Backflip: An Improved QC-MDPC Bitflipping Decoder

Nicolas Sendrier [1] Valentin Vasseur [1 2]

[1]Inria, Paris [2]Université Paris Descartes, Sorbonne Paris Cité

Results (BIKE parameters for 256 bits of security: $w = 274$ and $t = 264$ are fixed)

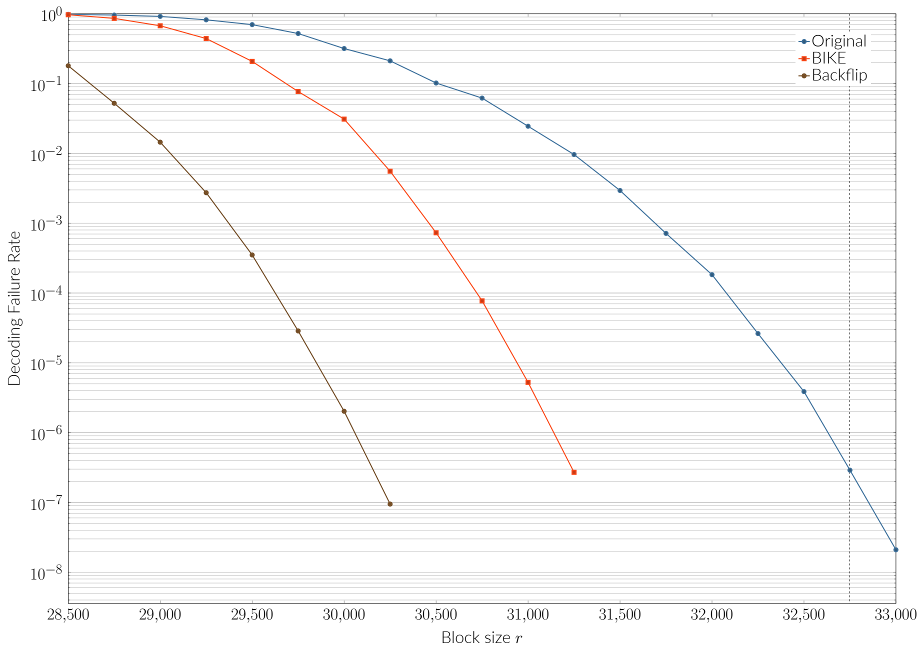

Quasi-Cyclic Moderate Density Parity Check codes [MTSB13]

- Similar construction to that of LDPC codes but with denser matrices
- Allow a McEliece-like public-key encryption scheme with a quasi-cyclic structure
  - Reasonable key sizes
  - Reduction to generic hard problems over quasi-cyclic codes

BIKE cryptosystem

- Key encapsulation mechanism using QC-MDPC codes
- 2nd round candidate to the NIST post-quantum cryptography standardization process

Focus:

- IND-CCA variant
  - Need a small Decoding Failure Rate (DFR), *e.g.* 2^{-128}
  - Need a proof of its DFR

Syndrome decoding for MDPC

H: moderately dense parity check matrix of size $r \times n$ (Hamming weight w of a row in $O(\sqrt{n})$)
e: error pattern (length n, Hamming weight t in $O(\sqrt{n})$)
s: corresponding syndrome
$\mathsf{s} = \mathsf{eH}^T$

Problem:

Knowing H and s, find e

Known algorithms:

- **Hard decoders:** bitflipping algorithm and its variants
- **Soft decoders:** sum-product algorithm and its variants

Main idea of the bitflipping algorithm

For a position j, compute its **counter**: the number of unverified equations it is involved in.

$|\mathsf{s} \cap \mathsf{h}_j|$: counter
If $j \notin \mathsf{e}$, $|\mathsf{s} \cap \mathsf{h}_j|$ is small
If $j \in \mathsf{e}$, $|\mathsf{s} \cap \mathsf{h}_j|$ is big

Original bitflipping algorithm

```
Input
    H ∈ {0,1}^{r×n}
    s = eH^T ∈ {0,1}^r
    |e| ≤ t
Output
    e ∈ {0,1}^r
e ← 0
while |s − eH^T| ≠ 0 do
    s' ← s − eH^T
    T ← threshold(context)
    for j ∈ {0, ..., n − 1} do
        if |s' ∩ h_j| ≥ T then
            e_j ← 1 − e_j
return e
```

Problem of the original algorithm

Algorithm sometimes takes **bad decisions** (adding errors instead of removing them)

- Bad flips are not always easy to detect
- Too many bad flips hinder progress of the algorithm and can lock it

Ideas of our variant

- Approach **soft decoding** by adjusting the duration of a flip in function of its reliability
- Regularly and systematically cancel oldest flips to avoid locking
- Each flip has a **time-to-live** (from 1 to 5 iterations)
- Most reliable flips (higher counters) live longer
- Threshold selection rule should be adapted

Small added cost of our variant

- For each flip, a time-to-live is computed
- F is a vector storing the time-of-death of each flipped position
- At the beginning of every iteration, obsolete flips are canceled

Backflipping algorithm

```
Input
    H ∈ {0,1}^{r×n}; s = eH^T ∈ {0,1}^r
    |e| ≤ t
Output
    e ∈ {0,1}^r
e ← 0;   F ← 0;   now ← 1
while |s − eH^T| ≠ 0 do
    for each j such that F_j = now do
        e_j ← 1 − e_j;   F_j ← 0
    now ← now + 1
    s' ← s − eH^T
    T ← threshold(context)
    for j ∈ {0, ..., n − 1} do
        if |s' ∩ h_j| ≥ T then
            e_j ← 1 − e_j
            if F_j ≥ now then
                F_j ← 0
            else
                F_j ← now + ttl(context)
return e
```

Time-to-live: ttl(δ)

δ: difference between the counter and the threshold
ttl: saturating affine function in δ

$$\mathrm{ttl}(\delta) = \max(1, \min(\mathrm{max\_ttl}, \lfloor A\,\delta + B \rfloor))$$

Using optimization methods to minimize the DFR:

security	max_ttl	A	B
128	5	0.45	1.1
192	5	0.36	1.41
256	5	0.45	1

BIKE-1 and BIKE-2

Thresholds: threshold($|\mathsf{s}|, |\mathsf{e}|$)

From [Cha17], a good threshold is the smallest T such that

$$|\mathsf{e}|\, f_{d,\pi_1}(T) \geq (n - |\mathsf{e}|) f_{d,\pi_0}(T)\,.$$

with

$$\pi_0 = \frac{\bar{\sigma}_{\mathrm{corr}}}{d} = \frac{(w-1)\,|\mathsf{s}| - X}{d(n - |\mathsf{e}|)} \quad \text{and} \quad \pi_1 = \frac{\bar{\sigma}_{\mathrm{err}}}{d} = \frac{|\mathsf{s}| + X}{d\,|\mathsf{e}|}$$

and $f_{d,\pi}$ is the probability mass function of a random variable following a binomial distribution of parameters d and π

π_0 and π_1 depend on

- |s| which we can know,
- |e| which we cannot.

Assume that $|\mathsf{e}| = t - |F|$

- true if no error was added,
- gives a more conservative threshold otherwise.

Estimating the DFR for BIKE parameters

- In [SV18] a simplified bitflipping algoritm is defined and a model is proposed
- A small difference is observed between the DFR obtained in the model and the DFR obtained by simulation, but the same behaviour is observed
- Other bitflipping algorithms also follow the same behaviour
- In the model, at worst $\log(\mathrm{DFR})$ is an affine function of the block size r
- DFR values for BIKE parameters are estimated by reducing the block size r to measure failures by simulation and then extrapolated assuming the above behaviour

BIKE-1 and BIKE-2 parameters for IND-CCA security using backflip

Achieving a DFR of $2^{-\lambda/2}$ where λ is the security parameter

security	Original r	Revised r	Ratio
128	10 163	10 253	1.009
192	19 853	21 059	1.061
256	32 749	34 939	1.067

Achieving a DFR of $2^{-\lambda}$ where λ is the security parameter

security	Original r	Revised r	Ratio
128	10 163	11 779	1.159
192	19 853	24 821	1.250
256	32 749	40 597	1.240

References

[Cha17] Julia Chaulet. "Étude de cryptosystèmes à clé publique basés sur les codes MDPC quasi-cycliques". PhD thesis. University Pierre et Marie Curie, Mar. 2017. URL: `https://tel.archives-ouvertes.fr/tel-01599347`.

[MTSB13] Rafael Misoczki et al. "MDPC-McEliece: New McEliece variants from Moderate Density Parity-Check codes". In: 2013, pp. 2069–2073.

[SV18] Nicolas Sendrier and Valentin Vasseur. *On the Decoding Failure Rate of QC-MDPC Bit-Flipping Decoders*. Cryptology ePrint Archive, Report 2018/1207. `https://eprint.iacr.org/2018/1207` - To appear in PQCrypto 2019. 2018.