

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

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

- Compressing: $|Y| < |X|$

- Collision-Resistant: For all efficient adversaries A :

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

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 $\text{SIS}_{n, m, q, \beta}$ 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 hash function is universal - this will be a very useful property (in conjunction with the leftover hash lemma)

Definition. Let $H: K \times X \rightarrow Y$ be a keyed hash function. We say H is ϵ -universal if for all $x_0, x_1 \in X$ where $x_0 \neq x_1$, $\Pr[k \xleftarrow{R} K : H(k, x_0) = H(k, x_1)] \leq \epsilon$. When $\epsilon = 1/|Y|$, we say it is universal.

Lemma. The SIS hash function $H: \mathbb{Z}_q^{n \times m} \times \{0, 1\}^m \rightarrow \mathbb{Z}_q^n$ is universal.

Proof. Take any $x_0, x_1 \in \{0, 1\}^m$ with $x_0 \neq x_1$. If $H(A, x_0) = H(A, x_1)$, then $A(x_0 - x_1) = 0$. Let $a_1, \dots, a_m \in \mathbb{Z}_q^n$ be columns of A . Then,

$$A(x_0 - x_1) = \sum_{i \in [m]} a_i (x_{0,i} - x_{1,i})$$

Since there exists some $j \in [m]$ where $x_{0,j} \neq x_{1,j}$, the above relation holds only if

$$a_j = \underbrace{(x_{1,j} - x_{0,j}) \sum_{i \neq j} a_i (x_{0,i} - x_{1,i})}_{\text{independent of } a_j}$$

Thus, $\Pr[A \xleftarrow{R} \mathbb{Z}_q^{n \times m} : A(x_0 - x_1) = 0]$

$$= \Pr[a_1, \dots, a_m \xleftarrow{R} \mathbb{Z}_q^n : a_j = (x_{1,j} - x_{0,j}) \sum_{i \neq j} a_i (x_{0,i} - x_{1,i})]$$

$$= \frac{1}{q^n}$$

Note: When q is prime, this argument also extends to any domain that is subset of \mathbb{Z}_q^n . Namely $H: \mathbb{Z}_q^{n \times m} \times \mathbb{Z}_q^m \rightarrow \mathbb{Z}_q^n$ is universal.

Definition. Let X be a random variable taking on values in a finite set S . We define the guessing probability of X to be

$$\max_{s \in S} \Pr[X=s]$$

We define the min-entropy of X to be

$$H_{\infty}(X) = -\log \max_{s \in S} \Pr[X=s]$$

Intuitively: if a random variable has k bits of min-entropy, then its most likely outcome occurs with probability at most 2^{-k} (i.e., there exists at least 2^k possible values for X)

Definition. Let D_0, D_1 be distributions with a common support S . Then, the statistical distance between D_0 and D_1 is defined to be

$$\Delta(D_0, D_1) = \frac{1}{2} \sum_{s \in S} |\Pr[t \leftarrow D_0 : t=s] - \Pr[t \leftarrow D_1 : t=s]|$$

If D_0 and D_1 are ϵ -close, then no adversary can distinguish with advantage better than ϵ

↳ When ϵ is negligible, we say the two distributions are statistically indistinguishable

denoted $D_0 \stackrel{\epsilon}{\approx} D_1$

↳ Contrast with computational indistinguishability which says no efficient adversary can distinguish

denoted $D_0 \stackrel{c}{\approx} D_1$

Theorem (Leftover Hash Lemma). Let $H: K \times X \rightarrow Y$ be an ϵ -universal hash function. Suppose $x \in X$ is a random variable with k bits of min-entropy. Then, define the following two distributions:

$$D_0: k \stackrel{R}{\leftarrow} K, y \leftarrow H(k, x); \text{ output } (k, y)$$

$$D_1: k \stackrel{R}{\leftarrow} K, y \stackrel{R}{\leftarrow} Y; \text{ output } (k, y)$$

The statistical distance between D_0 and D_1 is at most

$$\Delta(D_0, D_1) \leq \frac{1}{2} \sqrt{\frac{|Y|}{2^k} + (|Y| - 1)\epsilon}$$

When H is universal, this becomes

$$\Delta(D_0, D_1) \leq \frac{1}{2} \sqrt{|Y|/2^k} \quad (\text{since } \epsilon = 1/|Y|)$$

Typical setting: H is universal and $|Y| = 2^{k-2\lambda}$. By LHL, $(k, H(k, x)) \stackrel{\epsilon}{\approx} (k, y)$ where $y \stackrel{R}{\leftarrow} Y$.

This is an example of a "randomness extractor."

We have a source (x) with min-entropy, but not necessarily uniform.

We want to extract from it a uniform random value

LHL shows that universal hash functions can "smooth" out a non-uniform distribution

↓
Incurs loss of 2λ bits of entropy

Common application: extracting uniformly random cryptographic keys from non-uniform source

$$\rightarrow \text{consider } H: \mathbb{Z}_2^{n \times m} \times \{0, 1\}^m \rightarrow \mathbb{Z}_2^n$$

$$H(A, x) := Ax$$

could be binary representation of a group element

suitable for use as a symmetric key

Not typically used in practice because we need distribution with at least $n + 2\lambda$ bits of min-entropy (≥ 384 bits if $n = \lambda = 128$)
Practical heuristic: use random oracle

In lattices: If $A \stackrel{R}{\leftarrow} \mathbb{Z}_q^{n \times m}$ and $v \stackrel{R}{\leftarrow} \{0, 1\}^m$, then $Av \in \mathbb{Z}_q^n$ is uniform when $m > n \log q + 2\lambda$ typically, we just take n to be the security parameter and set $m = \Theta(n \log q)$

By a hybrid argument, if we sample $R \stackrel{R}{\leftarrow} \{0, 1\}^{m \times m}$, then AR is statistically close to uniform over $\mathbb{Z}_q^{n \times m}$

We will see this used in many constructions

Commitments from SIS (recall: commitment is a "sealed envelope")

- Setup (\mathcal{I}^λ) \rightarrow crs: Samples a common reference string
- Commit (crs, μ ; r) \rightarrow σ : Commits to a message μ with randomness r

} Here, opening can simply be the pair (m, r)
Verifier checks that $\sigma = \text{Commit}(\text{crs}, m; r)$

Useful building block for zero-knowledge proofs and other cryptographic protocols

- Setup (\mathcal{I}^λ): Let n, g be lattice parameters, and $m = \Theta(n \log g)$
Sample $A_1, A_2 \xleftarrow{r} \mathbb{Z}_g^{n \times m}$. Output crs = (A_1, A_2)
- Commit (crs, μ ; r): Output $\sigma = A_1 \mu + A_2 r$ where crs = (A_1, A_2)
 $\mu, r \in \{0, 1\}^m \rightarrow \sigma = [A_1 | A_2] \begin{bmatrix} \mu \\ r \end{bmatrix} \in \mathbb{Z}_g^n$


compressing when $m > n \log g$

Theorem (Statistically Hiding). If $m > 3n \log g$, then scheme is statistically hiding.

Proof. By the LHL, for $r \xleftarrow{r} \{0, 1\}^m$, $A_2 r \approx \text{Uniform}(\mathbb{Z}_g^n)$. Thus, $A_2 r$ acts as a one-time pad for $A_1 \mu$.

Theorem (Computational Hiding). Under $\text{SIS}_{n, 2m, g, \sqrt{2m}}$, the commitment scheme is computationally binding.

Proof. Suppose A can break the binding property. We use A to construct SIS adversary B:



If A is successful, then $\mu_1 \neq \mu_2$ and $[A_1 | A_2] \begin{bmatrix} \mu_1 \\ r_1 \end{bmatrix} = \sigma = [A_1 | A_2] \begin{bmatrix} \mu_2 \\ r_2 \end{bmatrix}$, which means $[A_1 | A_2] \begin{bmatrix} \mu_1 - \mu_2 \\ r_1 - r_2 \end{bmatrix} = 0$.
Since $\mu_1 \neq \mu_2$ this is a non-zero SIS solution with norm at most $\sqrt{2m}$.


Compare this with Pedersen commitments from discrete log:

Setup (\mathcal{I}^λ): Take a prime-order group $G \leftarrow \text{GroupGen}(\mathcal{I}^\lambda)$. Let p be the order of G.

Sample $g, h \xleftarrow{r} G$. Output crs = (g, h)

Commit (crs, μ ; r): Output $g^\mu h^r$.

$\mu, r \in \mathbb{Z}_p \rightarrow$



We will see many similar parallels between discrete log based systems and lattice-based systems