

Lecture 16:

Fully Homomorphic Encryption

CS 355 - Spring 2019

May 22, 2019

Henry Corrigan-Gibbs



Logistics

* HW4 Due Friday, May 24 at 5pm

* Vote for topics of bonus lectures

* HWS out now.

- * Events:
- Grubbs talk (5/24, 4:15pm in Gates 463A)
 - Aloni Cohen (5/29, 12pm in Gates 463A) *
 - Defense (5/31, 1pm in Packard 101)
 - David Lazar (6/4, 4:15pm in Gates 463A)
 - Michael Duff (6/5, 12pm in Gates 463A) *

* = Lunch!

Plan

* Reminder of LWE

* Defn of FHE

* Leveled HE (GSW)
↳ Construction

* Bootstrapping to FHE

Recap: Learning w/ Errors

NEW ASSUMPTION

LWE (n, m, q, χ_s) ... see diagram on next page

$$\left\{ (A, s^T A + e^T) \mid \begin{array}{l} A \leftarrow \mathbb{Z}_q^{n \times m} \\ s \leftarrow \mathbb{Z}_q^n \\ e \leftarrow \chi_B \end{array} \right\} \stackrel{\sim}{=} \left\{ u \mid u \leftarrow \mathbb{Z}_q^{(n+1) \times m} \right\}$$

Solving noisy overdetermined system of eqns is hard.

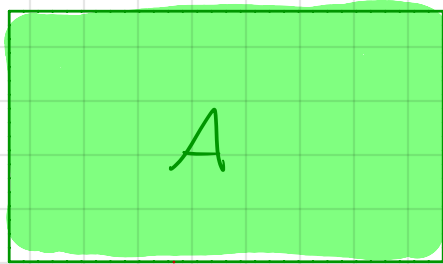
Reyer Encryption

- PKE w/ sem. sec under LWE assumption
- Same ideas used to construct FHE scheme (today).

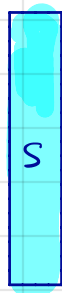
Recap : LWE

LWE assumption (params n, m, q, χ)

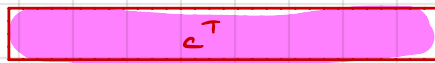
for



$$\leftarrow \sum_q^{n \times m} \mathcal{R}$$



$$\leftarrow \sum_q^n \mathcal{R}$$



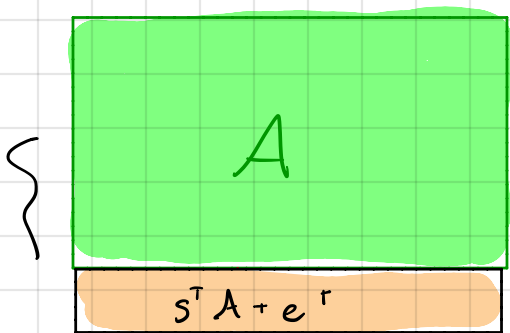
$$\leftarrow \chi \mathcal{R}$$

$$\lambda = 128$$

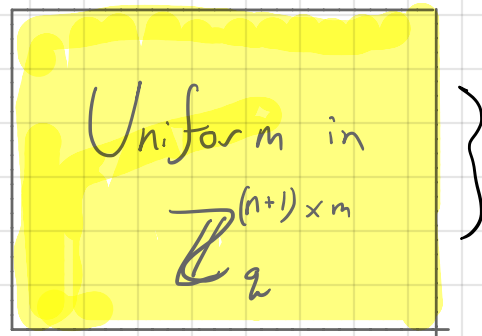
$$n = 600 - 800$$

$$m = n \log_2 q$$

$$q = n^2 \sim 2^{\sqrt{n}}$$

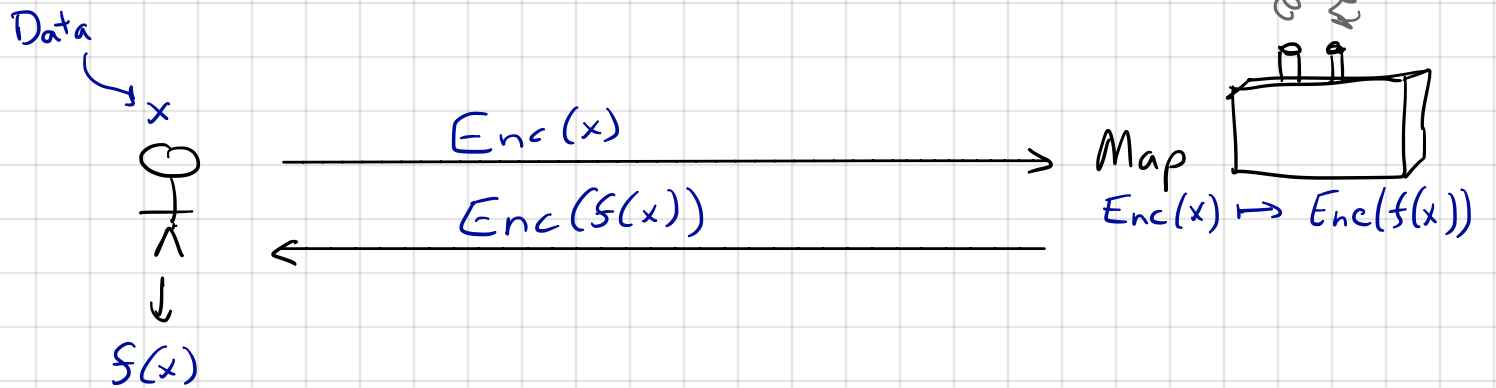


$$\approx_c$$



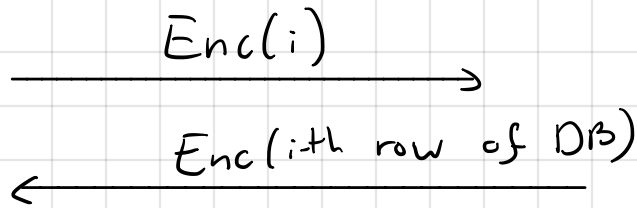
Fully Homomorphic Encryption (FHE)

Idea: Outsource computation without revealing inputs.

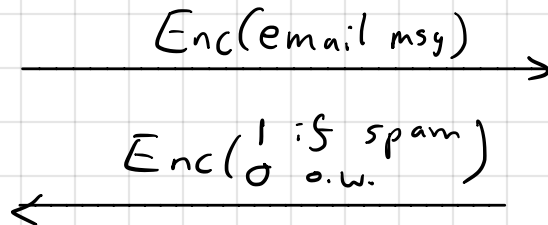


Examples:

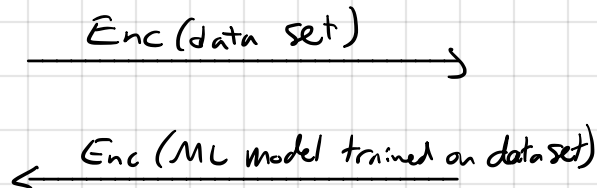
* Private lookup (PIR)



* Private spam filtering



* Private ML



* And so on...

N.B. FHE alone isn't powerful enough to do everything we'd like to do (e.g. secure comp on joint inputs), but it's still very powerful.

N.B. These schemes aren't used in practice yet...

Some history...

1978 - Rivest, Adleman, DeMillo (1978) introduce notion of FHE (under other name)

↳ We've had partially hom. enc. for a long time (RSA, Paillier, ElGamal, ...)

↳ No unbroken candidates

2009 - Craig Gentry (Stanford PhD student in Dan's group) gives first candidate FHE scheme

I think it's fair to call this a major breakthrough

↳ From new assumption (non-standard)

↳ Introduced crucial "bootstrapping" idea

2011 - Schemes based on LWE (Brakerski, Vaikuntanathan)
+ "circular-security" assumption

2013 - "Third-gen scheme" that we'll see today
(Gentry, Sahai, Waters)

Syntax of FHE (Based on notes from Sam Kim.)

$$\text{KeyGen}(1^\lambda) \rightarrow sk$$

$$\text{Enc}(sk, \mu) \rightarrow ct$$

$$\text{Eval}(C, ct_1, \dots, ct_e) \rightarrow \tilde{ct}$$

$$\text{Dec}(sk, ct) \rightarrow \mu$$

(Can define a PHE notion
of FHE as well...)

Note: The computation is represented by a boolean ckt. This is the source of some of the ineff. in practice, but is wlog wrt polynomial-time comps.

Properties

① Correctness. Let $C: \{0,1\}^{\ell} \rightarrow \{0,1\}$

$\forall sk$ output by KeyGen , $\forall C, \forall m_1, \dots, m_\ell$

$$\text{Dec}(sk, \text{Eval}(C, \text{Enc}(sk, m_1), \dots, \text{Enc}(sk, m_\ell))) = C(m_1, \dots, m_\ell).$$

② Semantic Security.

$\forall sk$ output by KeyGen , $\forall m_0, m_1$

$$\{\text{Enc}(sk, m_0)\} \stackrel{c}{\approx} \{\text{Enc}(sk, m_1)\}$$

③ Compactness. Dec. time is indep of $|C|$.

$\forall C, ct_1, \dots, ct_\ell$ output by $\text{Enc}(sk, \cdot)$

if $\tilde{ct} \leftarrow \text{Eval}(C, ct_1, \dots, ct_\ell)$

then $|\tilde{ct}| = \text{poly}(1)$ \leftarrow independent of $|C|, \ell$

Claim: w/o compactness, any sem. sec. enc scheme is also an FHE.

Idea behind FHE

All FHE schemes we have use a two-step approach

① Construct "leveled" HE

- ↳ Eval handles ccts of bounded depth.
- ↳ Intuition: There's some noise/error in cts that grows with each \times gate in ckt. Eventually, error gets too large and breaks correctness.

② Use bootstrapping to convert leveled HE \rightarrow FHE

- ↳ This is the mind-blowing idea from Gentry (2009) ... so simple, so nice

↳ Idea: Publish $ct_{sk} = \text{Enc}(pk, sk)$ as part of public info.

(Think of a public-key FHE...)

Compute $ct_{ct} \leftarrow \text{Enc}(pk, ct_{\mu})$

"noisy" encryption of μ

$\hat{ct} \leftarrow \text{Eval}(\text{Dec}, ct_{sk}, ct_{ct})$

"Clean" encryption of ct_{μ}

$\hat{ct} = \text{Enc}(pk, \text{Dec}(sk, ct_{ct}))$

Homomorphically decrypt the noisy ct_{μ}

$= \text{Enc}(sk, \mu)$

"clean encryption of μ "

Digression:

Two good ideas in CS.
Both from Turing.
FHE uses both...
st excellent...

\Rightarrow Can remove noise by bootstrapping.
Can compute arbitrary ccts on encrypted data.

Two comments about "bootstrapping"

1. Publishing $\text{Enc}(sk, sk)$ might be risky.

↳ Semantic security of Enc alone doesn't imply that this is safe.

(Ex. Construct a sem. sec. enc scheme for which publishing $\text{Enc}(sk, sk)$ leads to a total break of the scheme.)

When using bootstrapping we make a "circular security" assumption --- just assume that this is safe.

2. We need to run $\text{Eval}(\hat{\text{Dec}}, \dots)$

↳ Our SWHE scheme must support cts of depth $\geq \text{depth}(\hat{\text{Dec}})$.

↳ Want a SWHE scheme with a "shallow" decryption circuit.

Preliminaries

We'll use the "Gadget matrix" $G \in \mathbb{Z}_q^{n \times n \log q}$

$$G = \begin{pmatrix} 1 & 2 & 4 & \dots & 2^{q/2} & & & \\ & & & & & 1 & 2 & 4 & \dots & 2^{q/2} \\ & & & & & & & & & \dots \\ & & & & & & & & & & 1 & 2 & 4 & \dots & 2^{q/2} \end{pmatrix}$$

Bit decomposition

For a matrix $\tilde{C} \in \mathbb{Z}_q^{n \times n \log q}$, define the bit decomposition operation

$$C = \text{BitDecomp}(\tilde{C}) \in \mathbb{Z}_q^{m \times m}$$

that splits each \mathbb{Z}_q element x in \tilde{C} into a vector

$$x \in \mathbb{Z}_q \rightarrow \begin{pmatrix} x_1 \\ x_2 \\ \vdots \\ x_{\log q} \end{pmatrix} \in \{0, 1\}^{\log q \times 1} \subseteq \mathbb{Z}_q^{\log q \times 1}$$

Then:

$$\begin{array}{ccc} \begin{array}{c} m = n \log q \\ n \end{array} \boxed{\tilde{C}} & \Rightarrow & \begin{array}{c} m \\ \begin{array}{cc} c_{11} & c_{12} \\ c_{21} & \\ \vdots & \\ \end{array} \\ m \end{array} \boxed{C} \end{array}$$

Key Idea: $\tilde{C} = G \cdot \text{BitDecomp}(\tilde{C})$

Now, we construct leveled HE scheme...

GSW FHE (secret key)

← Shares ideas w/ Paillier encryption

$$\text{Key Gen}(1^\lambda): \quad \tilde{s} \leftarrow^R \mathbb{Z}_q^{n-1}$$
$$\text{sk} = \begin{pmatrix} \tilde{s} & -1 \end{pmatrix} \in \mathbb{Z}_q^n$$

$$\text{Encrypt}(s, \mu): \quad A \leftarrow^R \mathbb{Z}_q^{(n-1) \times n \log_2 q}$$
$$e \leftarrow^R \chi$$
$$\tilde{c} = \underbrace{\begin{pmatrix} A \\ \tilde{s}^T A + e^T \end{pmatrix}}_{\text{pseudorandom by LWE}} + \mu \cdot G \in \mathbb{Z}_q^{n \times n \log_2 q}$$

$$\text{output } C \leftarrow \text{BitDecompose}(\tilde{c}) \in \mathbb{Z}_q^{m \times m}$$

Decrypt (s, C) :

$$\text{Write } \tilde{s}^* \leftarrow s^T G \in \mathbb{Z}_q^{1 \times 2}$$

$$\text{Then } (\tilde{s}^*)^T C = s^T G \cdot C$$
$$= s^T \tilde{C}$$

$$= s^T \left[\begin{pmatrix} A \\ \tilde{s}^T A + e^T \end{pmatrix} + \mu \cdot G \right]$$

$$= (\tilde{s}^T - 1) \begin{pmatrix} A \\ \tilde{s}^T A + e^T \end{pmatrix} + \mu s^T G$$

$$= \tilde{s}^T A - (\tilde{s}^T A + e^T) + \mu s^T G$$

$$= \mu s^T G - \underbrace{e^T}_{\text{small noise}}$$

$$= \begin{cases} \text{"small"} & \text{if } \mu = 0 & (< \frac{q}{4}) \\ \text{"big"} & \text{if } \mu = 1 & (\geq \frac{q}{4}) \end{cases}$$

Now to compute hom. operations...

It's enough to show how to go from $C_{\mu_1}, C_{\mu_2} \mapsto C_{\mu_1 \text{ NAND } \mu_2}$
since NAND is a universal gate.

Homomorphic NAND:

$$C_{\mu_1}, C_{\mu_2} \mapsto \mathbb{I}_m - (C_{\mu_1} \cdot C_{\mu_2})$$

$$\begin{aligned} \text{Then } (s^*)^T [\mathbb{I}_m - C_{\mu_1} \cdot C_{\mu_2}] \\ &= s^T G - (s^T G C_{\mu_1}) \cdot C_{\mu_2} \\ &= s^T G - (s^T_{\mu_1} G + e_1^T) C_{\mu_2} \\ &= s^T G - \mu_1 s^T G C_{\mu_2} + e_1^T C_{\mu_2} \\ &= s^T G - \mu_1 (\mu_2 s^T G + e_2^T) + e_1^T C_{\mu_2} \\ &= s^T G - \mu_1 \mu_2 s^T G - \underbrace{\mu_1 e_2^T + e_1^T}_{\text{noise}} C_{\mu_2} \\ &= (1 - \mu_1 \mu_2) s^T G + \langle \text{noise} \rangle \end{aligned}$$

Small noise since $|\mu_i| \leq 1$
 $\|C_{\mu_2}\|_{\infty} \leq 1$

Large if $\mu_1 \text{ NAND } \mu_2 = 1$

Small if $\mu_1 \text{ NAND } \mu_2 = 0$

The catch...

With each homomorphic NAND we perform, the noise grows.

→ If we perform many NANDs, noise will overwhelm the "signal".

After computing a ckt of depth L , noise will be bounded by

$$(m+1)^{\text{depth}} \cdot B$$

↖ original noise bound

We need $(m+1)^{\text{depth}} \cdot B \ll q$

$$(m+1)^{\text{depth}} \ll q/B$$

$$\text{depth} \cdot \log(m+1) \ll \log(q/B)$$

$$\text{depth} \ll \log(q/B) / \log(m+1)$$

↖ If we take $q = 2^{\sqrt{n}}$, can handle ckt of \log depth.

↖ Regev-style encryption has a decryption ckt of $O(\log n)$ depth... it's just a matrix-vector product plus rounding.

↖ We can support bootstrapping, and we're done!

When will this be practical?