the two previous questions we have

$$\begin{aligned}\Pr_A\big(\lambda_1(\Lambda(A)) \leq \rho\big) &\leq \sum_{\substack{y \in \mathcal{B}_\rho \cap \mathbb{Z}^m \\ x \in (\mathbb{Z}/q\mathbb{Z})^n \setminus \{0\}}} \Pr_A\big(Ax = y \bmod q\big) \\ &= \sum_{\substack{y \in \mathcal{B}_\rho \cap \mathbb{Z}^m \\ x \in (\mathbb{Z}/q\mathbb{Z})^n \setminus \{0\}}} q^{-m} \\ &= |\mathcal{B}_\rho \cap \mathbb{Z}^m| \cdot q^n \cdot q^{-m},\end{aligned}$$

as desired.

5. Show that $|\mathcal{B}_\rho \cap \mathbb{Z}^m| \leq (2\rho+1)^m$.
(Hint: the ball $\mathcal{B}_\rho$ is contained in $\mathcal{B}_\rho^\infty := \{z \in \mathbb{R}^m \mid \|z\|_\infty \leq \rho\}$).

A: Note that for any $x \in \mathbb{R}^m$, we have that $\|x\|_\infty \leq \|x\|_2$. Hence, if $x \in \mathcal{B}_\rho$, we have that $\|x\|_\infty \leq \|x\|_2 \leq \rho$ and so $x \in \mathcal{B}_\rho^\infty$: the hint is correct. This implies that $|\mathcal{B}_\rho \cap \mathbb{Z}^m| \leq |\mathcal{B}_\rho^\infty \cap \mathbb{Z}^m|$. Counting the number of points in $\mathcal{B}_\rho^\infty \cap \mathbb{Z}^m$ is easier: we want the $x \in \mathbb{Z}^m$ such that $|x_i| \leq \rho$ for all coordinates x_i of x. We have at most $(2\rho+1)$ possible choice for each x_i, i.e., at most $(2\rho+1)^m$ choices for x.

6. Assume that m, n and q satisfy $q^{1-n/m} \geq 4$ (e.g., if $m \geq 2n$ and $q \geq 16$). Conclude that

$$\Pr_A\left(\lambda_1(\Lambda(A)) \leq \frac{q^{1-n/m}}{4} - 1\right) \leq 2^{-m}.$$

A: Combining the previous two questions, we see that

$$\Pr_A\left(\lambda_1(\Lambda(A)) \leq \rho\right) \leq (2\rho+1)^m \cdot q^{n-m}.$$

Instantiating this equation with $\rho = \frac{q^{1-n/m}}{4} - 1$, we see that the upper bound becomes

$$\begin{aligned}(2\rho+1)^m \cdot q^{n-m} &= \left(\frac{q^{1-n/m}}{2}\right)^m \cdot q^{n-m} \\ &= \frac{q^{m-n}}{2^m} \cdot q^{n-m} = 2^{-m}.\end{aligned}$$

Remark: *there is a gap of the order of $O(\sqrt{m})$ between the upper bound on $\lambda_1(L)$ from exercise 3 and the lower bound we just proved. This gap can be reduced by being more careful when upper bounding the quantity $|\mathcal{B}_\rho \cap \mathbb{Z}^m|$.*

7. Conclude that if $q^{1-n/m} \geq 16 \cdot B$ (e.g., if $q \geq 2^{16}$, $B = \sqrt{q}$ and $m \geq 4n$), then with probability at least $1 - 2^{-m}$ over the choice of A, it is possible to solve the LWE instance (A, b) by solving a γ-BDD instance in the lattice $L := \Lambda(A)$ with $\gamma = \frac{q^{1-n/m}}{8B}$.

A: Let $q^{1-n/m} \geq 16 \cdot B$. The previous question shows that, with probability at least $1 - 2^{-m}$ over the choice of A, it holds that $\lambda_1(\Lambda(A)) \geq \frac{q^{1-n/m}}{4} - 1 \geq 4B - 1$. Moreover, we can assume without loss of generality that $B \geq 1$ (otherwise, $\|e\| \leq B < 1$ and so we must have $e = 0$ since it has integer coordinates). Hence, we have $4B - 1 > 2B$. We have seen in question 1 that one could solve the LWE instance by recovering the lattice point $v_0 \in \Lambda(A)$ which is at distance $\leq B$ from the target t. We have just seen that $B < \lambda_1(\Lambda(A))/2$, hence this is a BDD instance (and v_0 is unique). The approximation factor γ is this BDD instance is

$$\begin{aligned}\gamma &= \|t - v_0\|_2^{-1} \cdot \lambda_1(\Lambda(A)) \\ &\geq B^{-1} \cdot \left(\frac{q^{1-n/m}}{4} - 1\right) \\ &\geq B^{-1} \cdot \frac{q^{1-n/m}}{8} = \frac{q^{1-n/m}}{8B}.\end{aligned}$$

6 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.

7 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 8.

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.

8 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 7.

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 7)

   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 7 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 7.)

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 7 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.

9 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

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

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 $\mathbb{R}^d$. (⋆⋆)

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 $\|\Sigma(s_i)\| = \|\Sigma(s)\|$. (⋆⋆)

   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 8)

   A: This is done by combining the previous question with exercise 8. Using the previous question, we construct d linearly independent vectors in $\Sigma(I)$ with the same euclidean norm as v. Then, using exercise 8, 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).

10 Lagrange-Gauss algorithm (⋆ ⋆ ⋆)

Recall the Lagrange-Gauss algorithm: given as input a basis (b_1, b_2) of a lattice in $\mathbb{R}^2$, the algorithm finds $x \in \mathbb{Z}$ that minimizes $\|b_2 - xb_1\|$ and replaces b_2 by $b_2 - xb_1$ (finding x efficiently is done by computing the QR factorization of the basis B, this step is not important for this exercise). The algorithm then switches b_1 and b_2 and starts again. The algorithm stops when no progress is made for two consecutive iterations (which means that we cannot reduce b_1 by b_2 nor b_2 by b_1 anymore).

1. Let b_1 and b_2 be two non-zero vectors in $\mathbb{R}^2$. Show that if $\|b_1\| \leq \|b_1 + b_2\|$, then for any $\alpha \in (1, +\infty)$ it holds that $\|b_1 + b_2\| \leq \|b_1 + \alpha b_2\|$. (⋆⋆)

   A: Let us consider the function

$$\begin{aligned} f : \mathbb{R} &\to \mathbb{R}_+ \\ \alpha &\mapsto \|b_1 + \alpha b_2\| \end{aligned}$$

   A drawing shows that this is a convex function which has a unique minimum at α_0, is decreasing on $(-\infty, \alpha_0]$ and increasing on $[\alpha_0, +\infty)$. Since $\|b_1\| \leq \|b_1 + b_2\|$ (i.e., $f(0) \leq f(1)$) by assumption, then it must be that $\alpha_0 \leq 1$. From this we conclude that f is increasing on $[1, +\infty)$ which implies that $\|b_1 + b_2\| \leq \|b_1 + \alpha b_2\|$ for any $\alpha \geq 1$.

2. Show that if the Lagrange-Gauss algorithm terminates, then either b_1 or b_2 is a shortest non-zero vector of $\mathcal{L}$. (Hint: you may want to consider a shortest non-zero vector $s = x_1 b_1 + x_2 b_2$ and write it as $s = x_1 \cdot (b_1 + \alpha b_2)$ with $\alpha = x_2/x_1$ if $x_1 \neq 0$.) $(\star\star\star)$

A: Without loss of generality, assume that $\|b_1\| \leq \|b_2\|$. Let's use the hint and take $s = x_1 b_1 + x_2 b_2$ be a shortest non-zero vector in $\mathcal{L}$ (with x_1 and x_2 integers). Without loss of generality, we can assume that $x_1, x_2 \geq 0$ (otherwise we can multiply b_1 and/or b_2 by -1, which does not change their size nor the fact that the algorithm cannot reduce them anymore).

If $x_1 = 0$, then we must have $x_2 \geq 1$ (since $x_2 \neq 0$ is a non-negative integer). Then we have $\|b_1\| \leq \|b_2\| \leq \|x_2 b_2\| = \|s\|$, from which we conclude that b_1 is a shortest non-zero vector of $\mathcal{L}$.

Similarly, if $x_2 = 0$, then $\|b_1\| \leq \|x_1 b_1\| = \|s\|$ and so b_1 is a shortest non-zero vector.

Let us now assume that x_1 and x_2 are both non-zero. Assume that $x_1 \geq x_2$, then $s = x_2 \cdot (\alpha b_1 + b_2)$, with $\alpha = x_1/x_2 \geq 1$. Since the algorithm terminated, we know that b_2 cannot be reduced anymore be adding to it multiples of b_1, which implies in particular that $\|b_2\| \leq \|b_2 + b_1\|$. From the previous question, we conclude that $\|b_2 + b_1\| \leq \|b_2 + \alpha b_1\| \leq \|s\|$ (since $x_2 \geq 1$). We finally conclude that $\|b_1\| \leq \|b_2\| \leq \|b_2 + b_1\| \leq \|s\|$ as desired.

If $x_1 \leq x_2$, the situation is very similar. We have

$$\begin{aligned} \|b_1\| &\leq \|b_1 + b_2\| \\ &\leq |x_1| \cdot \|b_1 + b_2\| \\ &\leq |x_1| \cdot \|b_1 + (x_2/x_1) \cdot b_2\| \\ &= \|s\|. \end{aligned}$$