Homomorphic Encryption and Lattices, Spring 2011	Instructor: Shai Halevi

Quadratic-Homomorphic Encryption from LWE

May 5, 2011 — Scribe: Ron Rothblum

We present the GHV encryption scheme [GHV10]. This scheme, based on the hardness of learning with errors (LWE), supports homomorphic operations that can be expressed as quadratic forms (similarly to the BGN cryptosystem [BGN05]).

1 Background

The decision-LWE problem. The D-LWE$[n, \alpha, q]$ assumption asserts that it is infeasible to distinguish the distribution $\mathsf{LWE}^*_{\vec{s}} = \{(\vec{a}, c) : \vec{a} \in_R \mathbb{Z}_q^n,\ e \leftarrow \mathcal{N}(0, \alpha q),\ c = \langle \vec{s}, \vec{a} \rangle + e \bmod q\}$ for a random $s \in_R \mathbb{Z}_q^n$ from the uniform distribution on $\mathbb{Z}_q^n \times [0, q)$, even when the distinguisher can get any (polynomial) number of samples from these distributions that it wants. This implies that it is also infeasible to distinguish $\mathsf{LWE}_{\vec{s}} = \{(\vec{a}, c) : \vec{a} \in_R \mathbb{Z}_q^n,\ e \leftarrow \mathcal{N}(0, \alpha q),\ c = \langle \vec{s}, \vec{a} \rangle + \lceil e \rfloor \bmod q\}$ from uniform on $\mathbb{Z}_q^{n+1}$.

In particular, for any polynomial $m = m(n)$, the distribution

$$\mathsf{LWE}[m] = \{(A, \vec{c})\ :\ A \in_R \mathbb{Z}_q^{n \times m},\ s \in_R \mathbb{Z}_q^n,\ \vec{e} \leftarrow \mathcal{N}(0, \alpha q)^m,\ \vec{c} = \vec{s}A + \lceil \vec{e} \rfloor \bmod q\}$$

is indistinguishable from the uniform distribution on $\mathbb{Z}_q^{(n+1)\times m}$. By an easy hybrid argument, we get that the distribution

$$\mathsf{LWE}[m \times m] = \{(A, C)\ :\ A \in_R \mathbb{Z}_q^{n \times m},\ S \in_R \mathbb{Z}_q^{m \times n},\ E \leftarrow \mathcal{N}(0, \alpha q)^{m \times m},\ C = SA + \lceil E \rfloor \bmod q\}$$

is indistinguishable from the uniform distribution on $\mathbb{Z}_q^{(n+m)\times m}$.

Trapdoors. On the other hand, the trapdoor constructions (e.g., [AP11] or [MP11]) let us generate a nearly-uniform matrix $A \in \mathbb{Z}_q^{n \times m}$ together with a trapdoor T_A such that given T_A we can invert the function

$$\mathbf{lwe}_A(\vec{s}, \vec{e}) = \vec{s}A + \vec{e} \bmod q$$

where $\vec{s} \in \mathbb{Z}_q^n$, $\vec{e} \in \mathbb{Z}^m$, and $|\vec{e}|_\infty < q/8m$ (say).

In particular, the Alwen-Peikert trapdoor from [AP11] is a full-rank integer matrix T such that $AT = 0 \bmod q$ and all the entries in T are at most 3 in absolute value. Hence $(\vec{s}A + \vec{e}) \times T = \vec{e} \times T \pmod q$, but $|\vec{e} \times T|_\infty \le |\vec{e}|_\infty \times |T|_\infty \times m \le \frac{q}{8m} \times 3 \times m < q/2$. This means that $((\vec{s}A + \vec{e}) \times T \bmod q) = \vec{e} \times T$ over the integers, so

$$((\vec{s}A + \vec{e}) \times T \bmod q) \times T^{-1}\ =\ (\vec{e} \times T) \times T^{-1}\ =\ \vec{e}.$$

2 The Gentry-Halevi-Vaikuntanathan Cryptosystem

Key-generation. Run the Alwen-Peikert trapdoor construction to get $A \in \mathbb{Z}_q^{n \times m}$ and the corresponding trapdoor T_A. The public key is A and the secret key is T_A.

Encryption$_A(B)$. The plaintext is a binary matrix $B \in \{0, 1\}^{m \times m}$.
1. Choose at random $S \in_R \mathbb{Z}_q^{m \times n}$ and $E \leftarrow \mathcal{N}(0, \alpha q)^{m \times m}$;
2. The ciphertext is a matrix over $\mathbb{Z}_q^{m \times m}$, $C = SA + 2\lceil E \rfloor + B \bmod q$.

Decryption$_{T_A}(C)$**.** Note that each row of C is of the form $\vec{c_i} = \vec{s_i}A + (2\lceil\vec{e_i}\rfloor + \vec{b_i}) \bmod q$. Use the trapdoor T_A to recover the "error vector" $\vec{x_1} = (2\lceil\vec{e_i}\rfloor + \vec{b_i})$, then reduce modulo 2 to get $\vec{b_i}$.

2.1 Correctness

If $\alpha \le 1/nm$ (say), then the probability of having any entry in $\vec{e}$ larger than $q/17m$ in absolute value is bounded by some $\exp(-n)$. Therefore the "error-vectors" $\vec{x} = (2\lceil\vec{e_i}\rfloor + \vec{b_i})$ satisfy $|\vec{x_i}|_\infty < q/8m$, and so we can recover it using the trapdoor.

Below we will need also a stronger bound: For any parameters k, m, q and α and any fixed unit vector $\vec{u} \in \mathbb{R}^m$, when we choose $\vec{e} \leftarrow \mathcal{N}(0, \alpha q)^m$, then the probability that $|\langle\vec{u}, \vec{e}\rangle| > \alpha q \cdot k$ is bounded by $\exp(-k^2/2)$.

2.2 Security

We show that when q is odd, then a successful chosen-plaintext attacker $\mathcal{A}$ against the scheme implies a distinguisher $\mathcal{D}$ between $\mathsf{LWE}[m \times m]$ and uniform.

The distinguisher gets (A, C) and it needs to decide if $C = SA + E \bmod q$ or C is uniform in $\mathbb{Z}_q^{m\times m}$. It runs the attacker $\mathcal{A}$ with public key A, and the attacker gives it two matrices B_0, B_1. Then $\mathcal{D}$ chooses at random $i \in_R \{0, 1\}$ and provides the attacker $\mathcal{A}$ with the "ciphertext" $C^* = 2C + B_i$. Then $\mathcal{A}$ outputs a guess i', if $i' = i$ then $\mathcal{D}$ outputs 1 (i.e., it guesses that the input distribution is $\mathsf{LWE}[m \times m]$), and otherwise it outputs 0 (i.e., it guesses that the distribution is uniform).

If (A, C) is taken from the uniform distribution then C^* is uniform (since q is odd), regardless of i, hence the probability of $i' = i$ is exactly 1/2.

If (A, C) is taken from $\mathsf{LWE}[m \times m]$ then $C = SA + E \bmod q$ and therefore $C^* = 2C + B_i = (2S)A + 2E + B_i \bmod q$. Since q is odd and S is uniform over $\mathbb{Z}_q$ then so is $2S \bmod q$, hence C^* is distributed exactly the same as a random encryption of B_i. It follows that in this case we have $i' = i$ with probability noticeably larger than 1/2.

2.3 Additive Homomorphism

Assume that we set $\alpha \le 1/mk$ for some parameter k, and consider a set of ℓ plaintext matrices $B_1, \ldots, B_\ell$ and their encryption $C_1, \ldots, C_\ell$, where $\ell \le o(k^2/\sqrt{logn})$. We claim that with overwhelming probability, the matrix $\sum_{i=1}^{\ell} C_i \bmod q$ will be decrypted to the binary sum $\sum_{i=1}^{\ell} B_i \bmod 2$. This is because

$$\sum_{i=1}^{\ell} C_i = \underbrace{\left(\textstyle\sum_i S_i\right)}_{S} A + 2\underbrace{\left(\textstyle\sum_i E_i\right)}_{E} + \underbrace{\left(\textstyle\sum_i B_i\right)}_{B} \;=\; SA + 2E + B \pmod{q}$$

and since each entry in E is a sum of ℓ independent Gaussians with variance $(\alpha q)^2$, then each such entry is itself a Gaussian with variance $\ell(\alpha q)^2$. From $\alpha \le 1/mk$ and $\ell \le o(k^2/\sqrt{logn})$ it follows that with overwhelming probability each entry in E is $o(q/m)$ and in particular smaller than $q/16m$, as needed for our trapdoor to work.

2.4 Multiplicative Homomorphism

Let $C_1 = S_1A + 2E_1 + B_1 \bmod q$ and $C_2 = S_2A + 2E_2 + B_2 \bmod q$, and let $C = C_1 \times C_2^t \bmod q$. Then $TCT^t = T(2E_1 + B_1) \times (2E_2 + B_2^t)T^t \pmod{q}$. If α is chosen small enough so that all the entries in E_1, E_2 are $o(\sqrt{q}/m^{1.5})$, then all the entries in $T(2E_1 + B_1)$ and $T(2E_2 + B_2)$ are smaller than than $o(\sqrt{q/m})$, and so all the entries in $T(2E_1+B_1)\times(2E_2+B_2^t)T^t$ are smaller than $m\times o(\sqrt{q/m})\times o(\sqrt{q/m})=o(q)$. Therefore

$$TCT^t \bmod q \;=\; T(2E_1+B_1)\times(2E_2+B_2^t)T^t$$

over the integers, and so we get

$$T^{-1}(TCT^t \bmod q)(T^{-1})^t \;=\; (2E_1+B_1)\times(2E_2+B_2^t) \;=\; B_1B_2^t \pmod 2$$

We can therefore multiply two ciphertext matrices, and be able to decrypt the product of the two plaintext binary matrices from the resulting product ciphertext.

References

[AP11] Joël Alwen and Chris Peikert. Generating shorter bases for hard random lattices. *Theory Comput. Syst.*, 48(3):535–553, 2011.

[BGN05] Dan Boneh, Eu-Jin Goh, and Kobbi Nissim. Evaluating 2-DNF formulas on ciphertexts. In *Theory of Cryptography - TCC'05*, volume 3378 of *Lecture Notes in Computer Science*, pages 325–341. Springer, 2005.

[GHV10] Craig Gentry, Shai Halevi, and Vinod Vaikuntanathan. A Simple BGN-type Cryptosystem from LWE. In *Advances in Cryptology - EUROCRYPT'10*, volume 6110 of *Lecture Notes in Computer Science*, pages 506–522. Springer, 2010. Full version available on-line from `http://eprint.iacr.org/2010/145`.

[MP11] Daniele Micciancio and Chris Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. Manuscript, 2011.