Round5 with ring lifting

CWG, September 14, 2018

Sauvik Bhattacharya, Scott Fluhrer, Oscar Garcia-Morchon,
Mike Hamburg, Thijs Laarhoven, Ronald Rietman,
Markku-Juhani O. Saarinen, Ludo Tolhuizen, Zhenfei Zhang

Outline

Introducing Round5

- NIST Post-Quantum Cryptography Standardization project
- NIST asked to merge proposals
- We looked for merge combinations with low bandwidth/communication requirements
- Round5 combines Round2 and HILA5

Round2: (R)LWR-based KEM and PKE with sparse ternary secrets

Round2: KEM and PKE based on (Ring) Learning with Rounding

- No explicit noise generation required, less calls to random
- Smaller alphabet sizes for public key and ciphertext
- Prime cyclotomic polynomial $\phi_{n+1}(x) = 1 + x + \ldots + x^n$ with n prime, and $\phi_{n+1}(x)$ irreducible modulo two
- Sparse ternary secrets

Round2 description

Alice | **Bob**

$a \xleftarrow{\$} \mathbb{Z}_q[x]/\phi(x)$, $s \xleftarrow{\$} \mathcal{S}$

$$\xrightarrow{a, b = \langle \lfloor \frac{p}{q} \langle as \rangle_\phi \rceil \rangle_p}$$

$r \xleftarrow{\$} \mathcal{S}$;

$$\xleftarrow{u = \langle \lfloor \frac{p}{q} \langle ar \rangle_\phi \rceil \rangle_p}$$

$$\xleftarrow{v = \langle \frac{t}{2} m + S_\mu(\lfloor \frac{t}{p} \langle br \rangle_\phi \rceil) \rangle_t}$$

$$w = \langle \frac{q}{t} v - \frac{q}{p} S_\mu(\langle us \rangle_\phi) \rangle_q$$
$$\hat{m} = \langle \lfloor \frac{2}{q} w + \frac{1}{2} \rfloor \rangle_2$$

$\mathcal{S}$ is a subset of all balanced ternary polynomials of Hamming weight h;
$\phi(x) = 1 + x + \ldots + x^n$ $S_\mu(f)$: μ highest order coefficients of f.

HILA5: RLWE- based KEM with error correction

HILA5: KEM based on Ring Learning with Errors

- Failure probability reduction by error correcting code Xe5, resulting in smaller public keys and ciphertexts
- Decoding Xe5 avoids table-lookups and conditions altogether and therefore is resistant to timing attacks.
- Five error correction by majority voting. For each information bit m_i, there are disjoint sets $S^i_1, \ldots, S^i_{10}$ of parity bit indices such that

$$m_i = \sum_{j \in S^i_k} p_j \text{ for } 1 \leq k \leq 10.$$

  Information bit i is flipped iff six or more sums equal one.

Round5 = Round2 + HILA5

Alice		**Bob**
$a \xleftarrow{\$} \mathbb{Z}_q[x]/\phi(x)$, $s \xleftarrow{\$} \mathcal{S}$		
	$\xrightarrow{a, b = \langle \lfloor \frac{p}{q} \langle as \rangle_\phi \rceil \rangle_p}$	
		$r \xleftarrow{\$} \mathcal{S}$;
	$\xleftarrow{u = \langle \lfloor \frac{p}{q} \langle ar \rangle_\phi \rceil \rangle_p}$	
		$c = \text{Encode}(m)$
	$\xleftarrow{v = \langle \frac{t}{2} c + S_\mu(\lfloor \frac{t}{p} \langle br \rangle_\phi \rceil) \rangle_t}$	
$w = \langle \frac{q}{t} v - \frac{q}{p} S_\mu(\langle us \rangle_\phi) \rangle_q$ $\hat{c} = \langle \lfloor \frac{2}{q} w + \frac{1}{2} \rfloor \rangle_2$ $\hat{m} = \text{Decode}\ (\hat{c})$		

$\mathcal{S}$ is a subset of all balanced ternary polynomials of Hamming weight h;
$\phi(x) = 1 + x + \ldots + x^n$ $S_\mu(f)$: μ highest order coefficients of f.

Benefit of error-correction

Round5 combines the LWR-based approach of Round2 and the error correcting code of HILA5.
Smaller public keys and ciphertext.
However, this assumes independence of errors...

From: Leo Ducas <leo.ducas1@gmail.com>
Sent: Tuesday, August 07, 2018 2:07 AM
To: pqc-forum
Subject: [pqc-forum] Re: OFFICIAL COMMENT: Round5 = Round2 + Hila5

Dear authors,

I note that the failure analysis assumes that ``bit failures occur independently'', but I'm unconvinced it would be the case, especially in the ring setting. I have searched for solution to this issue for a long time, and still don't know how to properly address this issue theoretically.

May I suggest to resort to experimental analysis to test how close or not to independent these failure events are, at least in a regime where failures are statistically measurable ?

Best regards
-- Leo Ducas

Simulation results: error values for prime cyclotomic ring

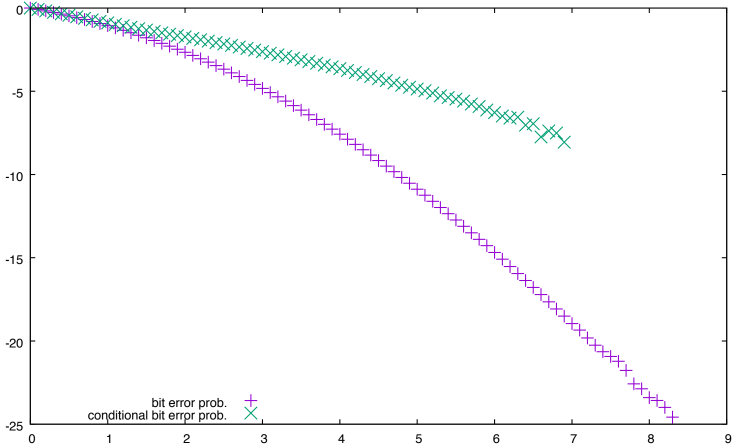

$\log_2(\text{Prob}(\text{error value} \geq x))$.

Issue with prime cyclotomic ring: correlated errors

One of the terms in the error in reconstruction is $\langle \langle se \rangle_\phi \rangle_q$.

$$\langle se \rangle_\phi = \sum_{k=0}^{n-1} [c_k(s, e) - c_n(s, e)] x^k,$$

where

$$c_j(s, e) = \sum_i s_i e_{\langle j-i \rangle_{n+1}}.$$

Hence, if $c_n(s, e)$ is large, then many coefficients of $\langle se \rangle_\phi$ may be large.

Round5 with ring lifting

Alice | **Bob**

Alice: $a \xleftarrow{\$} \mathbb{Z}_q[x]/\phi(x),\ s \xleftarrow{\$} \mathcal{S}$

Alice → Bob: $a, b = \langle \lfloor \frac{p}{q} \langle as \rangle_\phi \rceil \rangle_p$

Bob: $r \xleftarrow{\$} \mathcal{S};$

Bob → Alice: $u = \langle \lfloor \frac{p}{q} \langle ar \rangle_\phi \rceil \rangle_p$

Bob: $c = \text{Encode}(m)$

Bob → Alice: $v = \langle \frac{t}{2} c + S_\mu(\lfloor \frac{t}{p} \langle br \rangle_N \rceil) \rangle_t$

Alice:

$$w = \langle \frac{q}{t} v - \frac{q}{p} S_\mu(\langle us \rangle_N \rangle_q$$

$$\hat{c} = \langle \lfloor \frac{2}{q} w + \frac{1}{2} \rfloor \rangle_2$$

$$\text{m} = \text{Decode}\ (\hat{c})$$

$\mathcal{S}$ is a subset of all balanced ternary polynomials of Hamming weight h;

$\phi(x) = 1 + x + \ldots + x^n$ $\quad S_\mu(f)$: μ highest order coefficients of f.

$N(x) = (x-1)\phi(x) = x^{n+1} - 1$

Why this works (1)

$$\frac{q}{p}b = \langle as \rangle_\phi + e + q\lambda \text{ with } |e| \leq \frac{q}{2p}.$$

$$\frac{q}{p}br = \langle as \rangle_\phi r + er = (as + \lambda_1 \phi)r + er + q\lambda_2.$$

As $r(x) = (x-1)\rho(x) + r(1)$ for some $\rho \in \mathbb{Z}[x]$: $\langle \langle r \rangle_N \rangle q = 0$.

$$\frac{q}{p}br \equiv asr + er \pmod{N.q}$$

$$\frac{q}{p}us \equiv asr + e's \pmod{N, q}.$$

$$\frac{q}{p}(br - us) \equiv er - e's \pmod{N, q}.$$

Why this works (2)

$\frac{p}{t}\lfloor \frac{t}{p}\langle br\rangle_N \rceil = \langle br\rangle_N + e" \pmod{N, p}$

$$\frac{q}{t}v = \frac{q}{2}m + S_\mu(\frac{q}{p}\langle br\rangle_N + \frac{q}{p}e") \pmod{N, q}$$

$$w = \frac{q}{2}m + S_\mu(er - e's + \frac{q}{p}e") \pmod{N, q}.$$

So if $\langle er - e's\rangle_N + \frac{q}{p}e"$ is small (modulo q), then $w \approx \frac{q}{2}m$.
We got rid of the correlation between coefficients of $\langle er - e's\rangle_\phi$ caused by a large common term.

Simulation results: errors values in cyclic ring

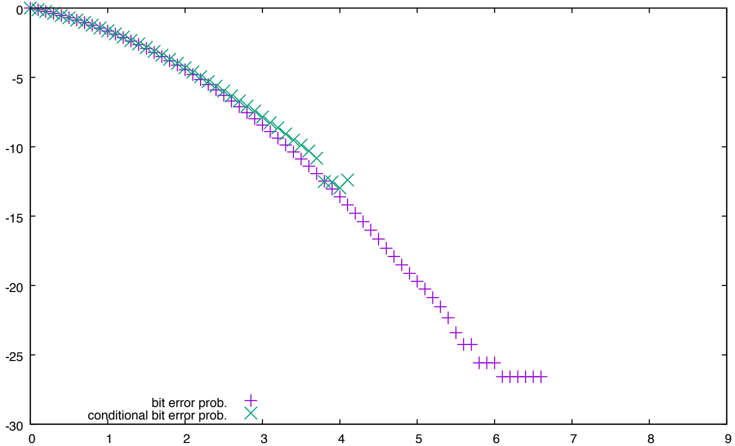

$\log_2(\text{Prob}(\text{error value} \geq x))$.

Simulation results: errors values for both rings

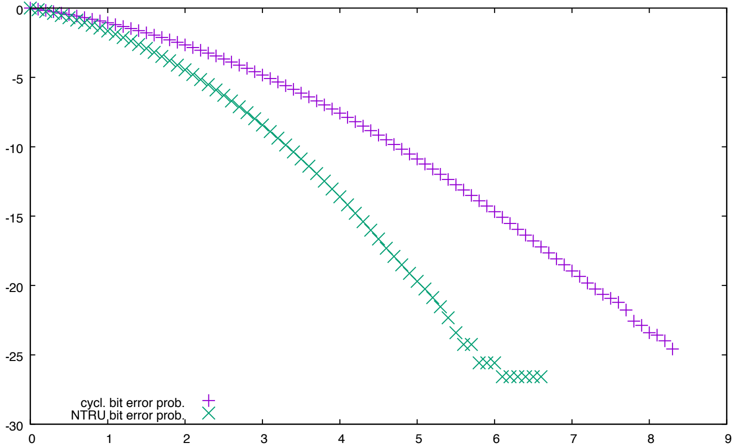

$\log_2(\text{Prob}(\text{error value} \geq x))$.

Benefits of ring lifting

Parameters (CCA NIST3)	**No FEC, cyclotomic**	**No FEC, ring lifting**	**Xef FEC, ring lifting**
d, n, h	852, 852, 212	820, 820, 254	756, 756, 242
q, p, t	$2^{12}, 2^9, 2^5$	$2^{12}, 2^9, 2^3$	$2^{12}, 2^8, 2^3$
$B, \bar{n}, \bar{m}, f$	1, 1, 1, 0	1, 1, 1, 0	1, 1, 1, 5
μ	192	192	$192 + 231$
Bandwidth	2087 B	1967 B	1720 B
Public key	984 B	948 B	781 B
Ciphertext	1103 B	1019 B	939 B
PQ Security	2^{181}	2^{176}	2^{181}
Classical	2^{193}	2^{192}	2^{193}
Failure rate	2^{-146}	2^{-162}	2^{-255}

S_μ stops the "evaluate at $x = 1$" attack

- "Evaluate at $x = 1$ attack" is a distinguishing attack
- Consider RLWE sample $(b, v = \langle br \rangle\rangle_N + e" + \frac{q}{2}m)\rangle_q$ with $\langle r(1) \rangle_q = 0$. As $(x-1)|N(x)$

$$v(1) \equiv e(1) + \frac{q}{2}m(1) \pmod{q}$$

so $\langle v(1) \rangle_q$ is not uniformly distributed.
If $\mu < n$, not all coefficients of $\langle br \rangle_N$ are available, so the evaluate at $x = 1$ attack does not apply.

CPA-Security proof for Round2 (1)

CPA: Chosen plaintext attack.

Adversary chooses two plaintexts, m_0 and m_1, after having seen a and b:

$$(m_0, m_1) = \mathcal{A}_1(a, b).$$

Adversary randomly chooses $k \in \{0, 1\}$ and encrypts m_k
Algorithm $\mathcal{A}_2$ runs on input (a, b, m_0, m_1, u, v) with output 0 or 1.
Output of game equals 1 if $\mathcal{A}_2(a, b, m_0, m_1, u, v) = k$ and zero otherwise. The advantage of $(\mathcal{A}_1, \mathcal{A}_2)$ equals

$$\mid \text{Prob[game output} = 1] - \frac{1}{2} \mid$$

where the probability over in the randomness in (a, b, u, v).

CPA-Security proof for Round2 (2)

- Sequence of CPA games. Gradual replacement of variables, ending with all variables being uniform.
- Two consecutive games can be used to construct a distinguisher between samples of the random variables in which these games differ.
  - Advantage of the constructed distinguisher equals the absolute value of the difference of the probabilities that the respective games output a 1.
- The advantage of the original CPA game is at most the sum of the advantages of the distinguishers for the replaced variables.
  - If the original CPA game has a large advantage, at least one of the distinghuishers has a large advantage.

Adapting the reduction for Round5 with ring lifting

The reduction proof from Round2 does not work for Round5 with ring lifting in the step where the distribution of

$$\begin{bmatrix} u \\ v' \end{bmatrix} = \begin{bmatrix} \lfloor \frac{z}{q} \langle ar \rangle_{\phi} \rceil \\ S_{\mu}(\lfloor \frac{z}{q} \langle br \rangle_{N} \rceil) \end{bmatrix}$$

is replaced by a uniform distribution.

With Round2, v' also involves rounding modulo ϕ, so $\begin{bmatrix} u \\ v' \end{bmatrix}$ has two R-LWR samples from the same ring.

With Round5 with ring lifting, $\begin{bmatrix} u \\ v' \end{bmatrix}$ has two R-LWR samples involving r from different rings.

Related result for lifted RLWE [1, Lemma 11]

Let $n+1$ be prime, and let q be relatively prime to $n+1$.
Assume that it is hard to distinguish samples $(a_i, b_i = a_i s + e_i) \in (\mathbb{Z}_q[x]/\Phi_{n+1}(x))^2$ from uniform,
Then the samples $(L_q(a_i), L_q((1-x)b_i)) \in S_{n+1,q}^2$ are also hard to distinguish from uniform.

Here $S_{n+1,q} = \{\sum_{i=0}^{n} a_i x^i \mid \sum_{i=0}^{n} a_i x^i \equiv 0 \pmod{q}\}$, and
$L_q(a(x)) = a(x) - (n+1)^{-1} \cdot a(1)\Phi_{n+1}(x)$.

[1] G. Bonnoron, L. Ducas and M. Fillinger, "Large FHE gates from Tensored Homomorphic Accumulator", iacr preprint Report 2017-996.

Applicability to Round5

In the proof in [1], the error polynomial e_i is lifted to $L_q((x-1)e_i(x)) = (x-1)e_i(x)$. Hence, if coefficients of each e_i are drawn independently, this is not true anymore for the coefficients after lifting.

Different even coefficients of $(x-1)f(x)$ do not contain a common coefficient from f. Hence, if $\mu < n/2$, we can let S_μ select μ even coefficients of a polynomial, and the dependence has been removed.

Can we generalize this RLWE result to RLWR?

Direction to make the proof work for R-LWR

From a discussion with Léo Ducas, we gathered the following possible way forward.

- Compute $b = \lfloor \frac{p}{q} \langle as \rangle_N \rfloor$.
- Transmit $\langle \tilde{b} \rangle_p$, where $\tilde{b}$ is the closest vector to b in the root lattice

$$\{(x_0, \ldots, x_n) \in \mathbb{Z}^{n+1} \mid \sum_{i=0}^{n} x_i = 0\}.$$

If $a(1) = 0$ or $s(1) = 0$, the noise introduced by transforming b to $\tilde{b}$ has a root at zero.
$\tilde{b}$ can be found in time $\mathcal{O}(n \log n)$, see MCKilliam, Clarkson, Quinn, "An algorithm to compute the nearest point in the lattice A_n^*, arxiv.org, Report 0801.1364, 2008.

Conclusions

- Applying error correction together with RLWR leads to solutions with small failure probability and small public keys and cipher texts, provided the error correlation is low.
- In Round5 with ring lifting, error correlation and failure probability are as small as with the cyclic ring.
- S_μ destroys ring structure, making the "evaluate at $x = 1$" attack infeasible.
- How to adapt the CPA-security proof of Round2 for ciphertext components computed in different rings?