

Average-Case Hardness

Instructor: Daniele Micciancio

UCSD CSE

The goal of this lecture is to show that the Shortest Integer Solution problem (SIS) is hard to solve on average, for randomly chosen instances. We recall that SIS is parametrized by positive integers n, m, q , and an instance is described by a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ chosen uniformly at random. The goal is to find a short nonzero vector $\vec{x} \in \{0, 1, -1\}^m$ such that $\mathbf{A}\vec{x} = \vec{0} \pmod{q}$. However, if $\mathbf{A}$ is chosen uniformly at random, it would be computationally hard to find, even with very small (but non-negligible) probability. More specifically, we will show that it is at least as hard as approximating the Shortest Independent Vectors Problem (SIVP) in the worst-case on n -dimensional lattices, for some approximation factor $\gamma(n)$ polynomial in the lattice dimension n . A few words about the parameters:

- The SIS dimension n will be the same as the SIVP lattice dimension, and acts as a security parameter. Typical values are around $n = 200$.
- By the pigeon-hole principle, in order to guarantee that the SIS problem is solvable, we need $m > n \log_2 q$. For concreteness we will set $m = 2n \log_2 q$, so that the SIS problem defines a collision resistant hash function that compresses the input by a factor 2.
- For efficiency reasons, we would like the modulus q (and therefore also m) to be as small as possible. Small values of q are useful both to allow the implementation of the SIS function using arithmetics on small integers, and also to keep the overall size of $\mathbf{A}$ relatively small. Typical values of q are between 8 and 16 bits.
- The approximation factor γ should also be as small as possible, as higher values of γ make SIVP easier to solve. The best known polynomial time algorithm to solve SIVP achieve an approximation factor γ exponential in the lattice dimension n . We it is reasonable to conjecture that no polynomial time algorithm exists when $\gamma(n) = n^{O(1)}$ is polynomial. In theory, one can easily make up for larger (but still polynomial) values of γ by using a larger dimension. But using small values of γ is useful to get concretely efficient functions.

In the rest of these notes we will assume we have access to an SIS oracle for a given set of parameters that solves random instances of the problem with non-negligible probability, and show how to use it to solve SIVP in the worst case in polynomial time, and with a polynomial number of calls to SIS, and with success probability exponentially close to 1. Note that we are solving SIVP in the worst case, for a fixed n -dimensional lattice Λ , and the probability is just over the internal randomness of the algorithm. So, any intermediate step of the algorithm that uses the SIS oracle and succeeds with non-negligible probability, can be made to work with probability exponentially close to 1 using standard repetition techniques. If SIS works with non-negligible probability $\varepsilon = 1/n^{O(1)}$, repeating the step with independent randomness a polynomial number of times $1/\varepsilon = n^{O(1)}$ will succeed with constant probability $\Omega(1)$, and using a larger polynomial number of repetitions $n/\varepsilon = n^{O(1)}$ lowers the failure probability to $(1 - \Omega(1))^n = 2^{-\Omega(n)}$.

Consider the Shortest Independent Vectors Problem (SIVP), i.e., the problem of finding n linearly independent vectors $\mathbf{B} = [\vec{b}_1, \dots, \vec{b}_n]$ in a lattice Λ of rank n . We may initially set $\mathbf{B}$ to a basis describing the lattice. If we apply the LLL basis reduction algorithm to $\mathbf{B}$ we can obtain a set of linearly independent vectors $\mathbf{B}$ of length $\|\mathbf{B}\| = \max_i \|\vec{b}_i\| \leq 2^{O(n)} \lambda_n$. The goal is to find a much shorter set of linearly independent lattice vectors $\|\mathbf{B}\| \leq \gamma(n) \lambda_n$, for some $\gamma(n)$ polynomial in n . We will achieve this with the help of an oracle that solves random instances of the SIS problem, through a sequence of successive improvements. Throughout this process, $\mathbf{B}$ will always be a set of n linearly independent vectors in Λ , but not necessarily a basis for Λ . In other words, $\mathbf{B}$ may only generate a sublattice of Λ .

At every step we replace a longest vector in $\mathbf{B}$ with a vector which is shorter by some constant factor $c < 1$. In order to bound the number of iterations, we use the potential function

$$\Delta = \prod_i \|\vec{b}_i\| / \det(\Lambda)$$

and make the following two observations:

- At every iteration Δ decreases at least by a factor c .
- If $\Lambda' = \mathbf{B} \mathbb{Z}^n \subseteq \Lambda$ is the lattice generated by $\mathbf{B}$, then by Hadamard inequality we have

$$\prod_i \|\vec{b}_i\| \geq \det(\Lambda') \geq \det(\Lambda).$$

So, we always have $\Delta \geq 1$.

It follows that the number of iterations is bounded by $\log_c \Delta$, where Δ is the initial value of the potential function. If the initial basis is LLL-reduced, then $\Delta \leq 2^{n(n-1)/4} \det(\Lambda)$, and the number of iterations is at most $\log_c 2^{n(n-1)/4} \leq n^2 / (4 \log_2 c) = O(n^2)$, polynomial in the lattice dimension.

We can now focus on a single iteration. Let $\beta = \|\mathbf{B}\|$ be the current maximum length of the vectors in $\mathbf{B} = [\mathbf{B}', \vec{b}_n]$ and assume without loss of generality that $\|\vec{b}_n\| = \beta$, i.e., $\vec{b}_n$ is a longest vector in $\mathbf{B}$. We want to find a lattice vector $\vec{v} \in \Lambda$ such that

- $\vec{v}$ is linearly independent from the vectors in $\mathbf{B}'$, and
- it has length at most $\|\vec{v}\| \leq c\beta$.

Let $\vec{t} = (\beta/2) (\vec{b}_n^* / \|\vec{b}_n^*\|)$ be a target vector orthogonal to $\vec{v}$ of length $\beta/2$. Note that using $\mathbf{B}$ it is easy to find a lattice vector within distance $(\sqrt{n}/2) \beta$ from $\vec{t}$ (e.g., using the nearest plane algorithm,) but this is not good enough for our purposes. We will show that (provided that β is sufficiently large) with the help of an SIS oracle, one can find a lattice vector much closer to $\vec{t}$, specifically, within distance $\|\vec{v} - \vec{t}\| \leq \beta/3$. It will follow that

- $\vec{v}$ is at distance at least $\|\vec{t}\| - \|\vec{v} - \vec{t}\| \geq \frac{\beta}{2} - \frac{\beta}{3} > 0$ from the linear span of $\mathbf{B}'$. In particular, $\vec{t}$ is linearly independent from the vectors in $\mathbf{B}'$, and
- it has length at most $\|\vec{v}\| \leq \|\vec{t}\| + \|\vec{v} - \vec{t}\| \leq \frac{\beta}{2} + \frac{\beta}{3} = \frac{5}{6} \beta$.

So, it satisfies the required properties with $c = 5/6 < 1$.

Now assume we have access to an oracle SIS that on input a uniformly random matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ outputs a nonzero vector $\vec{x} \in \{0, 1, -1\}^m$ such that $\mathbf{A} \vec{x} = \vec{0} \pmod{q}$, for $q = \gamma \eta_\varepsilon(\mathbb{Z})$. Our final goal is to find linearly independent vectors of length $\beta \leq \gamma \eta_\varepsilon(\Lambda) \approx \gamma \lambda_n(\Lambda)$. So, assume $\beta > \gamma \eta_\varepsilon(\Lambda)$, i.e., we have not already found a set of sufficiently short vectors, and let $s = \sqrt{2} \beta / \gamma$. (The $\sqrt{2}$ factor may not seem necessary at first, but it is used later to apply a discrete gaussian convolution theorem.) By construction this value satisfies:

$$s = \sqrt{2} \eta_\varepsilon(\mathbb{Z}) \frac{\beta}{q} > \sqrt{2} \eta_\varepsilon(\Lambda).$$

Using $\mathbf{B}$ we can build a lattice basis for Λ/q with orthogonalized length bounded by β/q , and then sample lattice vectors $\vec{w}_i \in \Lambda/q$ with discrete gaussian distribution of parameter s . Sample m such vectors, with discrete gaussian distribution centered around the origin, except for a randomly chosen index j , for which we center it around $\vec{t}$. Set $\vec{a}_i = \vec{w}_i \pmod{\Lambda}$. Since $s > \eta_\varepsilon(\Lambda)$, these values are distributed almost uniformly at random in $\frac{1}{q} \Lambda / \Lambda \equiv \mathbb{Z}_{q^n}$. So, on input $\mathbf{A} = [\vec{a}_1, \dots, \vec{a}_m]$, SIS finds a solution $\vec{x}$ to $\sum_i \vec{a}_i x_i = \vec{0} \pmod{q}$, or, equivalently,

$$\vec{v} = \sum_i \vec{w}_i x_i \in \Lambda.$$

Moreover, since the index j is statistically independent of $\mathbf{A}$, and $\vec{x} \neq \vec{0}$, we have $x_j \neq 0$ with probability at least $\frac{1}{m}$. So, after changing the sign of $\vec{x}$ if $x_j = -1$, we may assume that $x_j = 1$.

It remains to show that $\vec{v}$ is within distance $\beta/3$ from the target $\vec{t}$. We will show that the conditional distribution of $\vec{v}$ (given $\mathbf{A}$, j and $\vec{x}$) is statistically close to a discrete gaussian on Λ of parameter $\|\vec{x}\| s \leq \sqrt{m} s$ centered around $\vec{t}$. From this it will follow that

$$\|\vec{v} - \vec{t}\| \leq \sqrt{\frac{2nm}{6}} s = \frac{\sqrt{nm}}{\sqrt{3}\gamma} \beta$$

except with exponentially small probability. Setting $\gamma = \sqrt{3nm}$ gives $\|\vec{v} - \vec{t}\| \leq \beta/3$.

It remains to analyze the (conditional) distribution of $\vec{v}$. Let $\vec{w} = (\vec{w}_1, \dots, \vec{w}_m)$ be the concatenation of the samples $\vec{w}_i$, and $\vec{c} = \vec{e}_j \otimes \vec{t} = (\vec{0}, \dots, \vec{t}, \dots, \vec{0})$ the concatenation of the centers of their gaussian distributions. Then, we have $\vec{v} = \mathbf{T} \vec{w}$, where $\mathbf{T} = \vec{x}^t \otimes \mathbf{I}$. Moreover, $\vec{w}$ is distributed as a discrete gaussian of parameter s over $\Lambda^m + \vec{a}$ centered around $\vec{c}$. In order to apply the discrete gaussian convolution theorem,¹ we need to verify that $K = \ker(\mathbf{T}) \cap \Lambda^k$ has smoothing parameter $\eta(K) \leq s$. Write $K = (\mathbf{R} \mathbb{Z}^{k-1}) \otimes \Lambda$ where $\mathbf{R} \in \mathbb{Z}^{k \times (k-1)}$ is the matrix with columns $\vec{r}_i = \vec{e}_i - x_i \vec{e}_j$ indexed by $i \in \{1, \dots, k\} \setminus \{j\}$. Then,² we have

$$\eta(K) \leq \|\mathbf{R}^*\| \eta(\Lambda) \leq \sqrt{2} \eta(\Lambda) = s.$$

So, we can conclude that $\vec{v}$ is statistically close to a discrete gaussian distribution over

$$\mathbf{T}(\Lambda + \vec{a}) = \Lambda + \sum_i \vec{w}_i x_i = \Lambda$$

centered around $\mathbf{T} \vec{c} = \vec{c}_j = \vec{t}$. Moreover, the covariance parameter of the discrete gaussian distribution is $s \mathbf{T} (s \mathbf{T})^t = s^2 \|\vec{x}\|^2 \leq s^2 m$.

1. See (Genise, Micciancio, Peikert & Walter – PKC 2020, Lemma 3.2)

2. See (Micciancio & Peikert – Crypto 2014, Corollary 2.8)