Ring-SIS and Ideal Lattices

Noah Stephens-Davidowitz
(for Vinod Vaikuntanathan's class)

1 Recalling h_A, and its inefficiency

As we have seen, the SIS problem yields a very simple collision-resistant hash function that is provably secure if worst-case lattice problems are hard: $h_A(\boldsymbol{e}) = A\boldsymbol{e} \bmod q$ where the key $A \in [q]^{m \times n}$ is uniformly random and the input is $\boldsymbol{e} \in \{0,1\}^m$. (Here, we are being less general than before and assuming that our input set is simply $\{0,1\}^m$.) Recall that finding an h_A collision is equivalent to solving the SIS problem, whose definition we reproduce below.

Definition 1. *For parameters n, m, q, the (average-case) Short Integer Solutions (SIS) problem is defined as follows. The input is a uniformly random matrix $A \in [q]^{n \times m}$. The goal is to find a non-zero vector $\boldsymbol{e} \in \{-1,0,1\}^m$ such that $A\boldsymbol{e} = \boldsymbol{0} \bmod q$.*

h_A has a lot going for it as a hash function. It is remarkably simple—a linear collision-resistant hash function! And, we saw that it is provably secure under the assumption that certain well-studied worst-case lattice problems are hard. If those two things are not enough, h_A is also worthy of study because of its close relationship with LWE, the topic of this course and an extremely important problem for cryptographers.

Unfortunately, h_A is quite inefficient, since just reading the public key takes time roughly $nm \log q > n^2$ (where the inequality follows from the fact that we must have $m > n$ in order for h_A to be a compressing function). But, h_A is breakable in time $2^{O(m)}$ (even by brute-force search). *Ideal*ly, we would hope for a hash function that can be broken in time $2^{O(m)}$ to run in time roughly linear in $m \approx n$. Our goal is therefore to show a variant of h_A whose running time is in fact roughly linear in n.

2 The cyclic shift matrix, and the ring $\mathbb{Z}[x]/(x^n - 1)$

Since just reading the key of h_A requires time greater than n^2, any attempt to speed up the computation of h_A will presumably have to compress the key size. E.g., we could take some short uniformly random seed r (with bit length, say, $O(n)$) and set $A = H(r)$ for some suitable expanding function H. If H is a PRG modeled as a random oracle, then the resulting hash function $h_{H(r)}$ retains its security. (This idea is actually quite useful in practice [BCD+16] in the context of LWE.) However, if H is an arbitrary function, then we do not expect to be able to compute $h_{H(r)}(\boldsymbol{e})$ in time faster than n^2. So, though this idea immediately yields a hash function with a smaller key, we need to do more work to get a faster hash function.

In order to speed up our computation, we presumably need our matrix A to be a very special function of the seed. To that end, we take our short random seed to be $\ell = m/n$ uniformly random

vectors $\boldsymbol{a}_1, \ldots, \boldsymbol{a}_\ell \in [q]^n$, and we take the columns of our matrix A to be the vectors $\boldsymbol{a}_1, \ldots, \boldsymbol{a}_\ell$ together with all "cyclic rotations" of the $\boldsymbol{a}_i$. I.e., for $\boldsymbol{a} = (a_1, \ldots, a_n)^T \in \mathbb{Z}^n$, we define

$$\mathrm{Rot}(\boldsymbol{a}) := \begin{pmatrix} a_1 & a_n & \cdots & a_3 & a_2 \\ a_2 & a_1 & \cdots & a_4 & a_3 \\ a_3 & a_2 & \cdots & a_5 & a_4 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ a_{n-2} & a_{n-3} & \cdots & a_n & a_{n-1} \\ a_{n-1} & a_{n-2} & \cdots & a_1 & a_n \\ a_n & a_{n-1} & \cdots & a_2 & a_1 \end{pmatrix} \in \mathbb{Z}^{n \times n} ,$$

where each column is a simple cyclic permutation of the previous column.[1] Matrices of the form $\mathrm{Rot}(\boldsymbol{a})$ are sometimes referred to as "cyclic matrices" or "circulant matrices." We then take

$$A = (\mathrm{Rot}(\boldsymbol{a}_1), \mathrm{Rot}(\boldsymbol{a}_2), \ldots, \mathrm{Rot}(\boldsymbol{a}_\ell)) \in \mathbb{Z}^{n \times m} .$$

We claim that for A with this structure, we can compute $A\boldsymbol{e} \bmod q$ in time $n\ell \cdot \mathrm{polylog}(n, q)$. This is because the set of all integer cyclic matrices, $\widetilde{R} := \{\mathrm{Rot}(\boldsymbol{a}) \ : \boldsymbol{a} \in \mathbb{Z}^n\}$ is actually a very nice set with nice algebraic structure. In particular, we can write

$$\mathrm{Rot}(\boldsymbol{a}) = (\boldsymbol{a}, X\boldsymbol{a}, \ldots, X^{n-1}\boldsymbol{a}) ,$$

where

$$X := \begin{pmatrix} 0 & 0 & \cdots & 0 & 1 \\ 1 & 0 & \cdots & 0 & 0 \\ 0 & 1 & \cdots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \cdots & 1 & 0 \end{pmatrix} \in \{0,1\}^{n \times n}$$

is the "cyclic shift" matrix. Notice that $\widetilde{R} \subset \mathbb{Z}^{n \times n}$ is a lattice in $n \times n$ dimensions with rank n and basis $I_n, X, X^2, \ldots, X^{n-1}$. Indeed, for any $\boldsymbol{a} = (a_1, \ldots, a_n)^T \in \mathbb{Z}^n$, we can write

$$\mathrm{Rot}(\boldsymbol{a}) = a_1 I_n + a_2 X + \cdots + a_n X^{n-1} .$$

This identity immediately shows us that $\widetilde{R}$ is actually closed under (matrix) multiplication, and that multiplication is commutative over $\widetilde{R}$. I.e., $\widetilde{R}$ is a ring!

In fact, $\widetilde{R}$ is isomorphic to the polynomial ring $R := \mathbb{Z}[x]/(x^n - 1)$. I.e., R is the ring of polynomials in the variable x of degree at most $n-1$ and integral coefficients, with addition defined in the obvious way and multiplication defined by the distributive law together with the identity

$$x \cdot x^i = \begin{cases} x^{i+1} & i < n-1 \\ 1 & i = n-1 \end{cases} .$$

(The polynomial $x^n - 1$ is the characteristic polynomial of the cyclic shift matrix X, which is why it arises in this context.) To see that these two rings are isomorphic, one only needs to check that the map $X \mapsto x$ is a bijection that preserves addition and multiplication of basis elements.

[1] Notice that the definition of Rot does not depend at all on q. It is convenient to forget about q for now and to think of $\boldsymbol{a}$ as some arbitrary vector in $\mathbb{Z}^n$.

So, there's no reason to drag these $n \times n$ matrices around, and we can instead think of $\mathrm{Rot}(\boldsymbol{a}) \in \widetilde{R}$ as the corresponding polynomial $a \in R$ of degree at most $n-1$. (I.e., we change notation slightly.) We can therefore identify our matrix $A \in [q]^{n \times m}$ with a tuple of ring elements $(a_1, \ldots, a_\ell)^T \in R_{[q]}^\ell$, and similarly the input $\boldsymbol{e} \in \{0,1\}^m$ is a tuple of ring elements $(e_1, \ldots, e_\ell)^T \in R_{\{0,1\}}^\ell$, where we use the notation R_S to represent the set of polynomials in R with coefficients in S. Therefore, our hash function is now $h_{a_1,\ldots,a_\ell}(e_1, \ldots, e_\ell) = a_1 e_1 + \cdots + a_\ell e_\ell \bmod qR$.[2] For convenience, we abbreviate this by $h_a(e)$.

Now, to gain in efficiency, we simply recall that we can multiply two elements in $R_{[q]}$ in time $n \cdot \mathrm{polylog}(n, q)$ via the fast Fourier transform. Therefore, we can compute h_a in time $\ell n \cdot \mathrm{polylog}(n, q) = m \cdot \mathrm{polylog}(n, q)$, which is a tremendous speedup over the $nm \cdot \mathrm{polylog}(q)$ running time of the original h_A. Indeed, we typically think of $q = \mathrm{poly}(n)$ and $\ell = \mathrm{polylog}(n)$, so that this running time is quasilinear in n.

2.1 (In)security of this new hash function, and Ring-SIS

Of course, this is not very useful if h_a is not secure. In fact, Micciancio showed that h_a is secure *as a one-way function* (under a plausible worst-case lattice assumption) [Mic07]. I.e., with certain reasonable parameters, it is difficult to invert h_a on a random input. This result is really quite remarkable, but we will not state it formally, since we will soon see that another hash function appears to be a much better choice.

Unfortunately, h_a is *not* a collision-resistant hash function. To see this, it helps to define the Ring-SIS problem, which is the analogue of SIS in this setting.

Definition 2. *For a ring R, integer modulus $q \geq 2$, and integer $\ell \geq 1$, the (average-case) Ring-SIS problem is defined as follows. The input is $a_1, \ldots, a_\ell \in R_{[q]}$ sampled independently and uniformly at random. The goal is to output $e_1, \ldots, e_\ell \in R_{\{-1,0,1\}}$ not all zero such that $a_1 e_1 + \cdots + a_\ell e_\ell = 0 \bmod qR$.*

One can easily see that finding a collision in h_a is equivalent to solving Ring-SIS, just like finding a collision in h_A is equivalent to solving SIS. Unfortunately, Ring-SIS over $\mathbb{Z}[x]/(x^n - 1)$ is not hard. The issue is that the ring $\mathbb{Z}[x]/(x^n - 1)$ has non-trivial zero divisors (i.e., it is not an integral domain). To see this, let $\widetilde{e} = 1 + x + x^2 + \cdots + x^{n-1} \in \mathbb{Z}[x]/(x^n - 1)$, and notice that $(x-1)\widetilde{e} = x^n - 1 = 0$. (In terms of Rot and $\widetilde{R}$, this corresponds to the fact that $\mathrm{Rot}(\boldsymbol{u})$ is singular, where $\boldsymbol{u} = (1, 1, \ldots, 1)^T \neq \boldsymbol{0}$.) This leads to an attack.

Claim 2.1. *For any integer modulus $q \geq 2$ and integer $n \geq 1$, let $R := \mathbb{Z}[x]/(x^n - 1)$ and let $\widetilde{e} = 1 + x + x^2 + \cdots + x^{n-1} \in R_{-1,0,1}$. Then, $a\widetilde{e} = 0 \bmod qR$ with probability $1/q$ when $a \in R_{[q]}$ is sampled uniformly at random.*

In particular, $\widetilde{e}, 0, \ldots, 0 \in R_{\{-1,0,1\}}$ is a solution to Ring-SIS over R with probability $1/q$, and the hash function h_a can be broken efficiently with probability $1/q$.

Proof. Suppose that $a \in R_{[q]}$ is divisible by $x - 1$ modulo qR. I.e., $a = (x-1)a' \bmod qR$. Then, $\widetilde{e}a = \widetilde{e}(x-1)a' = 0 \bmod qR$. The result follows by noting that $a \in R_{[q]}$ is divisible by $x - 1$ modulo

[2] Here, we have chosen to think of the e_i as ring elements as well. This is formally justified by the identity

$$\mathrm{Rot}(\boldsymbol{a}) \cdot \mathrm{Rot}(\boldsymbol{e}) = \mathrm{Rot}\big(\mathrm{Rot}(\boldsymbol{a}) \cdot \boldsymbol{e}\big) .$$

The reduction mod qR simply means that we reduce the coefficients of the result of our polynomial multiplication modulo q.

qR with probability $1/q$. (Notice that being divisible by $x-1$ is equivalent to having coefficients that sum to zero mod q.) □

If our original hash function h_A is in fact $2^{\Omega(n)}$ secure, then this result makes h_a uninteresting as a collision-resistant hash function. In particular, in order for h_a to have a chance of matching this security, we would need to take $q = 2^{\Omega(n)}$, in which case h_a is actually a slower hash function than h_A.

One might complain that the above attack is a bit unsatisfying because it can be thwarted simply by throwing out "bad keys" (i.e., keys $a_1, \ldots, a_\ell$ such that $x-1$ divides one of the a_i modulo q). Plus, the above attack might not be too damaging if it turns out that SIS is not actually $2^{\Omega(n)}$ hard. So, for completeness, we note that Ring-SIS over $\mathbb{Z}[x]/(x^n-1)$ actually reduces to SIS with $n=1$ and $m=\ell$, which shows that h_a is completely useless as a collision-resistant hash function. We might as well take $n=1$, in which case $\mathbb{Z}[x]/(x^n-1)$ is just the integers!

Claim 2.2. *For any integer modulus $q \geq 2$ and integers $n, \ell \geq 1$, let $R := \mathbb{Z}[x]/(x^n-1)$. Then, Ring-SIS over R with parameters (q, ℓ) reduces efficiently to SIS with parameters $(q, n=1, m=\ell)$.*

I.e., h_a with $n > 1$ is no more secure (as a collision-resistant hash function) than h_a with $n = 1$!

Proof. For $i = 1, \ldots, \ell$, let $\alpha_i = a_i \bmod (q, x-1) \in [q]$. (I.e., α_i is the sum of the coordinates of a_i modulo q.) The reduction calls its SIS oracle on input $\alpha_1, \ldots, \alpha_\ell$ and receives as output $\varepsilon_1, \ldots, \varepsilon_\ell \in \{-1, 0, 1\}$ not all zero such that $\varepsilon_1\alpha_1 + \cdots + \varepsilon_\ell\alpha_\ell = 0 \bmod q$. Finally, the reduction simply outputs $\varepsilon_1\widetilde{e}, \cdots, \varepsilon_\ell\widetilde{e} \in R_{\{-1,0,1\}}$, where $\widetilde{e} = 1 + x + \cdots + x^{n-1}$.

To see that this reduction is correct, notice that $A := a_1\varepsilon_1 + \cdots + a_\ell\varepsilon_\ell$ is divisible by $x-1$ modulo q. Since $(x-1)\widetilde{e} = 0$, we immediately see that $A\widetilde{e} = 0 \bmod qR$, as needed. □

3 The ring $\mathbb{Z}[x]/(x^n+1)$, ideal lattices, and a secure collision-resistant hash function

Recall that our attack on h_a over $\mathbb{Z}[x]/(x^n-1)$ relied on the fact that x^n-1 has a nontrivial factor over the integers, $x^n - 1 = (x-1)(x^{n-1} + x^{n-2} + \cdots + 1)$. So, it is natural to try replacing x^n-1 with an irreducible polynomial. Indeed, one can easily show that $\mathbb{Z}[x]/(p(x))$ for some polynomial $p(x) \in \mathbb{Z}[x]$ is an integral domain if and only if p is irreducible.

We strongly prefer sparse polynomials with small coefficients (both because they are easy to work with and because this ensures that our ring has nice "geometric" properties). Since x^n-1 failed, we try x^n+1. This is irreducible over $\mathbb{Z}$ if and only if n is a power of two.[3] So, we take $R := \mathbb{Z}[x]/(x^n+1)$ for n some power of two. I.e., R is the ring of polynomials over $\mathbb{Z}$ of degree at most $n-1$ with addition defined in the obvious way and multiplication defined by

$$x \cdot x^i = \begin{cases} x^{i+1} & i < n-1 \\ -1 & i = n-1 \end{cases}.$$

[3] If $p > 1$ is a non-trivial odd factor of n, then $x^{n/p} + 1$ is a non-trivial factor of $x^n + 1$. If n has no odd factors, then $x^n + 1$ is the $2n$th cyclotomic polynomial—i.e., the minimal polynomial over $\mathbb{Z}$ of any primitive $2n$th root of unity.

From the matrix perspective of the previous section, this corresponds to taking

$$\mathrm{Rot}(\boldsymbol{a}) = (\boldsymbol{a}, X\boldsymbol{a}, \ldots, X^{n-1}\boldsymbol{a}) = \begin{pmatrix} a_1 & -a_n & \cdots & -a_3 & -a_2 \\ a_2 & a_1 & \cdots & -a_4 & -a_3 \\ a_3 & a_2 & \cdots & -a_5 & -a_4 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ a_{n-2} & a_{n-3} & \cdots & -a_n & -a_{n-1} \\ a_{n-1} & a_{n-2} & \cdots & a_1 & -a_n \\ a_n & a_{n-1} & \cdots & a_2 & a_1 \end{pmatrix} \in \mathbb{Z}^{n \times n} ,$$

where

$$X := \begin{pmatrix} 0 & 0 & \cdots & 0 & -1 \\ 1 & 0 & \cdots & 0 & 0 \\ 0 & 1 & \cdots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \cdots & 1 & 0 \end{pmatrix} \in \{0,1\}^{n \times n} .$$

Notice that X differs in just one entry from our choice in the previous section. Matrices of the form $\mathrm{Rot}(\boldsymbol{a})$ as above are occasionally called "anti-cyclic."

As before, we define our hash function $h_a(e) = a_1 e_1 + \cdots + a_\ell e_\ell \mod qR$, where the a_i are chosen uniformly $a_i \in R_{[q]}$ and $e_i \in R_{\{0,1\}}$. But, we stress that the underlying ring has changed from $\mathbb{Z}[x]/(x^n - 1)$ to $R = \mathbb{Z}[x]/(x^n + 1)$, so that this is not the same hash function as before. (Formally, we should include the ring as a parameter in h, i.e. $h_{a,\mathbb{Z}[x]/(x^n+1)}$, to distinguish it, but we prefer to keep the notation uncluttered.) As before, finding a collision for this hash function is equivalent to solving Ring-SIS, now over this new ring, $\mathbb{Z}[x]/(x^n + 1)$. As we will see, Ring-SIS is in fact hard over this ring, under a reasonable worst-case complexity assumption.

Remark. *The author feels compelled to note that the ring R is rather special; it is the ring of integers of the cyclotomic number field $\mathbb{Q}[x]/(x^n + 1)$. Number fields and their rings of integers are very well-studied and very interesting objects, and these notes stop just short of presenting some of the beautiful mathematics that is lurking beneath the surface here. (The fact that R is such a rich mathematical object also seems relevant for the security of h_a. In particular, we will mention later that there are algorithmic results for related problems that exploit rather deep properties of R [Ber14, CGS14, CDPR16, CDW17].)*

3.1 Ideal lattices

In order to present the worst-case hardness assumption that will imply the security of our hash function, we will need to introduce a special class of lattices known as *ideal lattices*. Recall that a lattice is an additive subgroup of $\mathbb{Z}^n$. I.e., a subset of $\mathbb{Z}^n$ closed under addition and subtraction. An ideal $\mathcal{I} \subseteq R$ is an additive subgroup of a ring R that is closed under multiplication by any ring element. I.e., $\mathcal{I}$ is closed under addition and subtraction, *and* for any $y \in \mathcal{I}$ and $r \in R$, we have $ry \in \mathcal{I}$.

For our choice of ring, we can view $\mathcal{I}$ as a lattice by embedding R in $\mathbb{Z}^n$ via the trivial embedding that maps x^i to the unit vector $\boldsymbol{e}_i$. So, $\mathcal{I}$ can equivalently be viewed as a lattice $\mathcal{I} \subseteq \mathbb{Z}^n$ that is invariant under the linear transformation X. I.e., $\mathcal{I} \subseteq \mathbb{Z}^n$ is a lattice such that $(y_1, \ldots, y_n)^T \in \mathcal{I}$ if

and only if $(-y_n, y_1, y_2, \ldots, y_{n-1})^T \in \mathcal{I}$. Such lattices are sometimes called "anti-cyclic," and the corresponding lattices over $\mathbb{Z}[x]/(x^n - 1)$ are often called "cyclic."

In particular, this embedding allows us to consider the geometry of an ideal $\mathcal{I}$, as a subset of $\mathbb{Z}^n$. E.g., we can define the ℓ_2 norm and the inner product over $\mathcal{I}$ by taking the ℓ_2 norm and the inner product over $\mathbb{Z}^n$.[4] We then see that ideal lattices $\mathcal{I}$ are a strange class of lattices in which non-zero lattice elements $y \in \mathcal{I}$ can be divided into groups of n linearly independent elements, $y, xy, x^2y, \ldots, x^{n-1}y$, all with the same length, $\|x^i y\| = \|x^j y\|$. In particular $\lambda_1(\mathcal{I}) = \lambda_n(\mathcal{I})$. (Notice that we move freely between the representation of R as $\mathbb{Z}^n$ and the representation of R as a polynomial ring. I.e., we can think of $y_1, y_2 \in R$ as scalars, written in plain font, as opposed to boldface vectors $\boldsymbol{y}_1, \boldsymbol{y}_2 \in \mathbb{Z}^n$. We can still talk about their norms $\|y_1\|, \|y_2\|$ and inner product $\langle y_1, y_2 \rangle$.)

Remark. *Ideals are very important objects in the study of rings, and they have a rich history that we do not discuss here. In fact, much of the early study of lattices was motivated by the study of the geometry of ideals, going back all the way to the seminal work of Minkowski, Dirichlet, and others in the middle of the 19th century.*

3.2 SVP over ideal lattices and worst-case hardness

For our purposes, this view of ideals as lattices is useful because it allows us to extend computational lattice problems to ideals. I.e., for some fixed ring R, we can define the computational problems γ-IdealSVP, γ-IdealSIVP, γ-GapIdealSVP, etc., as the corresponding computational problems restricted to ideal lattices. In fact, the above discussion shows that γ-IdealSVP and γ-IdealSIVP are equivalent over our ring $R = \mathbb{Z}[x]/(x^n + 1)$. A slightly more sophisticated argument shows that γ-GapIdealSVP is easy over R for $\gamma > \sqrt{n}$ because the length of the shortest vector in an ideal can be approximated up to a factor of $\sqrt{n}$ by the determinant. We therefore only present a formal definition of γ-IdealSVP.

Definition 3. *For a ring R (with an associated norm $\|\cdot\|$) and approximation factor $\gamma \geq 1$, γ-IdealSVP over R is the approximate search problem defined as follows. The input is (a basis for) an ideal lattice $\mathcal{I}$ over R. The goal is to output a non-zero element $y \in \mathcal{I}$ with $\|y\| \leq \gamma \lambda_1(\mathcal{I})$*

With this, we can present the worst-case to average-case hardness of Ring-SIS, which was discovered independently by Peikert and Rosen [PR06] and by Lyubashevsky and Micciancio [LM06].

Theorem 3.1 ([PR06, LM06])**.** *For any power of two n, integer $\ell \geq 1$, and integer modulus $q \geq 2n^2\ell$, γ-Ideal SVP over $R = \mathbb{Z}[x]/(x^n + 1)$ can be efficiently reduced to Ring-SIS over R, where $\gamma = \ell \cdot \text{poly}(n)$.*

3.3 Regarding the hardness of IdealSVP

Of course, Theorem 3.1 is only interesting if γ-IdealSVP is hard over the ring $\mathbb{Z}[x]/(x^n + 1)$. Until very recently, our best algorithms for this problem were essentially no better than our generic algorithms for γ-SVP over general n-dimensional lattices. However, very recently, polynomial-time

[4] For more general rings of integers over number fields, there is actually a different notion of geometry obtained via the "canonical embedding" of $\mathcal{I}$ into $\mathbb{C}^n$, which has very nice properties. E.g., in the canonical embedding, ring multiplication is coordinate-wise. For our very special choice of ring, $\mathbb{Z}[x]/(x^n + 1)$ for n a power of two, these two embeddings actually yield the same geometry.

quantum algorithms for γ-IdealSVP with the very large approximation factor $\gamma = 2^{\widetilde{O}(\sqrt{n})}$ were discovered in a series of works [Ber14, CGS14, CDPR16, CDW17]. (The best known algorithms for $2^{\sqrt{n}}$-SVP run in time roughly $2^{\sqrt{n}}$, even on a quantum computer. And, our best polynomial-time algorithms for γ-SVP only achieve an approximation factor of $\gamma = 2^{\widetilde{\Theta}(n)}$. So, this is a very big improvement.)

These algorithms are not known to extend to attacks on Ring-SIS for two reasons. First, the approximation factor $\gamma = 2^{\widetilde{O}(\sqrt{n})}$ is much larger than the approximation factors that are relevant to Ring-SIS. Second, Ring-SIS is not exactly an ideal lattice problem. Instead, notice that a solution to Ring-SIS consists of a *vector* of ring elements $(e_1, \ldots, e_\ell) \in R^\ell$. Indeed, Ring-SIS is technically a lattice problem over rank ℓ *modules*. It is therefore not currently known how to efficiently reduce Ring-SIS to IdealSVP.

As a result of all of this, the status of Ring-SIS is a bit unclear at the moment. The barriers mentioned in the previous paragraph seem to be quite hard to overcome, so perhaps this new line of research will not lead to an attack. As far as we know, Ring-SIS is just as hard as SIS, and indeed, as far as we know, it could yield a collision-resistant hash function that is computable in $\widetilde{O}(n)$ time and only breakable in time $2^{\Omega(n)}$.

3.4 The reduction

Finally, we present the worst-case to average-case reduction for Ring-SIS, which is a slight variant of the reduction that we have already seen for SIS. For simplicity, we leave out some details, sweep some technical issues under the rug, and assume that the reader is familiar with the presentation of the SIS reduction from an earlier lecture.

Proof of Theorem 3.1. The reduction receives as input some ideal $\mathcal{I} \subseteq R$. We may assume without loss of generality that it is also given some parameter s such that $s/2 \leq \sqrt{n}\lambda_n(\mathcal{I}) \leq s$ and some non-zero element $b \in \mathcal{I}$ in the ideal with not-too-large norm, say $\|b\| \leq 2^n \lambda_n(\mathcal{I})$. (Such an element b can be found by running the LLL algorithm, and we can simply try many different parameters $s_i = \sqrt{n}\|b\|/2^{i/2}$ for $i = 0, \ldots, 2n$ to obtain s.) As in the reduction for SIS, we will show a reduction that makes "slow progress." I.e., it finds a non-zero $b' \in \mathcal{I}$ with

$$\|b'\| \leq \ell n^2 \|b\|/q + \ell n^{1.5} s \ .$$

We may repeat this procedure, say, $10n$ times to eventually find a vector of length at most, say,

$$10\ell n^{1.5} s \leq 20 n^2 \lambda_n(\mathcal{I}) = 20 n^2 \ell \lambda_1(\mathcal{I}) \ ,$$

as needed.

For simplicity, we first assume that $\mathcal{I}$ is a principal ideal generated by b. I.e., every element in $\mathcal{I}$ can be written as rb for some $r \in R$. This is definitely not true in general,[5] and we will sketch how to remove this assumption at the end of the proof.

The reduction samples $\boldsymbol{y}_1, \ldots, \boldsymbol{y}_\ell \in \mathbb{R}^n$ from the (continuous) Gaussian distribution with parameter s. Let $\boldsymbol{b} \in \mathbb{Z}^n$ be the vector of coefficients of b, and let $\boldsymbol{y}'_i \in \mathcal{I}/q$ be $\boldsymbol{y}_i$ with its coordinates in $\mathrm{Rot}(\boldsymbol{b})$ rounded to the nearest integer multiple of $1/q$. Let $y'_i \in R/q$ be the associated polynomial with coefficients given by the vector $\boldsymbol{y}'_i$. Finally, let $a_i \in R_{[q]}$ such that

[5] Most ideals are not principal at all—i.e., there is no element b that generates the ideal. Even if our ideal is principal, our specific b will typically not be a generator.

$ba_i = qy'_i \bmod qR$. The reduction calls the Ring-SIS oracle on input $a_1, \ldots, a_\ell$, receiving as output non-zero $e_1, \ldots, e_\ell \in R_{\{-1,0,1\}}$ such that $a_1 e_1 + \cdots + a_\ell e_\ell = 0 \bmod qR$. (We might actually need to repeat this procedure many times to receive valid output from the oracle, but we ignore this here.) The reduction outputs

$$b' = y'_1 e_1 + \cdots + y'_\ell e_\ell \ .$$

We first note that the input $a_1, \ldots, a_\ell \in R_{[q]}$ is statistically close to random because of our choice of $s > \sqrt{n}\lambda_n(\mathcal{I})$, just like in the reduction for SIS. So, the input to the Ring-SIS oracle is distributed correctly.

We next notice that $b' \in \mathcal{I}$. In particular, we see from the definition of a_i that

$$b' = (ba_1 e_1 + \cdots + ba_\ell e_\ell)/q \bmod \mathcal{I} \ .$$

Since the right-hand side is in $\mathcal{I}$, b' is as well.

We next study the length of b'. To do this, we need the inequality $\|ye\| \leq \sqrt{n}\|y\|\|e\|$ for any $e, y \in R$, which follows from the definition of polynomial multiplication together with the Cauchy-Schwarz inequality. Therefore,

$$\|b'\| \leq \sum_{i=1}^{\ell} \|y'_i e_i\| \leq \sqrt{n} \sum_{i=1}^{\ell} \|y'_i\|\|e_i\| \leq n \sum_{i=1}^{\ell} \|\boldsymbol{y}'_i\| \leq \ell n^2 \|b\|/q + n \sum_{i=1}^{\ell} \|\boldsymbol{y}_i\| \ ,$$

where the last inequality follows from the fact that $\|\boldsymbol{y}'_i - \boldsymbol{y}_i\| \leq \sum \|x^j b\|/q = n\|b\|/q$. And, since $\|y_i\|$ was sampled from a Gaussian with parameter s, we have $\|\boldsymbol{y}_i\| \leq \sqrt{n} \cdot s$ except with negligible probability. The result follows.

So, b' is a short element in the ideal, but we must still show that it is non-zero. Notice that the oracle's input depends only on the coset $y_i \bmod \mathcal{I}$, and if $e_i \neq 0$, there is at most one value of y_i in this coset that yields $b' = 0$. (Notice that this fact is true over an integral domain, when there are no non-trivial zero divisors in our ring. If we used the ring $\mathbb{Z}[x]/(x^n - 1)$ instead, then there could potentially be many values of y_i that cause $b' = 0$, such as when $e_i = \widetilde{e}$ from our earlier attack.) Just like in the SIS case, our choice of parameter $s > \sqrt{n}\lambda_n(\mathcal{I})$ guarantees that y_i has high entropy, even conditioned on its coset. Therefore, b' will often be non-zero.

Finally, we sketch how to remove the assumption that $\mathcal{I}$ is principal. The basic idea is just to still do the reduction but to work with the lattice $\mathcal{I}'$ generated by b. The problem is that we might have $\lambda_n(\mathcal{I}') > s/\sqrt{n}$. (Indeed $\mathcal{I}'$ might not have any vectors shorter than b.) This causes issues in two steps of the proof: when we argue that a_i is uniformly random, and when we argue that b' is non-zero.

To ensure that a_i is still uniformly random, we add to each $\boldsymbol{y}_i$ an element $\boldsymbol{v}_i \in \mathcal{I}$ that is uniformly random mod $\mathcal{I}'$. We can then "subtract out" $\boldsymbol{v}_i$ later to ensure that it does not increase the length of b'. This effectively gives us a short vector in a coset $\mathcal{I}' + v$, where v is an R-linear combination of the $\boldsymbol{v}_i$. In fact, once we have done this, we can show that the distribution over $\boldsymbol{v}_i$ of the input to the oracle depends only on the coset of $\boldsymbol{y}_i$ modulo $\mathcal{I}$. It follows that our output vector b' has high entropy and is therefore rarely zero. □

References

[BCD+16] Joppe W. Bos, Craig Costello, Léo Ducas, Ilya Mironov, Michael Naehrig, Valeria Nikolaenko, Ananth Raghunathan, and Douglas Stebila. Frodo: Take off the ring! practical, quantum-secure key exchange from lwe. In *CCS*, 2016.

[Ber14] Daniel J. Bernstein. A subfield-logarithm attack against ideal lattices, 2014.

[CDPR16] Ronald Cramer, Léo Ducas, Chris Peikert, and Oded Regev. Recovering short generators of principal ideals in cyclotomic rings. In *EUROCRYPT*, 2016.

[CDW17] Ronald Cramer, Léo Ducas, and Benjamin Wesolowski. Short stickelberger class relations and application to ideal-svp. 2017.

[CGS14] Peter Campbell, Michael Groves, and Dan Shepherd. Soliloquy: a cautionary tale. ETSI 2nd Quantum-Safe Crypto Workshop, 2014.

[LM06] Vadim Lyubashevsky and Daniele Micciancio. Generalized compact knapsacks are collision resistant. In *ICALP (2)*, 2006.

[Mic07] Daniele Micciancio. Generalized compact knapsacks, cyclic lattices, and efficient one-way functions. *Computational Complexity*, 16(4), 2007.

[PR06] Chris Peikert and Alon Rosen. Efficient collision-resistant hashing from worst-case assumptions on cyclic lattices. In *TCC*, 2006.