FHE-MPC Notes

Lecturer: Nigel Smart Lecture #3
Scribe: Peter Scholl 24/10/11

In this lecture the LWE and ring-LWE problems were introduced, and it was described how these can be used to create cryptosystems based on lattice problems. Some possible parameter choices for ensuring the security of these schemes are discussed.

1 LWE problem

The *learning with errors* (LWE) problem is to efficiently distinguish vectors created from a 'noisy' set of linear equations between uniformly random vectors. Given a matrix $A \in \mathbb{Z}_q^{m \times n}$ and a vector $\mathbf{v} \in \mathbb{Z}_q^m$, the goal is to determine whether $\mathbf{v}$ has been sampled uniformly at random from $\mathbb{Z}_q^m$ or whether $\mathbf{v} = A\mathbf{s} + \mathbf{e}$ for some random $\mathbf{s} \in \mathbb{Z}_q^m$ and $\mathbf{e} \in \chi^m$, where χ is a small 'noise' distribution over $\mathbb{Z}_q$.

Observe that this is phrased as a decision problem. The *search* version of LWE (to recover the vector $\mathbf{s}$, given polynomially many samples of the form $A\mathbf{s} + \mathbf{e}$) can be shown to be equivalent.

The problem is very closely related to coding theory. If we choose the parameter $q = 2$, this becomes the well-studied *learning parity with noise* (LPN) problem, which is believed to be hard. Recovering the key from the more general LWE problem is essentially equivalent to decoding a noisy linear code, also a long established difficult problem in coding theory. However, for modern cryptographic purposes it is more important to ensure indistinguishability of encryptions rather than just security against key recovery. For this purpose it helps to look at the problem from a lattice-based perspective.

The vector $\mathbf{v} = A\mathbf{s} + \mathbf{e}$ can be seen as an element of the q-ary lattice $\Lambda_q(A^T)$ with a small perturbation vector added. The task here is to distinguish this from a uniformly random vector. In 2005, Regev [Reg05] formalised this relationship by giving a reduction from worst-case lattice problems to LWE (for certain parameter choices):

Theorem: If there exists an efficient algorithm that solves LWE then there exists an efficient algorithm that approximates the decision version of the shortest vector problem.

Regev's original paper only gave a quantum reduction, but Peikert later made his algorithm classical at the expense of more restrictive parameter choices [Pei09].

For the above theorem to work, it is necessary that χ is chosen to be a discrete multivariate Gaussian distribution. This means that sampling from LWE involves taking a lattice point and perturbing it by a small, normally distributed quantity, the idea being that this will look close enough to a uniform distribution if the standard deviation is large enough. Sampling from this discrete Gaussian is simply accomplished by sampling each component from a normal distribution and rounding to the nearest integer. The standard deviation is chosen to be $\alpha q / \sqrt{2\pi}$ for some $\alpha > 0$, and the mean is 0.

Note that in modern usage (e.g. the Brakerski/Vaikuntanathan homomorphic encryption schemes), it is required that the problem is hard when $\mathbf{s}$ is also sampled from the distribution χ, rather than just uniformly. It may be useful for these schemes to instead choose $\mathbf{s}$ from $\{-1, 0, 1\}^n$, in order to simplify the bootstrapping procedure that can be needed. However, the security of the corresponding LWE problem in this setting has not been studied.

2 The LWE cryptosystem

Here we present the original public-key cryptosystem based on LWE, as described by Regev. The parameters are a modulus q, the dimensions m, n, an integer l and the error distribution χ.
Secret key: Choose $S \in \mathbb{Z}_q^{n \times l}$ uniformly at random.
Public key: Choose $A \in \mathbb{Z}_q^{m \times n}$, $E \in \chi^{m \times l}$. The public key is given by the pair $(A, P := AS + E)$.
Encryption: To encrypt a message $m \in \mathbb{Z}_t^l$, first choose $a \in \{-r, \dots, r\}$. Then output the ciphertext

$$c = (A^T a, P^T a + \lceil m\frac{q}{t} \rfloor)$$

Decryption: We first compute the value

$$\begin{aligned}\mathbf{d} &= c_1 - S^T c_0 \\ &= P^T \mathbf{a} + \lceil \mathbf{m}\frac{q}{t} \rfloor - S^T A^T \mathbf{a} \\ &= E^T \mathbf{a} + \lceil \mathbf{m}\frac{q}{t} \rfloor.\end{aligned}$$

Since E is chosen from the relatively small noise distribution, we can recover the bits of $\mathbf{m}$ by simply rounding $\frac{t}{q}\mathbf{d}$ to the nearest integer.

Under the LWE assumption, the public key and any ciphertexts are all indistinguishable from random, and so the scheme can be shown to be IND-CPA secure.

2.1 Choosing the parameters

To choose parameters for the scheme, it helps to think about the LWE problem in terms of lattices. That is, to determine whether a given point $\mathbf{v}$ is in the lattice $\Lambda_q(A^T)$ or is uniformly random. A possible approach for this problem is to find a short vector $\mathbf{w} \in \Lambda_q(A^T)^*$. If $\mathbf{v}$ is random then $\langle \mathbf{v}, \mathbf{w} \rangle$ will also be random, whereas if $\mathbf{v} = A\mathbf{s} + \mathbf{e}$ then $\langle \mathbf{v}, \mathbf{w} \rangle$ will be close to an integer. To thwart this attack, we only need to ensure that the noise added in the direction of $\mathbf{w}$ much bigger than $1/\mathbf{w}$. Since the standard deviation of this noise is given by $\frac{\alpha q}{\sqrt{2\pi}}$, Regev suggests that to ensure security it is sufficient that

$$\frac{\alpha q}{\sqrt{2\pi}} > 1.5 \frac{1}{\|\mathbf{w}\|}.$$

If the BKZ algorithm is used to find this short vector $\mathbf{w}$ then its length is predicted to be $\min\{1, \delta^m q^{n/m-1}\}$, where δ is the Hermite root factor discussed previously (generally taken to be 1.005 or 1.01). So this leads us to the constraint:

$$\frac{\alpha q}{\sqrt{2\pi}} > 1.5 \max\{1, \delta^{-m} q^{1-n/m}\}.$$

Since an attacker has the capability to generate as many LWE instances as they wish (by simply encrypting messages), they effectively have complete freedom over the choice of the parameter m. This is therefore taken to be the optimal value for lattice-based attacks on LWE, which is given by $m = \sqrt{n \log q / \log \delta}$.

3 Ring LWE

The *ring-LWE* problem is an algebraic variant of LWE, which operates over elements of polynomial rings instead of vectors. A polynomial $f(x) \in \mathbb{Z}[x]$ is chosen, defining the ring $R = \mathbb{Z}[x]/f$. We then choose $a, s \in \mathbb{Z}[x]$ uniformly at random and an integer modulus q. The problem is to distinguish between random elements of $\mathbb{Z}_q[x]$ and elements of the form $v = as + e \mod q, f(x)$, where e is chosen from a small 'error' distribution χ.

This is an instance of the standard LWE problem, since polynomials in the ring R can be represented by matrices, with multiplication given by multiplication of a matrix and a vector, as in LWE. The security of this problem has been related to hardness problems on *ideal* lattices, rather than ordinary lattices, and it is unknown whether these problems could be any easier than standard LWE. However, ring-LWE is still an attractive alternative, as its algebraice structure implies that operations can be far more efficient than those in LWE.

References

[Pei09] Chris Peikert. Public-Key Cryptosystems from the Worst-Case Shortest Vector Problem. In *Proceedings of the 41st annual ACM symposium on Theory of computing*, pages 333–342. ACM, 2009.

[Reg05] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In *Proceedings of the 37th annual ACM symposium on Theory of Computing*, pages 84–93. ACM, 2005.