Improved Security for a Ring-Based Fully Homomorphic Encryption Scheme

Michael Naehrig
Cryptography Research Group
Microsoft Research

Joint work with Joppe W. Bos (MSR), Jake Loftus (University of Bristol), and Kristin Lauter (MSR)

PQCrypto 2013
Limoges, 5 June 2013

Fully Homomorphic Encryption (FHE)

Enables unlimited computation on encrypted data
Need scheme with unlimited add and mult capability

- Idea: Rivest, Adleman, Dertouzos (1978)
- Boneh, Goh, Nissim (2005): unlimited add + 1 mult
- Breakthrough: Gentry (2009) showed such schemes exist
- A lot of progress since then
- Gentry, Halevi, Smart (2012): homomorphic evaluation of AES
  5 minutes per block (16 bytes)

Totally and utterly impractical!

Totally impractical!

Homomorphic Encryption from RLWE

Encryption from RLWE

- RLWEencrypt (Lyubashevsky, Peikert, Regev 2010)
- secureNTRU (Stehlé, Steinfeld 2011)

Homomorphic encryption schemes from (R)LWE

- RLWE FHE: BV (Brakerski, Vaikuntanathan 2011)
- Leveled HE: BGV (Brakerski, Gentry, Vaikuntanathan 2012)
- Multi-key scheme from NTRU
  (López-Alt, Tromer, Vaikuntanathan 2012)
- Scale-invariant HE from LWE (Brakerski 2012)
- Scale-invariant HE from RLWE (Fan, Vercauteren 2012)

This talk

Rather theoretical result:
A fully homomorphic encryption scheme

- Based on secureNTRU
  with security based only on RLWE
  (and a circular security assumption)
- no need for the SPR assumption
  (from NTRU-based multi-key FHE)

This talk

More practical result:
A leveled homomorphic encryption scheme

- Based on NTRU
  with security based on RLWE
  and SPR assumption (as in NTRU-based multi-key FHE)
- Using “Regev-style” encryption [B12]
  i.e. scale invariant without modulus switching
- Ciphertexts have only one element (half the size of BGV)
- Parameters comparable to BGV

In this talk

there will be No Bootstrapping!
only leveled homomorphic encryption

In “practice”, one tries to avoid bootstrapping

A Ring R

Let Φ_d be the d-th cyclotomic polynomial for $d > 0$.

- Define
$$R = \mathbf{Z}[X]/(\Phi_d(X))$$
represented by the set of polynomials with integer coefficients of degree less than $n = \deg(\Phi_d) = \varphi(d)$

- $a = \sum_{i=0}^{n-1} a_i X^i \in R$, $\|a\|_\infty = \max_i\{|a_i|\}$
- For an integer modulus q let $R_q = R/qR$

For example: $d = 2^k, n = \frac{\varphi(d)}{2} = 2^{k-1}, R = \mathbf{Z}[X]/(X^n + 1)$

A Discrete Noise Distribution χ

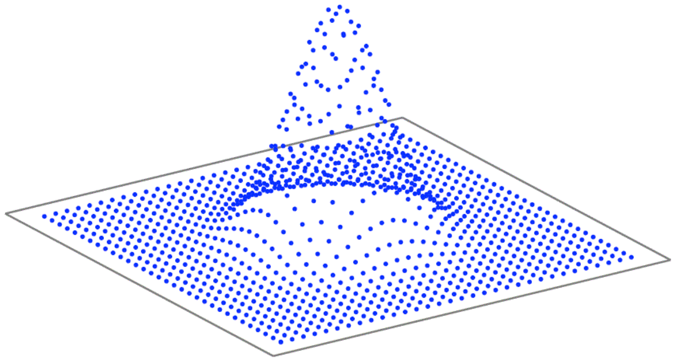

Let χ be a probability distribution on R
that samples small elements $a \leftarrow \chi$ with high probability
e.g. a discrete Gaussian distribution

- For example: If $d = 2^k, n = 2^{k-1}, R = \mathbf{Z}[X]/(X^n + 1)$, can take
$$\chi = D_{Z^n,\sigma}$$
- i.e. each coefficient is sampled independently from a one-dimensional discrete Gaussian with standard deviation σ

- probability proportional to $\exp(-\pi|x|^2/\sigma^2)$ for each $x \in \mathbf{Z}$

Ring Learning With Errors (RLWE)

(Lyubashevsky, Peikert, Regev 2010)

Given the Ring R, modulus q, $R_q = R/qR$, and the probability distribution χ on R

Problem: distinguish between two distributions

1. Uniform distribution $(a, b) \in R_q^2$
2. The distribution that for a fixed $s \leftarrow \chi$ samples $a \leftarrow R_q$ uniformly, an error $e \leftarrow \chi$ and outputs $(a, a \cdot s + e)$

Assumption: The RLWE problem is hard, i.e.

$(a, a \cdot s + e) \sim (a, b)$ looks uniform random

(Symmetric) Encryption from RLWE

Message $m \in R/2R$
$s \leftarrow \chi$ secret key

BV (Brakerski, Vaikuntanathan 2011) encryption:

Sample $a \leftarrow R_q$ uniform, $e \leftarrow \chi$ error/noise
$\mathrm{b} = m + a \cdot s + 2e \bmod q$, ciphertext $\mathrm{c} = (\mathrm{a}, \mathrm{b})$

$b - a \cdot s = m + 2e \bmod q$
decrypt: $(b - a \cdot s) \bmod 2$
decrypts correctly if $\|e\|_\infty < \frac{q}{2}$

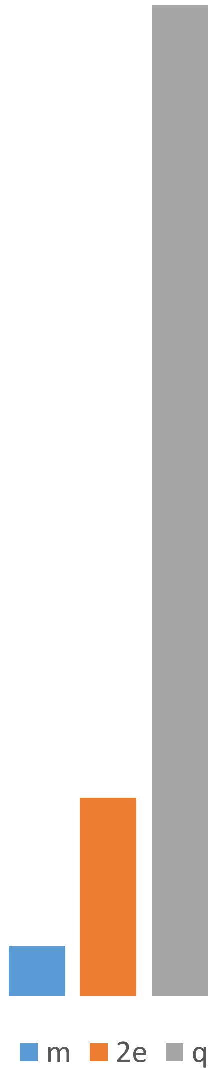

Homomorphic Addition

$$\mathrm{c}_1 = (a_1, b_1) = (\mathrm{a}_1, m_1 + a_1 \cdot s + 2e_1\ )$$
$$\mathrm{c}_2 = (a_2, b_2) = (\mathrm{a}_2, m_2 + a_2 \cdot s + 2e_2)$$

Addition:

$$\begin{aligned} \mathrm{c}_3 &= (\mathrm{a}_3, \mathrm{b}_3) \\ &= c_1 + c_2 = \left(a_1 + a_2, (m_1 + m_2) + (a_1 + a_2) \cdot s + 2(e_1 + e_2)\right) \end{aligned}$$

encrypts $(m_1 + m_2)\ \mathrm{mod}\ 2$, i.e. sum in R_2

Homomorphic Multiplication

$\mathrm{c}_1 = (a_1, b_1) = (\mathrm{a}_1, m_1 + a_1 \cdot s + 2e_1)$

$\mathrm{c}_2 = (a_2, b_2) = (\mathrm{a}_2, m_2 + a_2 \cdot s + 2e_2)$

Multiplication (BV):

$(b_1 - a_1 \cdot s)(b_2 - a_2 \cdot s) = (m_1 + 2e_1)(m_2 + 2e_2)$

$= m_1 m_2 + 2(m_1 e_2 + m_2 e_1 + 2e_1 e_2)$

$(b_1 - a_1 \cdot s)(b_2 - a_2 \cdot s) = b_1 b_2 - (b_1 a_2 + b_2 a_1)s + a_1 a_2 s^2$

New ciphertext: $\mathrm{c}_3 = (a_1 a_2, b_1 a_2 + b_2 a_1, b_1 b_2)$ now 3 elements!

Relinearization transforms it back to two elements (key switching)

Encrypts $(m_1 \cdot m_2)$ mod 2, i.e. product in R_2

Noise Growth

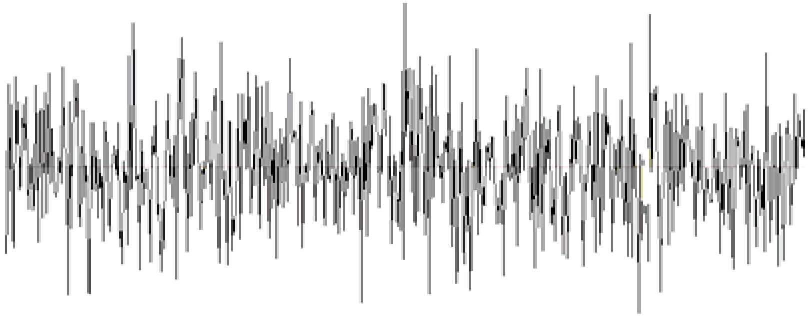

- Initial noise: B
- Addition: noise terms add up, $B \rightarrow 2B$
- Multiplication: noise terms are multiplied, $B \rightarrow B^2$

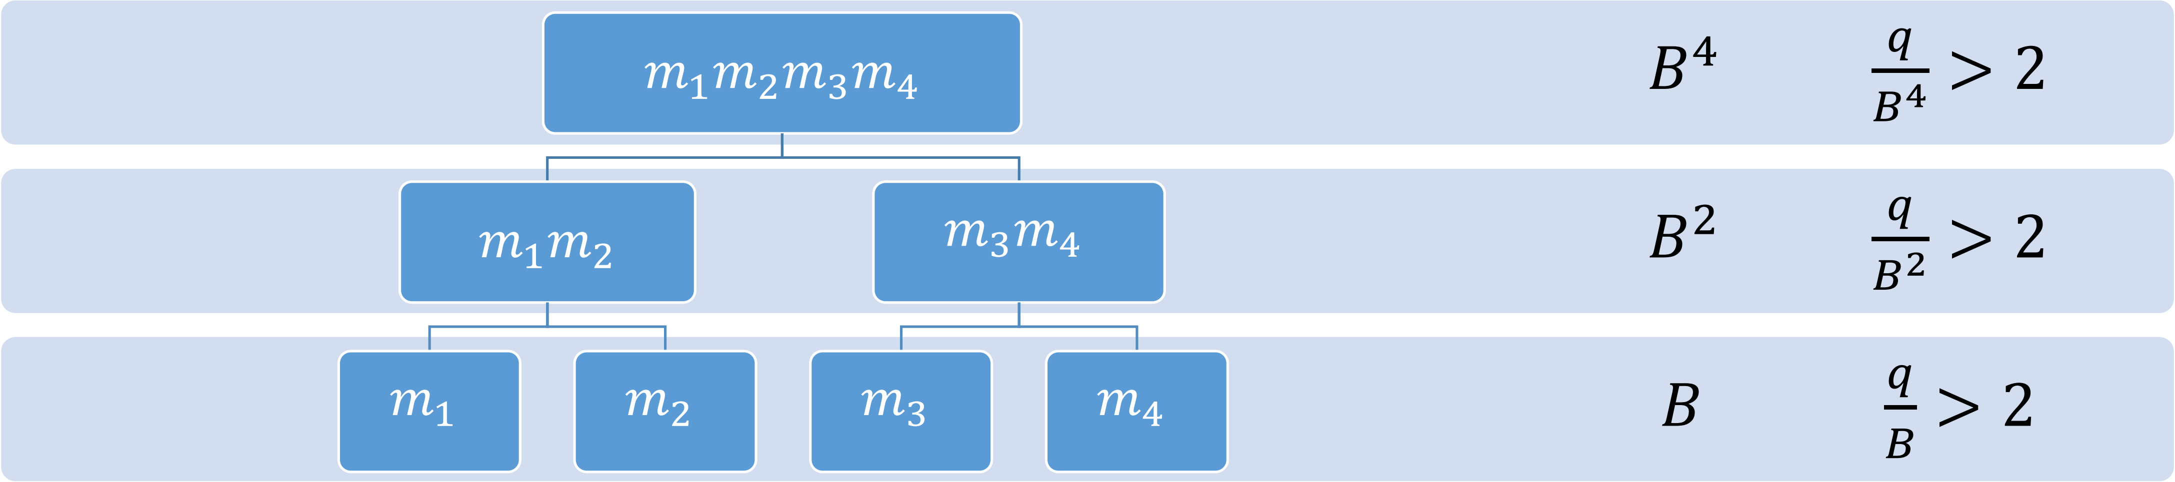

- $B^2 \rightarrow B^4, B^4 \rightarrow B^8, \ldots, B^{2^{L-1}} \rightarrow B^{2^L}$ (L levels of multiplications)

Modulus Switching

Brakerski, Gentry, Vaikuntanathan (BGV, 2012)

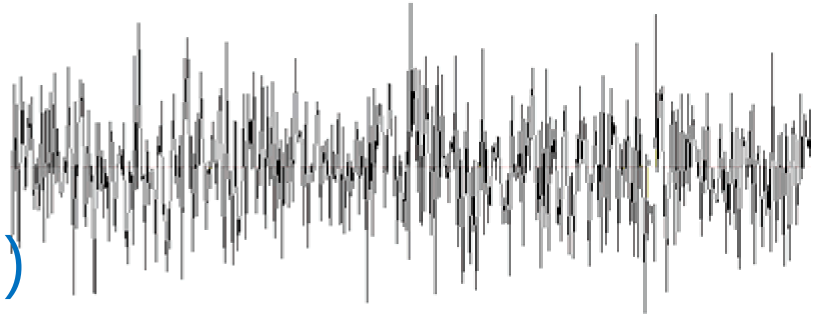

Switch (scale down) to a smaller modulus after each mult. level

- Need a chain of moduli $q = q_0, q_i \approx \frac{q_{i-1}}{B}$

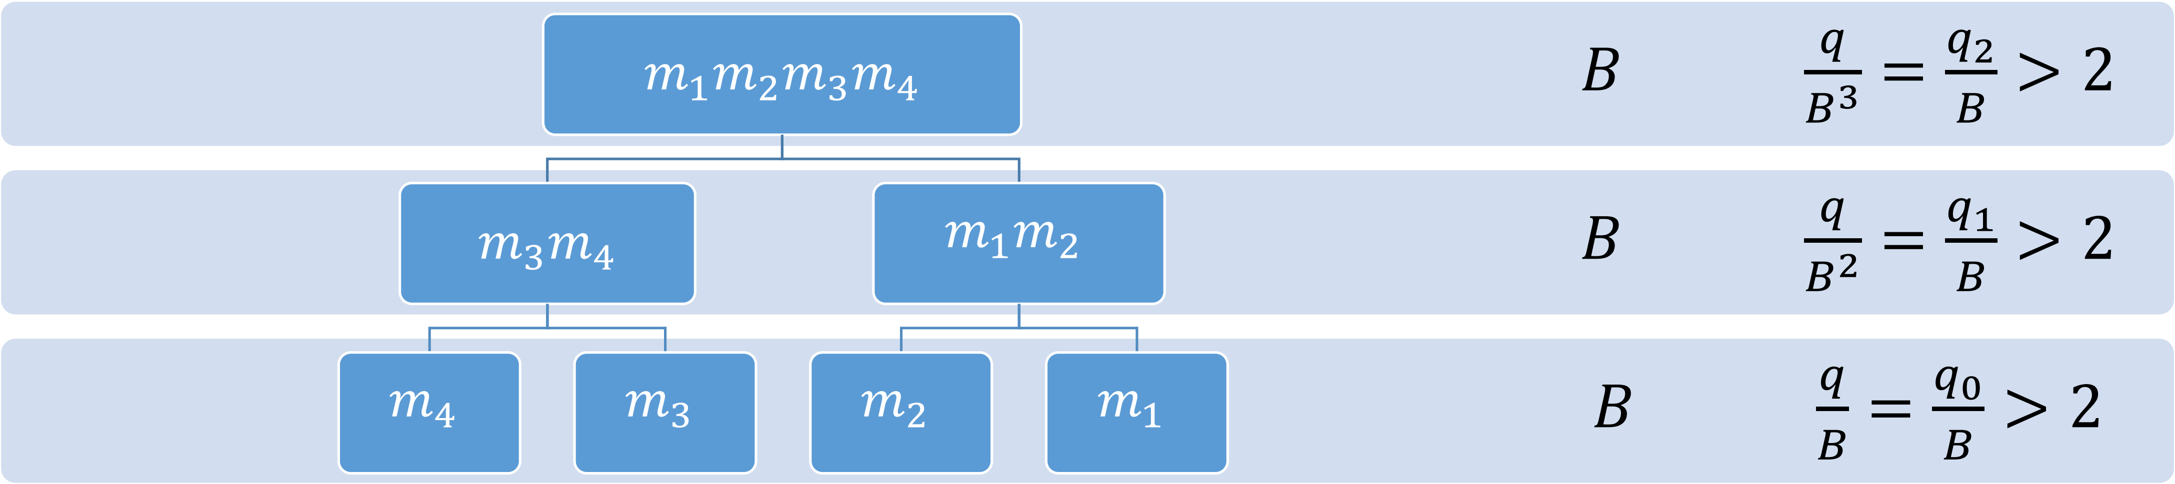

- $B^2 \to B^3 \to B^4, \ldots, \to B^L$ (L levels of mult)
- Leveled homomorphic encryption

Avoiding Modulus Switching

Message $m \in R/2R$

$s \leftarrow \chi$ secret key

Regev (2005) encryption for RLWE (Fan, Vercauteren 2012):

Sample $a \leftarrow R_q$ uniform, $e \leftarrow \chi$ noise

$b = \left\lfloor \frac{q}{2} \right\rfloor m + a \cdot s + e \mod q$, ciphertext $\mathrm{c} = (\mathrm{a}, \mathrm{b})$

$b - a \cdot s = \left\lfloor \frac{q}{2} \right\rfloor m + e$, decrypt: $\left\lfloor \frac{2}{\mathrm{q}} (b - a \cdot s) \right\rceil$

decrypts correctly if $\|e\|_\infty < \frac{q}{4}$ because

$\left\lfloor \frac{q}{2} \right\rfloor \cdot 2 = q - (q \bmod 2)$, i.e. $\left\lfloor \frac{q}{2} \right\rfloor \cdot \frac{2}{q} = 1 - \frac{q \bmod 2}{q}$

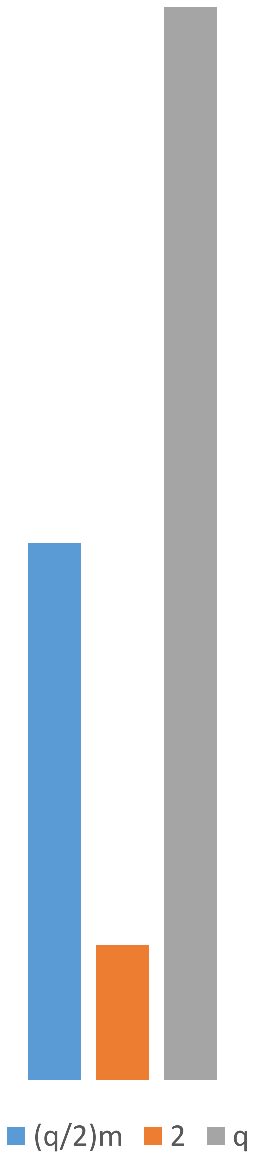

Scale-invariant Multiplication

Multiplication (FV):

- $(b_1 - a_1 \cdot s)(b_2 - a_2 \cdot s) = \left(\left\lfloor \frac{q}{2} \right\rfloor m_1 + e_1\right)\left(\left\lfloor \frac{q}{2} \right\rfloor m_2 + e_2\right)$
$$= \left\lfloor \frac{q}{2} \right\rfloor^2 m_1 m_2 + \left\lfloor \frac{q}{2} \right\rfloor (m_1 e_2 + m_2 e_1) + e_1 e_2$$
- $\frac{2}{q}(b_1 - a_1 \cdot s)(b_2 - a_2 \cdot s) = \left\lfloor \frac{q}{2} \right\rfloor m_1 m_2$
$$+ (m_1 e_2 + m_2 e_1) + \frac{2}{q} e_1 e_2 + \tilde{e}$$
- New noise term is of size $C \cdot B$, after L levels $C^L \cdot B$
C independent of B

Multi-key homomorphic encryption

López-Alt, Tromer, Vaikuntanathan (2012)

Message $m \in \{0,1\}$

Sample $f, g \leftarrow \chi$, $f = 1 + 2f'$ invertible mod q

secret key f, public key $h = \frac{2g}{f}$

NTRU-like encryption:

Encryption: Sample $s, e \leftarrow \chi$

$c = m + h \cdot s + 2e \bmod q$

Decryption: $m = (f \cdot c \bmod q) \bmod 2$, since

$f \cdot c = m + 2(gs + ef + mf')$,

decrypts correctly if $\|gs + ef + mf'\| < \frac{q}{2}$.

Multi-key homomorphic encryption

López-Alt, Tromer, Vaikuntanathan (2012)

$c_1 = m_1 + h_1 \cdot s + 2e_1 \quad f_1 \cdot c_1 = m_1 + 2(g_1 s_1 + f_1 e_1 + m_1 f_1') \bmod q$

$c_2 = m_2 + h_2 \cdot s + 2e_2 \quad f_2 \cdot c_2 = m_2 + 2(g_2 s_2 + f_2 e_2 + m_2 f_2') \bmod q$

Multiplication:

$(f_1 \cdot c_1)(f_2 \cdot c_2) = (m_1 + 2E_1)\,(m_2 + 2E_2)$

$= m_1 m_2 + 2(m_1 E_2 + m_2 E_1 + 2E_1 E_2)$

For $f_1 = f_2 = f$ (i.e. $g_1 = g_2 = g, h_1 = h_2 = h$):

Ciphertext $c_1 \cdot c_2 \bmod q$ decrypts under f^2 instead of f

Key switching transforms it back to a ciphertext that decrypts under f

Multi-key homomorphic encryption

López-Alt, Tromer, Vaikuntanathan (2012)

- Replaces uniform random $a \leftarrow R_q$ by polynomial ratio $h = \frac{2g}{f}$
- Security follows from RLWE if $h = \frac{2g}{f}$ looks uniform random

RLWE	**LATV12**
$a \leftarrow R_q$ uniform random Secret $s \leftarrow \chi$ Noise $e \leftarrow \chi$	PK: $h = \frac{2g}{f}$, SK: $f, g \leftarrow \chi$ Noise $s \leftarrow \chi$, $e \leftarrow \chi$
$b = a \cdot s + 2e$	c= $h \cdot s + 2e + m$

Modified NTRU

Stehlé, Steinfeld (2011)

LATV12 make an additional assumption, the
Small Polynomial Ratio (SPR) assumption:

- $\frac{g}{f}$ looks uniform random in R_q

Theorem (Stehlé, Steinfeld 2011):
If $d = 2^k, n = 2^{k-1}, R = \mathbf{Z}[X]/(X^n + 1)$, $\chi = D_{Z^n,\sigma}$
then the SPR assumption holds if $\sigma > \mathrm{poly}(n) \cdot \sqrt{q}$.

LATV12 conclude that such σ is too large for homomorphism

Observation

- The distribution for sampling f, g needs not be the same as that for sampling s, e
- Choose different distributions $f, g \leftarrow \chi_{\text{key}}$ and $s, e \leftarrow \chi_{\text{err}}$ with different standard deviations σ_{key} and σ_{err}

RLWE	**LATV12**
$a \leftarrow R_q$ uniform random Secret $s \leftarrow \chi_{\text{err}}$ Noise $e \leftarrow \chi_{\text{err}}$	PK: $h = \frac{2g}{f}$, SK: $f, g \leftarrow \chi_{\text{key}}$ Noise $s, e \leftarrow \chi_{\text{err}}$
$b = a \cdot s + 2e$	c= $h \cdot s + 2e + m$

Basic Encryption Scheme

- KeyGen: $f, g \leftarrow \chi_{\text{key}}$, $f = 1 + \mathrm{t}f'$ invertible mod q
  SK: f, PK: $h = \frac{\mathrm{t}g}{f}$
- Encrypt: $m \in R/tR$, $s, e \leftarrow \chi_{\text{err}}$, $c = \left\lfloor \frac{q}{t} \right\rfloor m + hs + e$
- Decrypt: $m = \left\lfloor \frac{t}{q} (f \cdot c \bmod q) \right\rceil \bmod t$
- $f \cdot c \equiv \left(\left\lfloor \frac{q}{t} \right\rfloor m + v \right) \bmod q$, v is the noise level in c
  Decryption is correct, if $\|v\|_\infty < \left(\left\lfloor \frac{q}{t} \right\rfloor - t \right) /2$
- Noise in a fresh ciphertext is $\|v\|_\infty < \delta t B_{\text{key}} (2B_{\text{err}} + t/2)$,
  where B_{key} and B_{err} are bounds on the norms of the noise polys

Homomorphic Multiplication

- First step: $\widetilde{c_3} = \left\lfloor \frac{t}{q}(c_1 \cdot c_2) \right\rceil \bmod q$

  But this needs to be decrypted with f^2
- Use the following functions:

  $$P_w(f) = \left(f \cdot w^i \bmod q\right)_{\mathrm{i}=0}^{\ell-1}$$

  and $D_w(c)$ is the base w decomposition of c, i.e.

  $D_w(c) = (c_i)_{i=0}^{\ell-1}$, $c = \sum_{i=0} c_i w^i$.

  Then $\langle D_w(c), P_w(f)\rangle = fc \bmod q$.
- In key generation compute and publish evaluation key

  $\gamma = P_w(f) + \boldsymbol{e} + h\boldsymbol{s}$, where $\boldsymbol{e}, \boldsymbol{s} \leftarrow \chi_{err}^{\ell}$, $\ell = \lfloor \log_w(q) \rfloor + 2$
- KeySwitch: compute $\mathrm{c}_3 = \langle D_w(\widetilde{c_3}), \gamma \rangle$

Noise Growth in Homomorphic Multiplication

- Assume c_1 and c_2 have noise levels bounded by V
- and key and noise distribution are bounded by B_{key} and B_{err}, resp.

- $f c_3 = \left\lfloor \frac{q}{t} \right\rfloor m_1 m_2 + v \mod q$

$$\|v\|_\infty < \delta^2 t^2 B_{\mathrm{key}} V + \delta^2 t^2 B_{\mathrm{key}}^2 + \delta^2 t \ell w B_{\mathrm{err}} B_{\mathrm{key}}$$

- Indeed, if σ_{key} is as demanded by Stehlé and Steinfeld, then there is no guarantee that the noise is less than q

Avoiding the SPR assumption

Use tensor products of decompositions and powers (see Brakerski 2012)

- Change multiplication from $\widetilde{c_3} = \left\lfloor \frac{t}{q}(c_1 \cdot c_2) \right\rceil \bmod q$ to $\widetilde{c_3} = \left\lfloor \frac{t}{q} P_w(c_1) \otimes P_w(c_2) \right\rceil \bmod q \in R_q^{\ell^2}$
- This intermediate ciphertext decrypts under $D_w(f) \otimes D_w(f)$
- Adjust evaluation key to

$$\gamma = f^{-1} P_w\big(D_w(f) \otimes D_w(f)\big) + \boldsymbol{e} + h\boldsymbol{s} \bmod \mathrm{q} \in R_q^{\ell^3}$$

- Noise bound is now

$$\|v\|_\infty < \delta^2 t \,\mathrm{w} \log_w(tB_{\mathrm{key}})\, V + \delta^2 t^2\, w \log_w(tB_{\mathrm{key}}) + \cdots$$

Avoiding the SPR assumption

Noise growth small enough to use Stehlé, Steinfeld setting
$d = 2^k, n = 2^{k-1}, R = \mathbf{Z}[X]/(X^n + 1), \chi = D_{Z^n,\sigma}, \sigma > \mathrm{poly}(n) \cdot \sqrt{q}$.

- PK is indistinguishable from uniform random element in R_q
- Tensoring helps with noise growth, but is rather unnatural and annoying

For a "more practical" version:

- Need SPR assumption, take narrow key distribution
- Power and decomposition functions with varying base w
  give more flexibility trading size of evaluation key vs. noise growth
- Use distributions of different widths for different purpose

Parameters

- Correctness via noise bounds
- Security via estimating runtime of attack on scheme in time 2^{80} based on Lindner-Peikert analysis

q (bits)	Dimension n	Size of elt in R	t	Levels L
128	2^{12}	66 KB	2	3
			1024	1
256	2^{13}	262 KB	2	7
			1024	4
1024	2^{15}	4.2 MB	2	31
			1024	19

Implementation

We have implemented homomorphic encryption with
127-bit prime q, $n = 4096$, $w = 2^{32}$

- plain C, no assembly (yet), a lot potential for optimization

Operation	Encrypt	Decrypt	Add	Mul
Cycles/10^6	79.2	14.1	0.07	90.7
ms	27	5	0.03	31

Intel Core i7-3520M @ 2.893 GHz

We have not implemented AES yet!
(Due to lack of motivation for using AES as a benchmark for HE.)

Improved Security for a Ring-Based Fully Homomorphic Encryption Scheme

Michael Naehrig
Cryptography Research Group
Microsoft Research

Joint work with Joppe W. Bos (MSR), Jake Loftus (University of Bristol), and Kristin Lauter (MSR)

PQCrypto 2013
Limoges, 5 June 2013

Thank you!