BIKE

5th NIST PQC Standardization Workshop
April 10th, 2024

Nicolas Aragon, University of Limoges, France
Paulo L. Barreto, University of Washington Tacoma, USA
Slim Bettaieb, Worldline, France
Loïc Bidoux, Worldline, France
Olivier Blazy, University of Limoges, France
Jean-Christophe Deneuville, ENAC, Federal University of Toulouse, France
Philippe Gaborit, University of Limoges, France
Santosh Ghosh, Intel, USA
Shay Gueron, University of Haifa, and Meta, Israel & USA
Tim Güneysu, Ruhr-Universität Bochum & DFKI, Germany
Carlos Aguilar Melchor, University of Toulouse, France
Rafael Misoczki, Meta, USA
Edoardo Persichetti, Florida Atlantic University, USA
Jan Richter-Brockmann, Ruhr-Universität Bochum, Germany
Nicolas Sendrier, INRIA, France
Jean-Pierre Tillich, INRIA, France
Valentin Vasseur, Thales, France
Gilles Zémor, IMB, University of Bordeaux, France

https://bikesuite.org

Agenda

- BIKE Recap
- Proposed tweaks
  - Enhancing multi-target security
  - Constant-weight sampler
  - Decoder

BIKE Recap

KeyGen : $() \mapsto (h_0, h_1, \sigma), h$

Output: $(h_0, h_1, \sigma) \in \mathcal{H}_w \times \mathcal{M}$, $h \in \mathcal{R}$

$(h_0, h_1) \xleftarrow{\mathcal{D}} \mathcal{H}_w$ ▷ (1)

$h \leftarrow h_1 h_0^{-1}$

$\sigma \xleftarrow{\$} \mathcal{M}$

Encaps : $h \mapsto K, c$

Input: $h \in \mathcal{R}$

Output: $K \in \mathcal{K}$, $c \in \mathcal{R} \times \mathcal{M}$

$m \xleftarrow{\$} \mathcal{M}$

$(e_0, e_1) \leftarrow \mathbf{H}(m)$

$c \leftarrow (e_0 + e_1 h, m \oplus \mathbf{L}(e_0, e_1))$

$K \leftarrow \mathbf{K}(m, c)$

Decaps : $(h_0, h_1, \sigma), c \mapsto K$

Input: $((h_0, h_1), \sigma) \in \mathcal{H}_w \times \mathcal{M}$, $c = (c_0, c_1) \in \mathcal{R} \times \mathcal{M}$

Output: $K \in \mathcal{K}$

$e' \leftarrow \texttt{decoder}(c_0 h_0, h_0, h_1)$ ▷ $e' \in \mathcal{R}^2 \cup \{\perp\}$

$m' \leftarrow c_1 \oplus \mathbf{L}(e')$ ▷ with the convention $\perp = (0, 0)$

if $e' = \mathbf{H}(m')$ **then** $K \leftarrow \mathbf{K}(m', c)$ **else** $K \leftarrow \mathbf{K}(\sigma, c)$

Design

- Niederreiter-based KEM instantiated with QC-MDPC codes (faster polynomial inversion by [DGK'20]).
- Leverage Fujisaki-Okamoto Transform [DGKP'21].
- State-of-the-art QC-MDPC Decoding Failure Rate analysis.
- Black-Gray-Flip Decoder implemented in constant time.

NOTATION	
$\mathbb{F}_2$:	Binary finite field.
$\mathcal{R}$:	Cyclic polynomial ring $\mathbb{F}_2[X]/(X^r - 1)$.
$\mathcal{H}_w$:	Private key space $\{(h_0, h_1) \in \mathcal{R}^2 \mid \lvert h_0 \rvert = \lvert h_1 \rvert = w/2\}$
$\mathcal{E}_t$:	Error space $\{(e_0, e_1) \in \mathcal{R}^2 \mid \lvert e_0 \rvert + \lvert e_1 \rvert = t\}$
$\lvert g \rvert$:	Hamming weight of a binary polynomial $g \in \mathcal{R}$.
$u \xleftarrow{\$} U$:	Variable u is sampled uniformly at random from the set U.
$\oplus$:	exclusive or of two bits, componentwise with vectors

Functions

- $\mathbf{H} : \mathcal{M} \to \mathcal{E}_t$.
- $\mathbf{K} : \mathcal{M} \times \mathcal{R} \times \mathcal{M} \to \mathcal{K}$.
- $\mathbf{L} : \mathcal{R}^2 \to \mathcal{M}$

Parameters

r : block length
w : row weight
t : error weight
ℓ : shared secret size
$\mathcal{M}$: message space in $\{0,1\}^\ell$
$\mathcal{K}$: key space in $\{0,1\}^\ell$

BIKE Recap

Level 1

	Public key size (bytes)	Ciphertext size (bytes)	KeyGen (kilocycles)	Encaps (kilocycles)	Decaps (kilocycles)
BIKE	1,540	1,572	589	97	1,135
HQC	2,249	4,497	187	419	833
mceliece348864	261,120	128	140,870	46	137
Kyber-512	800	768	123	155	289

Level 3

	Public key size (bytes)	Ciphertext size (bytes)	KeyGen (kilocycles)	Encaps (kilocycles)	Decaps (kilocycles)
BIKE	3,082	3,114	1,823	223	3,887
HQC	4,522	9,042	422	946	1,662
mceliece460896	524,160	188	441,517	83	273
Kyber-768	1,184	1,088	213	249	275

BIKE performance numbers from Drucker, Gueron, Kostic, "Additional implementation of BIKE (Bit Flipping Key Encapsulation)". https://github.com/awslabs/bike-kem.

Security in the Multi-Target Setting

- In [WWW’23], a multi-target attack against the CCA variant of BIKE leveraging decryption failures was presented
  - The attack needs to first identify a key (out of many, e.g. 2^{87} keys) so that the gathering property is observed
  - Queries per target: ~2^{29} steps
  - Total complexity: ~2^{116} steps
- The attack is defeated by binding the public key to the ciphertext

[WWW’23]: Wang, Tianrui, Anyu Wang, and Xiaoyun Wang. "Exploring Decryption Failures of BIKE: New Class of Weak Keys and Key Recovery Attacks." *Annual International Cryptology Conference*. Cham: Springer Nature Switzerland, 2023.

Security in the Multi-Target Setting

$(sk, \sigma, h) \xleftarrow{\$} \texttt{Keygen}()$

1. Generate $\sigma \xleftarrow{\$} \{0,1\}^{256}$
2. $sk = (h_0, h_1) \xleftarrow{\$} \mathcal{R}^2$ with $wt(h_0) = wt(h_1) = w$ odd
3. $h = h_1 h_0^{-1}$
4. Return (sk, σ, h)

$(C, K) \xleftarrow{\$} \texttt{Encaps}(h)$

1. Generate a message $m \xleftarrow{\$} \mathcal{M}$
2. Compute error vectors $(e_0, e_1) = \mathbf{H}(f_i(m, h))$ with $wt(e_0, e_1) = t$ and $e_0, e_1 \in \mathcal{R}$.
3. Compute the ciphertext $C = (c_0, c_1) = (e_0 + e_1 h, m \oplus \mathbf{L}(e_0, e_1))$
4. Compute the shared key $K = \mathbf{K}(m, C)$

$m = \texttt{Decaps}(sk, \sigma, h, C)$

1. $m' = \text{Decode}(sk, C)$ // Or $\perp$ on decoding failure.
2. If $((m' \neq \perp)$ and $(C == \text{ReEncrypt}(m', h)))$ return $\mathbf{K}(m', C)$
   ▷ ReEncrypt uses $\mathbf{H}(f_i(m', h))$ instead of $\mathbf{H}(m')$
3. Else return $\mathbf{K}(\sigma, C)$

- In [DGK'21], a fix has been proposed, performance-studied and implemented.
- BIKE will adopt this protection moving forward, thus defeating such multi-target attacks.

$$f_1 : \mathcal{M} \times \mathcal{PK} \longrightarrow \{0,1\}^{256}$$
$$(m, pk) \longmapsto H(m \mid\mid pk)$$

$$f_2 : \mathcal{M} \times \mathcal{PK} \longrightarrow \{0,1\}^{256}$$
$$(m, pk) \longmapsto H(m \mid\mid H'(pk))$$

[DGK'21]: Drucker, Nir, Shay Gueron, and Dusan Kostic. "Binding BIKE errors to a key pair." International Symposium on Cyber Security Cryptography and Machine Learning. Cham: Springer International Publishing, 2021.

Constant Weight Sampling

- Variant of Fisher-Yates algorithm was vulnerable to timing attacks [GHJ'22]]. Latest spec fixed this: a constant-time variant (slightly biased output).
- Using a biased sampler in Encaps/Decaps has no impact to security [Sen'23]. To avoid code duplication, we also use the biased sampler in KeyGen.
- However, in KeyGen there would be an impact to the security reduction [DHK'23] (assumption $h = h^{-1}_0 h_1$ must be indistinguishable from random when h_0 and h_1 are produced by the biased sampler).
- Proposed tweak: Revert BIKE KeyGen to unbiased constant-weight sampler.

[GHJ'22]: Qian Guo, Clemens Hlauschek, Thomas Johansson, Norman Lahr, Alexander Nilsson, and Robin Leander Schröder. Don't reject this: Key-recovery timing attacks due to rejection-sampling in HQC and BIKE. IACR Trans. Cryptogr. Hardw. Embed. Syst., 2022(3):223–263, 2022.
[Sen'23]: Nicolas Sendrier. Secure sampling of constant-weight words – Application to BIKE. Cryptology ePrint Archive, Report 2021/1631, August 2023.
[DGK'23]: Drucker, N., Gueron, S., Kostic, D.: To Reject or Not Reject: That Is the Question. The Case of BIKE Post Quantum KEM. In: Latifi, S. (ed.) ITNG 2023 20th International Conference on Information Technology-New Generations. pp. 125–131. Springer International Publishing, Cham (2023).

New Bit-Flipping Decoder

- Simpler algorithm (no black/gray iterations) with modified threshold schedule.
- State-of-the-art extrapolation techniques [SV’20] for the waterfall region predict a DFR at most 2^{-180} for level 1 with blocklength 12,323.
- Better resistance to weak key attacks [WWW’23]. In the region where simulation is possible, the attack requires 2^{168} decapsulation queries instead of 2^{116}.

[SV’20: Sendrier, Nicolas and Vasseur Valentin. “About Low DFR for QC-MDPC Decoding.” *PQCrypto 2020.*
[WWW’23]: Wang, Tianrui, Anyu Wang, and Xiaoyun Wang. "Exploring Decryption Failures of BIKE: New Class of Weak Keys and Key Recovery Attacks." *Annual International Cryptology Conference*. Cham: Springer Nature Switzerland, 2023.

New Bit-Flipping Decoder

Algorithm 1 `New BIKE Decoder`

Input: $s \in \mathbf{F}_2^r$, $H \in \mathbf{F}_2^{r \times n}$

$\tilde{e} \leftarrow 0^n$; $\tilde{s} \leftarrow s$
for $i = 1, \dots, \mathrm{NbIter}$ **do**
  $T \leftarrow \textsc{threshold}(i, s, \tilde{s})$
  for $j = 0, \dots, n-1$ **do**
    $\sigma_j \leftarrow \mathtt{ctr}(H, \tilde{s}, j)$
  for $j = 0, \dots, n-1$ **do**
    if $\sigma_j \geq T$ **then**
      $\tilde{e}_j \leftarrow \tilde{e}_j \oplus 1$
      $\tilde{s} \leftarrow \tilde{s} - \mathrm{col}(H, j)$
return $\tilde{e}$

$\mathtt{ctr}(H, \tilde{s}, j)$ number of unsatisfied equations involving position j

function $\textsc{threshold}(i, s, \tilde{s})$
  $T' \leftarrow f_t(|s|)$ ▷ optimal
  $M \leftarrow (d+1)/2$ ▷ majority
  if $i = 1$ **then** $T \leftarrow T' + \delta$
  if $i = 2$ **then** $T \leftarrow (2T' + M)/3 + \delta$
  if $i = 3$ **then** $T \leftarrow (T' + 2M)/3 + \delta$
  if $i \geq 4$ **then** $T \leftarrow M + \delta$
  return $\max(f_t(|\tilde{s}|), T)$

$f_t(x) = 0.006258 \cdot x + 11.094, \delta = 3$ (level 1)

References

- [DGK]: Drucker, Gueron, Kostic, "Additional implementation of BIKE (Bit Flipping Key Encapsulation)". https://github.com/awslabs/bike-kem.
- [DGK’20]: Drucker, N., Gueron, S., Kostic, D.: Fast Polynomial Inversion for Post Quantum QC-MDPC Cryptography. Cyber Security Cryptography and Machine Learning. pp. 110–127. Springer International Publishing, Cham (2020). https://doi.org/10.1007/978-3-030-49785-9_8
- [DGK’23]: Drucker, N., Gueron, S., Kostic, D.: To Reject or Not Reject: That Is the Question. The Case of BIKE Post Quantum KEM. In: Latifi, S. (ed.) ITNG 2023 20th International Conference on Information Technology-New Generations. pp. 125–131. Springer International Publishing, Cham (2023).
- [DGKP’21] Drucker, N., Gueron, S., Kostic, D., Persichetti, E.: On the applicability of the Fujisaki-Okamoto transformation to the BIKE KEM. Int. J. Comput. Math. Comput. Syst. Theory 6(4), 364–374 (2021). https://doi.org/10.1080/23799927.2021.1930176
- [WWW’23]: Wang, Tianrui, Anyu Wang, and Xiaoyun Wang. "Exploring Decryption Failures of BIKE: New Class of Weak Keys and Key Recovery Attacks." Annual International Cryptology Conference. Cham: Springer Nature Switzerland, 2023.
- [GHJ’22]: Qian Guo, Clemens Hlauschek, Thomas Johansson, Norman Lahr, Alexander Nilsson, and Robin Leander Schröder. Don’t reject this: Key-recovery timing attacks due to rejection-sampling in HQC and BIKE. IACR Trans. Cryptogr. Hardw. Embed. Syst., 2022(3):223–263, 2022.
- [Sen’23]: Nicolas Sendrier. Secure sampling of constant-weight words – Application to BIKE. Cryptology ePrint Archive, Report 2021/1631, August 2023.

Questions?

https://bikesuite.org

Constant Weight Sampling

Algorithm 3 WSHAKE256-PRF(seed, len, wt)

Require: seed (32 bytes), len, wt

Ensure: A list (wlist) of wt distinct elements in $\{0, \ldots, \mathsf{len} - 1\}$.

$\mathsf{wlist} \leftarrow ()$ ▷ empty list

$s_0, \ldots, s_{\mathsf{wt}-1} \leftarrow$ SHAKE256-Stream$(\mathsf{seed}, 32 \cdot \mathsf{wt})$
▷ parse as a sequence of wt non negative 32-bits integers

for $i = (\mathsf{wt} - 1), \ldots, 1, 0$ **do** ▷ i decreasing from $\mathsf{wt} - 1$ to 0

  $\mathrm{pos} \leftarrow i + \lfloor (\mathsf{len} - i) s_i / 2^{32} \rfloor$

  $\mathsf{wlist} \leftarrow \mathsf{wlist}, (\mathrm{pos} \in \mathsf{wlist})\ ?\ i\ :\ \mathrm{pos}$

return wlist