

Let's revisit Regan's encryption scheme. It turns out that it readily generalizes to give a fully homomorphic encryption scheme.

Abstractly: given encryption ct_x of value x under some public key, can we derive from that an encryption of $f(x)$ for an arbitrary function f ?

- So far, we have seen examples of encryption schemes that support one type of operation (e.g., addition) on ciphertexts

- ElGamal encryption (in the exponent): homomorphic with respect to addition

- Boneh-Goh-Nissim: addition + 1 multiplication

- For FHE, need homomorphism with respect to two operations: addition and multiplication

Major open problem in cryptography (dates back to late 1970s!) - first solved by Stanford student Craig Gentry in 2009

↳ revolutionized lattice-based cryptography!

↳ Very surprising this is possible: encryption needs to "scramble" messages to be secure, but homomorphism requires preserving structure to enable arbitrary computation

General blueprint: 1. Build somewhat homomorphic encryption (SWHE) - encryption scheme that supports bounded number of homomorphic operations

2. Bootstrap SWHE to FHE (essentially a way to "refresh" ciphertext)

Focus will be on building SWHE (has all of the ingredients for realizing FHE)

Starting point: Regan encryption

$$\text{pk: } A = \begin{bmatrix} \bar{A} \\ s^T \bar{A} + e^T \end{bmatrix} \in \mathbb{Z}_q^{n \times m} \quad \left. \vphantom{\begin{bmatrix} \bar{A} \\ s^T \bar{A} + e^T \end{bmatrix}} \right\} \text{Invariant: } s^T A = e^T$$

$$\text{sk: } s^T = \begin{bmatrix} -\bar{s}^T & | & 1 \end{bmatrix} \in \mathbb{Z}_q^n$$

$$\text{ct: } r \xleftarrow{R} \{0,1\}^m, c \leftarrow Ar + \begin{bmatrix} 0^{n-1} \\ \lfloor \frac{q}{2} \rfloor \cdot \mu \end{bmatrix} \in \mathbb{Z}_q^n$$

$$\hookrightarrow s^T c = s^T (Ar + \begin{bmatrix} 0^{n-1} \\ \lfloor \frac{q}{2} \rfloor \cdot \mu \end{bmatrix}) = e^T r + \lfloor \frac{q}{2} \rfloor \cdot \mu$$

as long as $e^T r$ is small, decryption succeeds

Essentially, with Regan encryption, the decryption invariant is

$$s^T c = \mu \cdot \lfloor \frac{q}{2} \rfloor + \text{error}$$

Suppose however that instead of encrypting μ , we encrypted the entries of $\mu \cdot s^T$ instead. And also ignore the scaling factor.

Then, the ciphertext would be a matrix $C \in \mathbb{Z}_q^{n \times n}$ where

$$s^T C = \mu \cdot s^T + \text{error} \in \mathbb{Z}_q^n$$

$$\uparrow \text{specifically: } C = AR + \mu \cdot I_n$$

$$\text{where } R \xleftarrow{R} \{0,1\}^{m \times n}$$

$$\begin{aligned} \xrightarrow{s^T A = e^T} s^T C &= s^T AR + \mu \cdot s^T \\ &= \underbrace{e^T R}_{\text{error}} + \mu \cdot s^T \end{aligned}$$

Observe: Suppose C_1 was a Regan encryption of $\mu_1 \cdot s^T$ and C_2 was Regan encryption of $\mu_2 \cdot s^T$. Then:

$$s^T C_1 C_2 = (\mu_1 \cdot s^T + e_1^T) C_2 = \mu_1 (\mu_2 \cdot s^T + e_2^T) + e_1^T C_2$$

$$= \mu_1 \mu_2 \cdot s^T + \mu_1 e_2^T + e_1^T C_2$$

This is basically an encryption of μ_1, μ_2 with new error term $\mu_1 e_2^T + e_1^T C_2$.

small since $\mu_1 \in \{0,1\}$ and e_2^T is small
 big because C_2 is a Regev ciphertext (has large entries over $\mathbb{Z}_q^{n \times n}$)

Due to the large noise, cannot recover the message anymore...

Need a way to avoid multiplying by something large.

- How to make something small? Binary decomposition!

First, we define the "gadget" matrix (there are actually many possible gadget matrices - here, we use a common one sometimes called the "powers-of-two" matrix):

$$G = \begin{pmatrix} 1 & 2 & 4 & 8 & \dots & 2^{\lceil \log_2 b \rceil - 1} & & & \\ & 1 & 2 & 4 & \dots & 2^{\lceil \log_2 b \rceil - 1} & & & \\ & & & & & & \ddots & & \\ & & & & & & & 1 & 2 & 4 & \dots & 2^{\lceil \log_2 b \rceil - 1} \end{pmatrix} = \underbrace{\begin{pmatrix} 1 & 2 & 4 & \dots & 2^{\lceil \log_2 b \rceil - 1} \end{pmatrix}}_{g^T} \otimes I_n = g^T \otimes I_n$$

Each row of G consists of the powers of two (up to $2^{\lceil \log_2 b \rceil - 1}$). Thus, $G \in \mathbb{Z}_q^{n \times n \lceil \log_2 b \rceil}$. Oftentimes, we will just write $G \in \mathbb{Z}_q^{n \times m}$ where $m > n \lceil \log_2 b \rceil$. Note that we can always pad G with all-zero columns to obtain the desired dimension.

Observation: given any $y \in \mathbb{Z}_q^n$, it is easy to find an $x \in \{0,1\}^m$ where $Gx = y$.

Let $y_{i, \lceil \log_2 b \rceil - 1}, \dots, y_{i,0}$ be the binary decomposition of y_i (the i 'th component of y). Then,

$$G \cdot \begin{pmatrix} y_{1,0} \\ y_{1,2} \\ \vdots \\ y_{1, \lceil \log_2 b \rceil - 1} \\ y_{2,0} \\ \vdots \\ y_{2, \lceil \log_2 b \rceil - 1} \\ \vdots \\ y_{n,0} \\ \vdots \\ y_{n, \lceil \log_2 b \rceil - 1} \end{pmatrix} = \begin{pmatrix} \sum_{j=0}^{\lceil \log_2 b \rceil - 1} 2^j y_{1,j} \\ \vdots \\ \sum_{j=0}^{\lceil \log_2 b \rceil - 1} 2^j y_{n,j} \end{pmatrix} = \begin{pmatrix} y_1 \\ y_2 \\ \vdots \\ y_n \end{pmatrix} = y$$

↑ Observe that this is a 0/1 vector (binary valued vector)

We will denote this "bit-decomposition" operation by the function $G^{-1}: \mathbb{Z}_q^n \rightarrow \{0,1\}^m$

↑ important: G^{-1} is not a matrix (even though G is)!

Then, for all $y \in \mathbb{Z}_q^n$, $G \cdot G^{-1}(y) = y$ and $\|G^{-1}(y)\| = 1$.

↑ ℓ_∞ -norm (max absolute value of component of the vector)

Returning to FHE:

Approach: instead of encrypting $\mu \cdot s^T$, we will encrypt $\mu \cdot s^T G$ instead.

Invariant: C is an encryption of μ if
$$s^T C = \mu \cdot s^T G \in \mathbb{Z}_q^m$$

We can construct C as

$$C = AR + \mu G \in \mathbb{Z}_q^{n \times m}$$

$$\text{Then } s^T C = s^T AR + \mu \cdot s^T G = e^T R + \mu \cdot s^T G$$

Suppose we have two ciphertexts C_1 and C_2 where

$$s^T C_1 = \mu_1 \cdot s^T G + e_1^T$$

$$s^T C_2 = \mu_2 \cdot s^T G + e_2^T$$

Then $C_1 + C_2$ is an encryption of $\mu_1 + \mu_2$:

$$s^T (C_1 + C_2) = (\mu_1 + \mu_2) \cdot s^T G + e_1^T + e_2^T \quad [\text{errors add}]$$

To multiply, we compute $C_1 G^{-1}(C_2)$:

$$\begin{aligned} s^T C_1 G^{-1}(C_2) &= (\mu_1 \cdot s^T G + e_1^T) G^{-1}(C_2) \\ &= \mu_1 \cdot s^T C_2 + e_1^T G^{-1}(C_2) \\ &= \mu_1 \mu_2 \cdot s^T G + \underbrace{\mu_1 e_2^T + e_1^T G^{-1}(C_2)} \end{aligned}$$

small since $\mu_1 e_2^T + e_1^T G^{-1}(C_2)$ is also small

To decrypt a ciphertext C , can compute $s^T C \cdot G^{-1}(\lfloor \frac{q}{2} \rfloor \cdot u_n)$ where $u_n = \begin{bmatrix} 0 \\ \vdots \\ 1 \end{bmatrix}$ since $s^T u_n = 1$.

As long as total error is less than $\frac{q}{4}$, decryption recovers message

This gives the Gentry-Sahai-Waters encryption scheme.

- Setup (1^λ): Sample $\bar{A} \xleftarrow{R} \mathbb{Z}_q^{(n-1) \times m}$ \rightarrow $pk = A = \begin{bmatrix} \bar{A} \\ s^T \bar{A} + e^T \end{bmatrix}$ ($s^T A = e^T$)
 $\bar{s} \xleftarrow{R} \mathbb{Z}_q^{n-1}$
 $e \leftarrow \chi^m$ $sk = s = [-\bar{s} \mid 1]$

- Encrypt (A, μ): $R \xleftarrow{R} \{0,1\}^{m \times n}$
 $C \leftarrow AR + \mu \cdot G \in \mathbb{Z}_q^{n \times m}$

- Decrypt (s, C): compute $s^T C G^{-1}(\frac{q}{2} \cdot I_n)$ and round as usual

Security is same argument as for Regw encryption:

Namely, by LWE, the public key is indistinguishable from a uniformly random matrix $A \xleftarrow{R} \mathbb{Z}_q^{n \times m}$
by LHL, (A, AR) is indistinguishable from (A, U) where $U \xleftarrow{R} \mathbb{Z}_q^{n \times m}$

$\Rightarrow U + \mu G$ perfectly hides μ .

Let's look at noise growth. Suppose $C_1 = AR_1 + \mu_1 G$

$$C_2 = AR_2 + \mu_2 G$$

$$\text{Then } s^T C_1 = s^T AR_1 + \mu_1 s^T G = \mu_1 s^T G + e^T R_1$$

← noise in the ciphertext: must be small relative to g in order to decrypt

Noise increases with each operation:

$$C_1 + C_2 = A(R_1 + R_2) + (\mu_1 + \mu_2)G \rightsquigarrow \text{new noise is } R_1 + R_2$$

$$C_1 G^{-1}(C_2) = AR_1 G^{-1}(C_2) + \mu_1 C_2$$

$$= A(R_1 G^{-1}(C_2) + \mu_1 R_2) + \mu_1 \mu_2 G \rightsquigarrow \text{new noise is } R_1 G^{-1}(C_2) + \mu_1 R_2$$

norm is bounded by $\|R_1\|_{\infty} \cdot m + \|R_2\|_{\infty}$ when $\mu_1, \mu_2 \in \{0, 1\}$.

After computing d repeated squarings: noise is $m^{O(d)}$. Will eventually overwhelm g . Thus, there is a bound on number of homomorphic operations the scheme supports.

Fully homomorphic encryption: support arbitrary number of computations.

From SWHE to FHE. The above construction requires imposing an a priori bound on the multiplicative depth of the computation.

To obtain fully homomorphic encryption, we apply Gentry's brilliant insight of bootstrapping.

High-level idea. Suppose we have SWHE with following properties:

1. We can evaluate functions with multiplicative depth d
2. The decryption function can be implemented by a circuit with multiplicative depth $d' < d$

Then, we can build an FHE scheme as follows:

- Public key of FHE scheme is public key of SWHE scheme and an encryption of the SWHE decryption key under the SWHE public key
- We now describe a ciphertext-refreshing procedure:
 - For each SWHE ciphertext, we can associate a "noise" level that keeps track of how many more homomorphic operations can be performed on the ciphertext (while maintaining correctness).
 - ↳ for instance, we can evaluate depth- d circuits on fresh ciphertexts; after evaluating a single multiplication, we can only evaluate circuits of depth- $(d-1)$ and so on ...
 - The refresh procedure takes any valid ciphertext and produces one that supports depth- $(d-d')$ homomorphism; since $d > d'$, this enables unbounded (i.e., arbitrary) computations on ciphertexts

Idea: Suppose we have a ciphertext ct where $\text{Decrypt}(sk, ct) = x$.

To refresh the ciphertext, we define the Boolean circuit $C_{ct} : \{0, 1\}^{n \log b} \rightarrow \{0, 1\}$ where $C_{ct}(sk) := \text{Decrypt}(sk, ct)$

and homomorphically evaluate C_{ct} on the encryption of sk

$$\text{Encrypt}(pk, sk) \rightarrow \text{Encrypt}(pk, C_{ct}(sk))$$

↑ fresh ciphertext that supports d levels

↑ homomorphic evaluation consumes d' levels

↑ x ← refreshed ciphertext still supports $d-d'$ levels of multiplication

Security now requires that the public key includes a copy of the decryption key

↳ Requires making a "circular security" assumption

Open question: FHE without circular security from LWE (possible from iO)

Can be shown that GSW is bootstrappable. [Decryption operation is linear, followed by rounding - can be implemented with low-depth circuit]