

Short Integer Solutions (SIS): The SIS problem is defined with respect to lattice parameters n, m, q and a norm bound β . The $SIS_{n,m,q,\beta}$ problem says that for $A \leftarrow \mathbb{Z}_q^{n \times m}$, no efficient adversary can find a non-zero vector $x \in \mathbb{Z}^m$ where $Ax = 0 \in \mathbb{Z}_q^n$ and $\|x\| \leq \beta$ [We will use l_∞ norm: $\|v\| = \|v\|_\infty = \max_i |v_i|$]

In lattice-based cryptography, the lattice dimension n will be the primary security parameter.

Notes: - The norm bound β should satisfy $\beta \leq q$.

- We need to choose m, β to be large enough so that a solution does exist.

↳ When $m = \Omega(n \log q)$ and $\beta \geq 1$, a solution always exists. In particular, when $m \geq \lceil n \log q \rceil$, there always exists $x \in \{-1, 0, 1\}^m$ such that $Ax = 0$: recall that we are using the l_∞ norm (unless otherwise noted)

- There are $2^m \geq 2^{n \log q} = q^n$ vectors $y \in \{0, 1\}^m$
 - Since $Ay \in \mathbb{Z}_q^n$, there are at most q^n possible outputs of Ay
 - Thus, if we set $x = y_1 - y_2 \in \{-1, 0, 1\}^m$, then $Ax = A(y_1 - y_2) = Ay_1 - Ay_2 = 0 \in \mathbb{Z}_q^n$
- } By a counting argument, there exist $y_1 \neq y_2 \in \{0, 1\}^m$ such that $Ay_1 = Ay_2$

SIS as a lattice problem: given $A \leftarrow \mathbb{Z}_q^{n \times m}$, find non-zero $x \in \mathbb{Z}_q^m$ such that $Ax = 0 \in \mathbb{Z}_q^n$ and $\|x\| \leq \beta$.

↳ Can be viewed as an average-case version of finding short vectors in a " q -ary" lattice:

$$L^\perp(A) = \{z \in \mathbb{Z}^m : Az = 0 \pmod{q}\}$$

Notice that by construction, $q\mathbb{Z}^m \subseteq L^\perp(A)$

↳ " q -ary" lattice (e.g., vectors where all entries are integer multiples of q)

Hardness of SIS: Ajtai first showed (in 1996) that average-case hardness of SIS can be based on worst-case hardness of certain lattice problems \Rightarrow long sequence of works understanding and improving the worst-case to average-case reductions

Typical statement: Let n be the lattice dimension. For any $m = \text{poly}(n)$, norm bound $\beta > 0$, and sufficiently large $q \geq \beta \cdot \text{poly}(n)$, then, the $SIS_{n,m,q,\beta}$ problem is at least as hard as solving GapSVP_γ on an arbitrary n -dimensional lattice for $\gamma = \beta \cdot \text{poly}(n)$.

↳ i.e., solving SIS is as hard as approximating GapSVP in the worst case!

We can use SIS to construct a collision-resistant hash function (CRHF).

Definition. A keyed hash family $H: \mathcal{K} \times \mathcal{X} \rightarrow \mathcal{Y}$ is collision-resistant if the following properties hold:

- Compressing: $|\mathcal{Y}| < |\mathcal{X}|$

- Collision-Resistant: For all efficient adversaries A :

$$\Pr[k \leftarrow \mathcal{K}; (x, x') \leftarrow A(1^\lambda, k) : H(k, x) = H(k, x') \text{ and } x \neq x'] = \text{negl}(\lambda).$$

We can directly appeal to SIS to obtain a CRHF:

$$H: \mathbb{Z}_q^{n \times m} \times \{0, 1\}^m \rightarrow \mathbb{Z}_q^n$$

where we set $m > \lceil n \log q \rceil$. In this case, domain has size $2^m > 2^{n \log q} = q^n$, which is the size of the output space. Collision resistance follows assuming $SIS_{n,m,q,\beta}$ for any $\beta \geq \sqrt{\lceil n \log q \rceil}$

The SIS hash function supports efficient local updates:

Suppose you have a public hash $h = H(x)$ of a bit-string $x \in \{0,1\}^m$. Later, you want to update $x \mapsto x'$ where x and x' only differ on a few indices (e.g., updating an entry in an address book). For instance, suppose x and x' differ only on the first bit (e.g., $x_1 = 0$ and $x'_1 = 1$). Then observe the following

$$h = H(k, x) = A \cdot x = \begin{pmatrix} | & | & & | \\ a_1 & a_2 & \dots & a_m \\ | & | & & | \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ \vdots \\ x_m \end{pmatrix} = \sum_{i \in [m]} x_i a_i = \sum_{i=2}^m x_i a_i \quad \text{since } x_1 = 0$$

$$h' = H(k, x') = A \cdot x' = \sum_{i \in [m]} x'_i a_i = x'_1 a_1 + \sum_{i=2}^m x'_i a_i = a_1 + \sum_{i=2}^m x_i a_i = a_1 + h \quad \text{since } x'_i = x_i \text{ for all } i \geq 2$$

Thus, we can easily update h to h' by just adding to it the first column of A without (re)computing the full hash function.

The SIS assumption can also serve as the basis for digital signatures — to develop this, we will first need to introduce lattice trapdoors. Will define them first and construct them later.

Inhomogeneous SIS: given $A \stackrel{R}{\leftarrow} \mathbb{Z}_q^{n \times m}$ and $y \stackrel{R}{\leftarrow} \mathbb{Z}_q^n$, find $x \in \mathbb{Z}_q^m$ such that $Ax = y \in \mathbb{Z}_q^n$ and $\|x\| \leq \beta$

It turns out that this can actually be used as a trapdoor function. Namely, there exist efficient algorithms

- $\text{TrapGen}(n, m, q, \beta) \rightarrow (A, \text{td}_A)$: On input the lattice parameters n, m, q , the trapdoor-generation algorithm outputs a matrix $A \in \mathbb{Z}_q^{n \times m}$ and a trapdoor td_A
- $f_A(x) \rightarrow y$: On input $x \in \mathbb{Z}_q^m$, computes $y = Ax \in \mathbb{Z}_q^n$
- $f_A^{-1}(\text{td}_A, y) \rightarrow x$: On input the trapdoor td_A and an element $y \in \mathbb{Z}_q^n$, the inversion algorithm outputs a value $\|x\| \leq \beta$

Moreover, for a suitable choice of n, m, q, β , these algorithms satisfy the following properties:

- For all $y \in \mathbb{Z}_q^n$, $f_A^{-1}(\text{td}_A, y)$ outputs $x \in \mathbb{Z}_q^m$ such that $\|x\| \leq \beta$ and $Ax = y$
- The matrix A output by TrapGen is statistically close to uniform over $\mathbb{Z}_q^{n \times m}$

in fact, more general: can sample a solution to this system from a discrete Gaussian distribution (sampling needed to ensure solution does not leak trapdoor)

Digital signatures from lattice trapdoors: We can use lattice trapdoors to obtain a digital signature scheme in the random oracle model (this is essentially an analog of RSA signatures):

- $\text{KeyGen}: (A, \text{td}_A) \leftarrow \text{TrapGen}(n, m, q, \beta)$
Output $\text{vk} = A$ and $\text{sk} = \text{td}_A$
- $\text{Sign}(\text{sk}, m)$: Output $\sigma \leftarrow f_A^{-1}(\text{td}_A, H(m))$. Here, $H: \{0,1\}^* \rightarrow \mathbb{Z}_q^n$ is modeled as a random oracle (ideal hash function)
- $\text{Verify}(\text{vk}, m, \sigma)$: Check that $\|\sigma\| \leq \beta$ and that $f_A(\sigma) = H(m)$.

just as in RSA-FDH

Hardness reduces to hardness of inhomogeneous SIS (similar proof as RSA-FDH). Intuition:

- Matrix A output by KeyGen is uniformly random (property of KeyGen)
- To forge signature on message m^* , adversary needs to find short x such that $Ax = y$ where $y = H(m^*)$, which is uniform
- Formally: need to rely on random oracle to embed inhomogeneous SIS challenge + respond to signing queries

essentially solving inhomogeneous SIS instance

(ask in office hours for more details)

Suppose we know $R \in \mathbb{Z}_f^{m \times m}$ such that $AR = G$. We can then define the inversion algorithm as follows:

- $f_A^{-1}(td_A = R, y \in \mathbb{Z}_f^n)$: Output $x = R \cdot G^{-1}(y)$.

Important note: When using trapdoor functions in a setting where the adversary can see trapdoor evaluations, we actually need to randomize the computation of f_A^{-1} .

We check two properties:

1. $Ax = AR \cdot G^{-1}(y) = G \cdot G^{-1}(y) = y$ so x is indeed a valid pre-image

2. $\|x\| = \|R \cdot G^{-1}(y)\| \leq m \cdot \|R\| \|G^{-1}(y)\| = m \cdot \|R\|$

Thus, if $\|R\|$ is small, then $\|x\|$ is also small (think of β as a large polynomial in n).

Otherwise, we leak the trapdoor. But this basic scheme illustrates the main ideas--

Remaining question: How do we generate A together with a trapdoor (and so that A is statistically close to uniform)?

Many techniques to do so; we will look at one approach using the "leftover hash lemma" (also used when arguing security of Regev's PKE scheme)

Sample $\bar{A} \xleftarrow{R} \mathbb{Z}_f^{n \times m}$ and $\bar{R} \xleftarrow{R} \{0,1\}^{m \times m}$.

Set $A = [\bar{A} \mid \bar{A}\bar{R} + G] \in \mathbb{Z}_f^{n \times 2m}$

Output $A \in \mathbb{Z}_f^{n \times 2m}$, $td_A = R = \begin{bmatrix} -\bar{R} \\ I \end{bmatrix} \in \mathbb{Z}_f^{2m \times m}$

By construction that $AR = -\bar{A}\bar{R} + \bar{A}\bar{R} + G = G$, and moreover $\|R\| = 1$.

By leftover hash lemma, for $m = O(n \log f)$, $(\bar{A}, \bar{A}\bar{R})$ is indistinguishable from (\bar{A}, U) where $U \xleftarrow{R} \mathbb{Z}_f^{n \times m}$