BIKE, Classic McEliece, HQC

Comparing and contrasting NIST PQC 4th Round KEMs

Introductions

- Panelist – Carlos Aguilar Melchor, on behalf of the HQC team
- Panelist – Edoardo Persichetti on behalf of the Classic McEliece team
- Panelist – Nicolas Sendrier on behalf of the BIKE team
- Moderator – Angela Robinson, NIST

Panel Overview

- Security
  - Cryptanalysis
  - Side-channel attacks
  - Open security questions
  - Other desirable security properties
- Performance
  - Memory requirements and computational costs
  - Use-cases
- Other considerations

Security

- Cryptanalysis
  - Please discuss any cryptanalysis that has been presented against your submission and what changes you have made to mitigate these attacks.
  - Have you made any changes or tweaks to achieve additional desirable security properties beyond IND-CCA2?

Security

- Cryptanalysis
  - Please discuss any cryptanalysis that has been presented against your submission and what changes you have made to mitigate these attacks.
  - Have you made any changes or tweaks to achieve additional desirable security properties beyond IND-CCA2?
- Other desirable security properties
  - Please discuss what sort of protections are needed for your algorithm to be secure against side-channel attacks.
  - Does your algorithm feature any other desirable security properties (resistance to misuse, multi-target attack resistance, etc.)?

Security

- Cryptanalysis
  - Please discuss any cryptanalysis that has been presented against your submission and what changes you have made to mitigate these attacks.
  - Have you made any changes or tweaks to achieve additional desirable security properties beyond IND-CCA2?
- Other desirable security properties
  - Please discuss what sort of protections are needed for your algorithm to be secure against side-channel attacks.
  - Does your algorithm feature any other desirable security properties (resistance to misuse, multi-target attack resistance, etc.)?
- Open security questions
  - What impact does the structure of your underlying code have on security?
  - What is the current state of DFR analysis for BIKE and HQC?

Performance

NIST Security Category 1, taken from algorithm specifications

- BIKE – AVX512
- HQC AVX2 optimized
- Classic McEliece AVX

ML-KEM performance included for comparison

	Public key size (bytes)	Private key (bytes)	Ciphertext size (bytes)	KeyGen (kilocycles)	Encaps (kilocycles)	Decaps (kilocycles)
BIKE	1,540	2,801	1,572	589	97	1,135
HQC	2,249	56	4,497	87	204	362
mceliece348864f	261,120	6,492	96	35,978	38	128
Kyber-512	800	32	768	123	155	289

Performance

NIST Security Category 3, taken from algorithm specifications

- BIKE – AVX512
- HQC AVX2 optimized
- Classic McEliece AVX

ML-KEM performance included for comparison

	Public key size (bytes)	Private key (bytes)	Ciphertext size (bytes)	KeyGen (kilocycles)	Encaps (kilocycles)	Decaps (kilocycles)
BIKE	3,082	418	3,114	1,823	223	3,887
HQC	4,522	64	9,042	204	465	755
mceliece460896f	524,160	13,608	156	117,301	81	264
Kyber-768	1,184	32	1,088	213	249	275

Performance

How does performance change when your algorithm is implemented in constant time?

Use cases

- Where do you envision your algorithm being used?
- For what use cases is your algorithm the best fit?
- For what use cases would your algorithm be ill-equipped?

Other considerations

- What benefit would there be to standardizing your KEM in addition to ML-KEM?
- Questions from the audience

Recent Attacks on BIKE

1. [Guo, Hlauschek, Johansson, Lahr, Nilsson, Schröder , CHES 2022]
   Don't Reject This: Key-Recovery Timing Attacks Due to Rejection-Sampling in HQC and BIKE

   → New sampler without rejection

2. [Wang, Wang, Wang, Crypto 2023]
   Exploring Decryption Failures of BIKE: New Class of Weak Keys and Key Recovery Attacks

   → Public-key binding as protection against multi-target key attacks

   → New decoder with reduced failure probability

   → Connections between distance spectrum multiplicities and failures

   → Preliminary results on error floor modeling

BIKE DFR Estimates

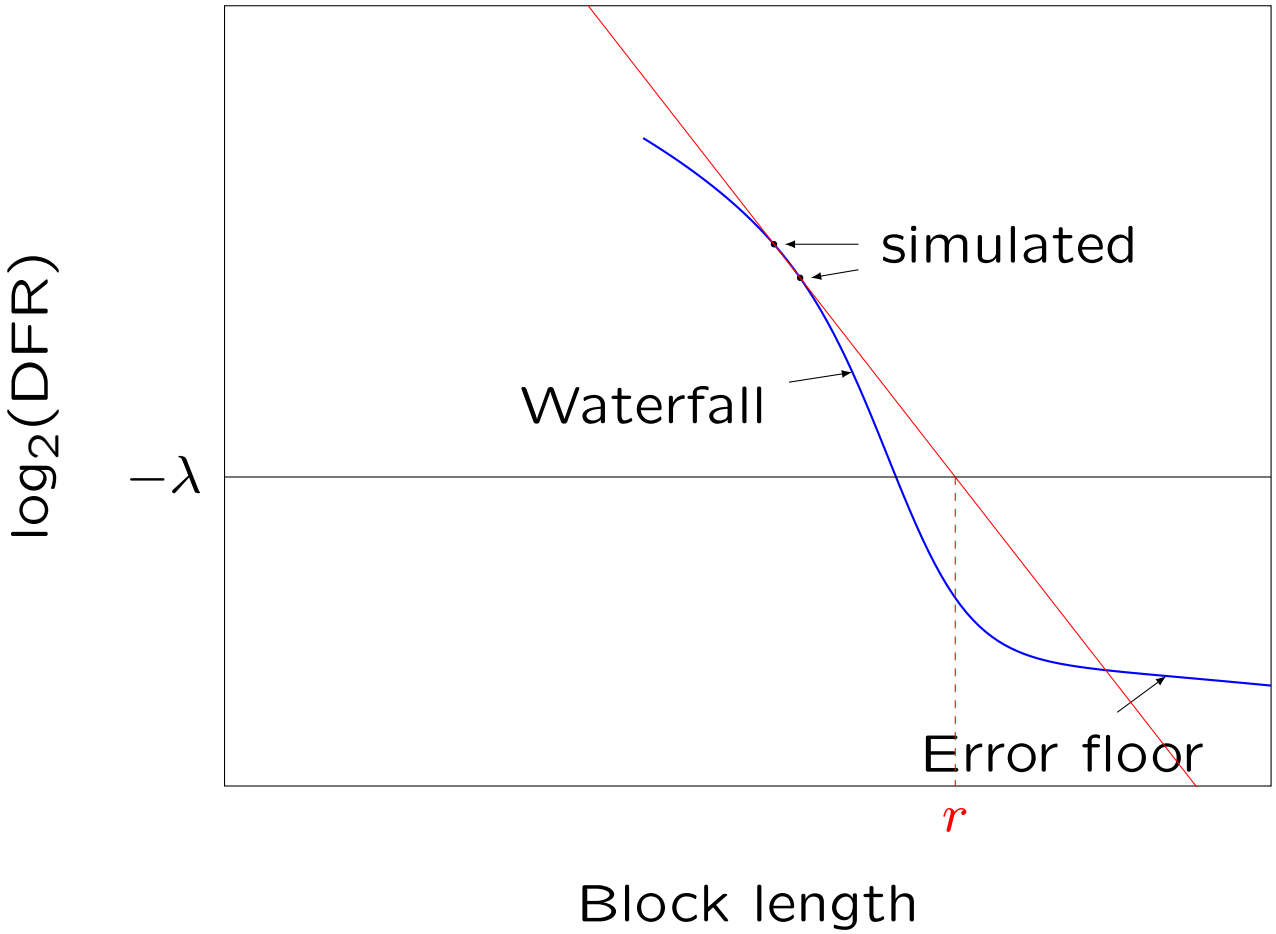

Distance Spectrum and Multiplicity

[Guo, Johansson, Stankovski, Asiacrypt 2016]

Coordinate distance

$$\mathsf{d}(i,j) = \min(i-j \bmod r, j-i \bmod r), \quad 0 \le \mathsf{d}(i,j) \le \lfloor r/2 \rfloor,$$

Distance spectrum

$$\mathsf{Sp}(h) = \{\mathsf{d}(i,j) \mid h_i = h_j = 1\}.$$

Multiplicity

$$\mu(\delta, h) = \left|\{(i,j) \mid h_i = h_j = 1, 0 \le i \le j < r, \mathsf{d}(i,j) = \delta\}\right|.$$

m-gathering Weak Keys

[Wang, Wang, Wang, Crypto 2023]

h_0 and h_1 of weight d in $\mathbf{F}_2[x]/(x^r - 1)$ a BIKE secret key

$\epsilon \in \{0, 1\}$

Generalize by applying the isomorphism $x \mapsto x^i$, $0 < i < r/2$

e.g. $(m, \epsilon) = (4000, 1)$, density is $2^{-87.28}$ and the DFR is $2^{-29.33}$

$\rightarrow$ contribution to the average DFR is $\geq 2^{-117}$

$\rightarrow$ successful attack with 2^{117} Decaps queries

Distance Spectrum Cardinality of m-gathering Keys

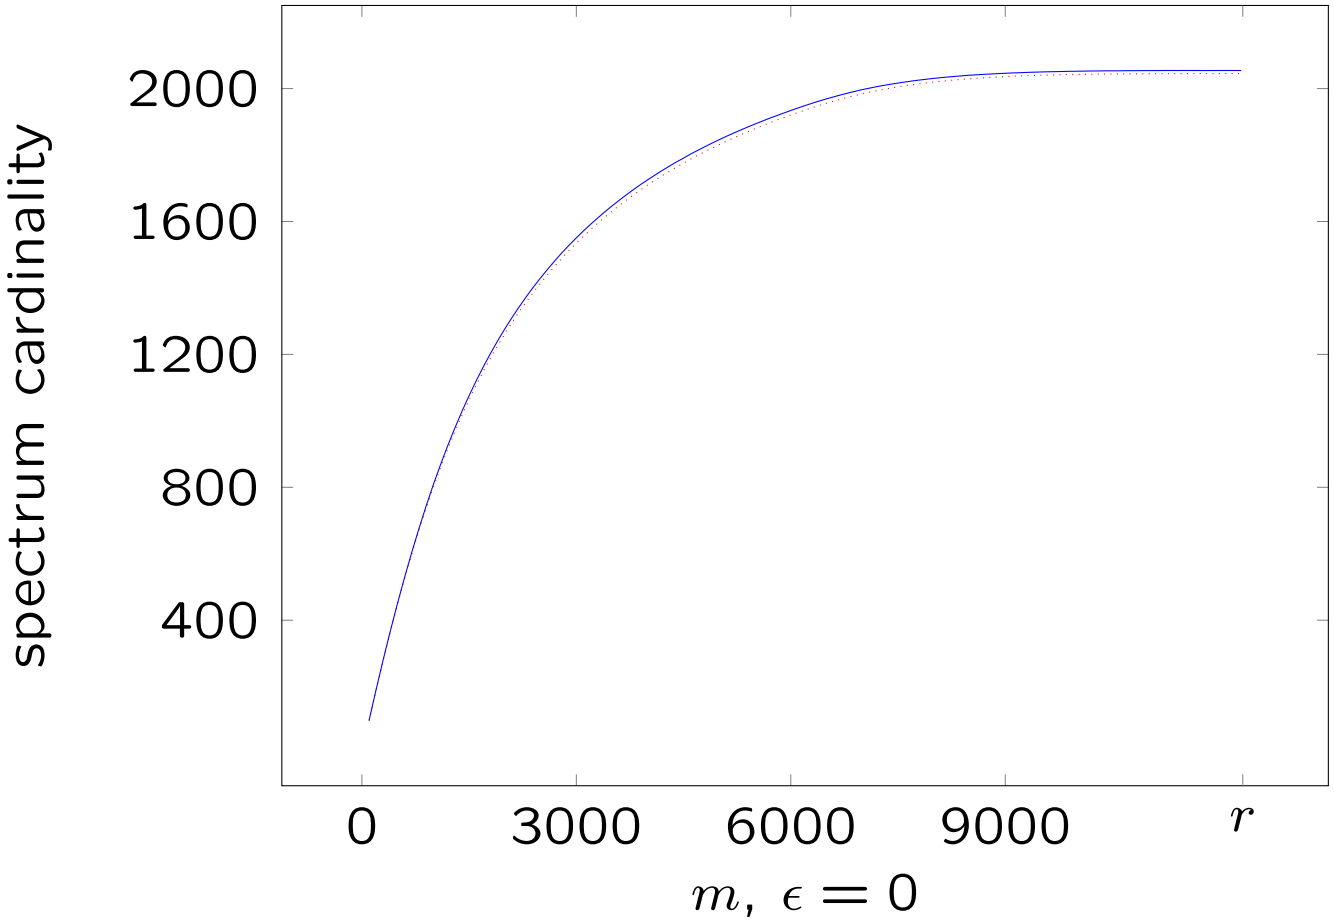

Spectrum Cardinality vs DFR for m-gathering Keys

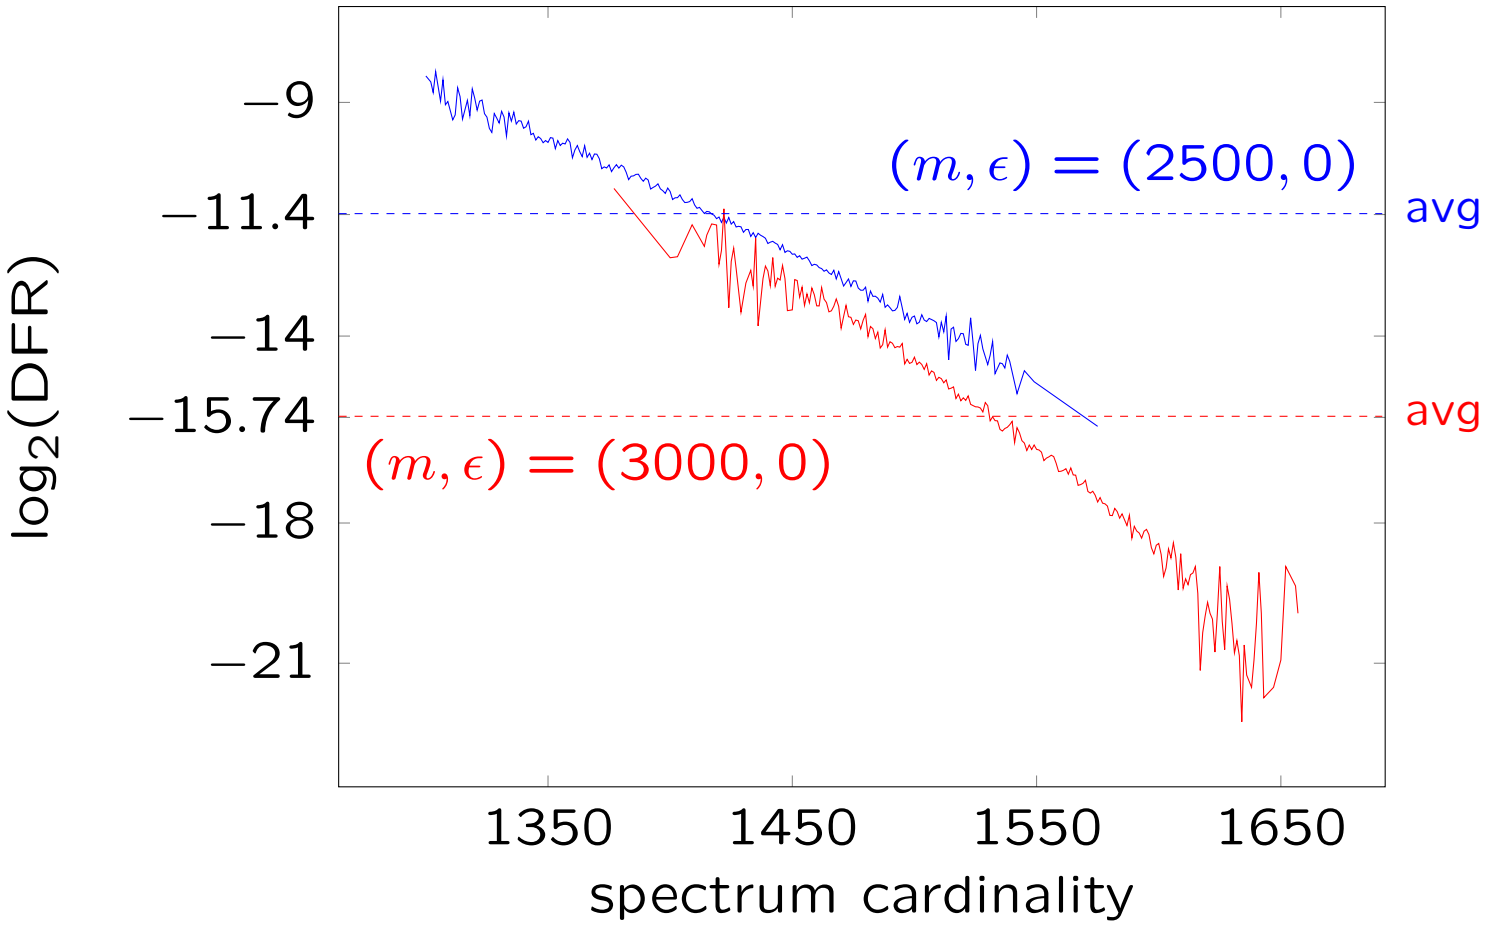

Distance Mutiplicity for m-gathering Keys

	(m, ϵ)				
μ	$(2000, 0)$	$(3000, 0)$	$(4000, 0)$	$(6000, 0)$	uniform
1	571.0	893.1	1 135.5	1 459.2	1 660.4
2	364.1	419.1	418.1	373.6	334.8
3	192.6	157.2	120.9	73.92	44.98
4	85.99	48.77	28.65	11.87	4.53
5	33.01	12.84	5.73	1.60	0.36
6	11.07	2.93	0.99	0.19	0.024
7	3.29	0.59	0.15	0.019	0.0014
8	0.88	0.11	0.020	0.0017	0.000071
$\mu > 0$	1262.2	1534.7	1710.1	1920.4	2045.1

Expected number of distances of multiplicity μ for various gathering parameters

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 $i = 1, \ldots,$ NbIter **do**
    $T \leftarrow \text{THRESHOLD}(i, s, \tilde{s})$
    for $j = 0, \ldots, n-1$ **do**
      $\sigma_j \leftarrow \texttt{ctr}(H, \tilde{s}, j)$
    for $j = 0, \ldots, n-1$ **do**
      if $\sigma_j \geq T$ **then**
        $\tilde{e}_j \leftarrow \tilde{e}_j \oplus 1$
        $\tilde{s} \leftarrow \tilde{s} - \text{col}(H, j)$
  return $\tilde{e}$

function $\text{THRESHOLD}(i, s, \tilde{s})$
  $T' \leftarrow f_t(|s|)$ ▷ 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)

NbIter $= 7$ (level 1)
$\texttt{ctr}(H, \tilde{s}, j)$ number of unsatisfied equations involving position j

New Decoder DFR – Waterfall

r	10 620	10 650	10 680	10 700
#samples	$4.4 \cdot 10^9$	$13.3 \cdot 10^9$	$56.2 \cdot 10^9$	$38.2 \cdot 10^9$
#failures	16 222	7 756	3 183	870
$\log_2$(DFR)	-18.04	-20.71	-23.46	-25.39

"Waterfall DFR extrapolation":

- $r = 12\,323 \rightarrow 2^{-181}$
- $r = 11\,768 \rightarrow 2^{-128}$

New Decoder DFR – Weak Keys (1/2)

(m, ϵ)	$(1600, 0)$	$(1700, 0)$	$(1800, 0)$	$(1900, 0)$	$(2000, 0)$	$(2100, 0)$
#samples	$9.5 \cdot 10^9$	$9.6 \cdot 10^9$	$9.6 \cdot 10^9$	$9.6 \cdot 10^9$	$15.7 \cdot 10^9$	$9.4 \cdot 10^9$
#failures	79913	32596	13153	5293	3383	763
$\log_2$(DFR)	-16.86	-18.16	-19.48	-20.79	-22.15	-23.56
$\log_2$(density)	-188.41	-182.15	-176.26	-170.69	-165.41	-160.40
$\log_2$(sum)	-205.27	-200.31	-195.74	-191.48	-187.56	-183.96

(m, ϵ)	$(2200, 0)$	$(2300, 0)$	$(2400, 0)$	$(2500, 0)$	$(2600, 0)$	$(2700, 0)$
#samples	$9.4 \cdot 10^9$	$15.8 \cdot 10^9$	$19.1 \cdot 10^9$	$25.5 \cdot 10^9$	$59.8 \cdot 10^9$	$7.9 \cdot 10^9$
#failures	270	177	81	40	31	0
$\log_2$(DFR)	-25.06	-26.41	-27.81	-29.25	-30.84	–
$\log_2$(density)	-155.62	-151.06	-146.70	-142.51	-138.50	-134.63
$\log_2$(sum)	-180.68	-177.47	-174.51	-171.76	-169.34	–

New Decoder DFR – Weak Keys (2/2)

(m, ϵ)	$(1600, 1)$	$(1700, 1)$	$(1800, 1)$	$(1900, 1)$	$(2000, 1)$
#samples	$9.3 \cdot 10^9$	$9.3 \cdot 10^9$	$9.3 \cdot 10^9$	$9.3 \cdot 10^9$	$13.6 \cdot 10^9$
#failures	17916	6904	2524	921	482
$\log_2$(DFR)	-18.98	-20.36	-21.82	-23.27	-24.75
$\log_2$(density)	-179.47	-173.31	-167.52	-162.05	-156.86
$\log_2$(sum)	-198.45	-193.67	-189.34	-185.31	-181.61

(m, ϵ)	$(2100, 1)$	$(2200, 1)$	$(2300, 1)$	$(2400, 1)$	$(2500, 1)$
#samples	$10.1 \cdot 10^9$	$20.1 \cdot 10^9$	$25.2 \cdot 10^9$	$25.3 \cdot 10^9$	$59.9 \cdot 10^9$
#failures	132	94	47	16	12
$\log_2$(DFR)	-26.19	-27.68	-29.00	-30.56	-32.22
$\log_2$(density)	-151.93	-147.24	-142.76	-138.47	-134.36
$\log_2$(sum)	-178.12	-174.91	-171.76	-169.03	-166.58

New Decoder DFR – m-gathering Simulation (1/2)

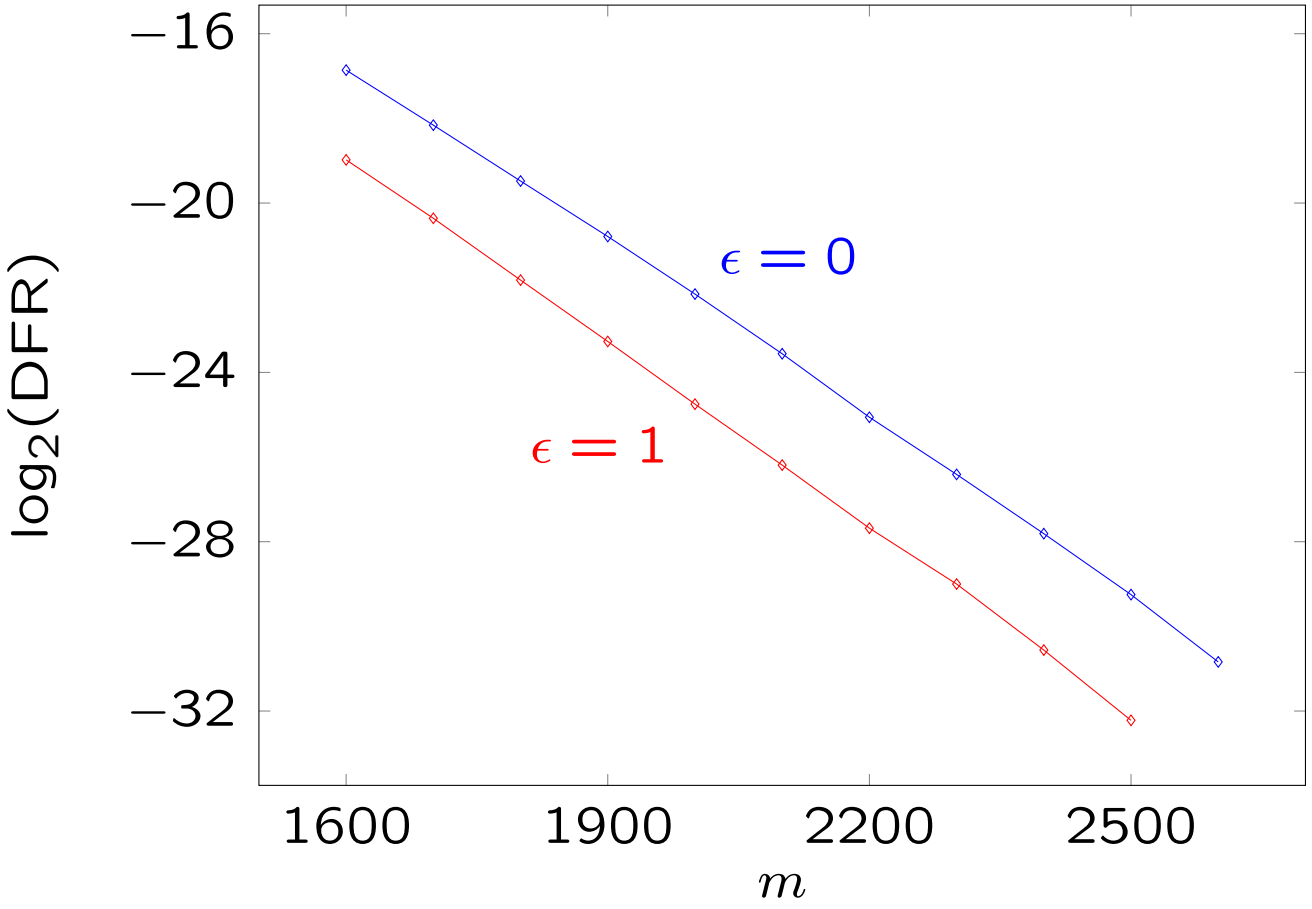

With the current decoder $(m, \epsilon) = (4000, 1) \rightarrow \text{DFR} = 2^{-29.33}$

New Decoder DFR – m-gathering Simulation (2/2)

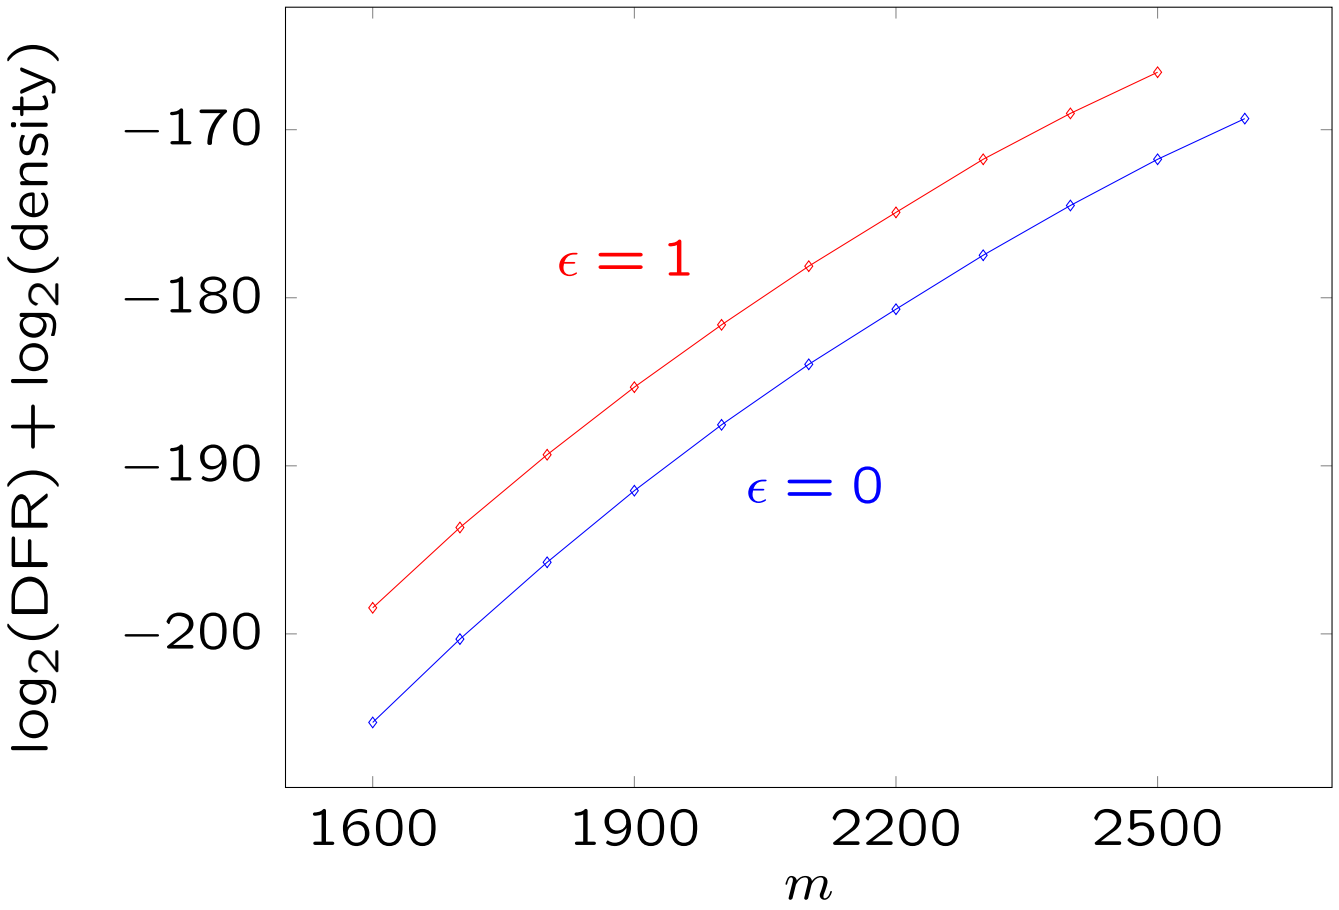

Preliminary Results on Error Floor Modeling

[Tillich, Vasseur, ongoing]

- Markovian model in [Sendrier, Vasseur, PQCrypto 2019]

  State = (syndrome weight, error weight)

  Predicts the waterfall but not the error floor

- The new Markovian model state includes the distance to the closest near-codeword (near-codeword as in [Vasseur, ePrint 2021/1458])

  For reduced length and perfect keys, the Markovian model matches with simulations and accurately predicts the error floor.

- Next step: adapt the formula of the transition probabilities to cover the case of all keys

(Perfect key: all distances in the spectrum have multiplicity 0 or 1)

Markovian Model Matching the Error Floor (Perfect Keys)

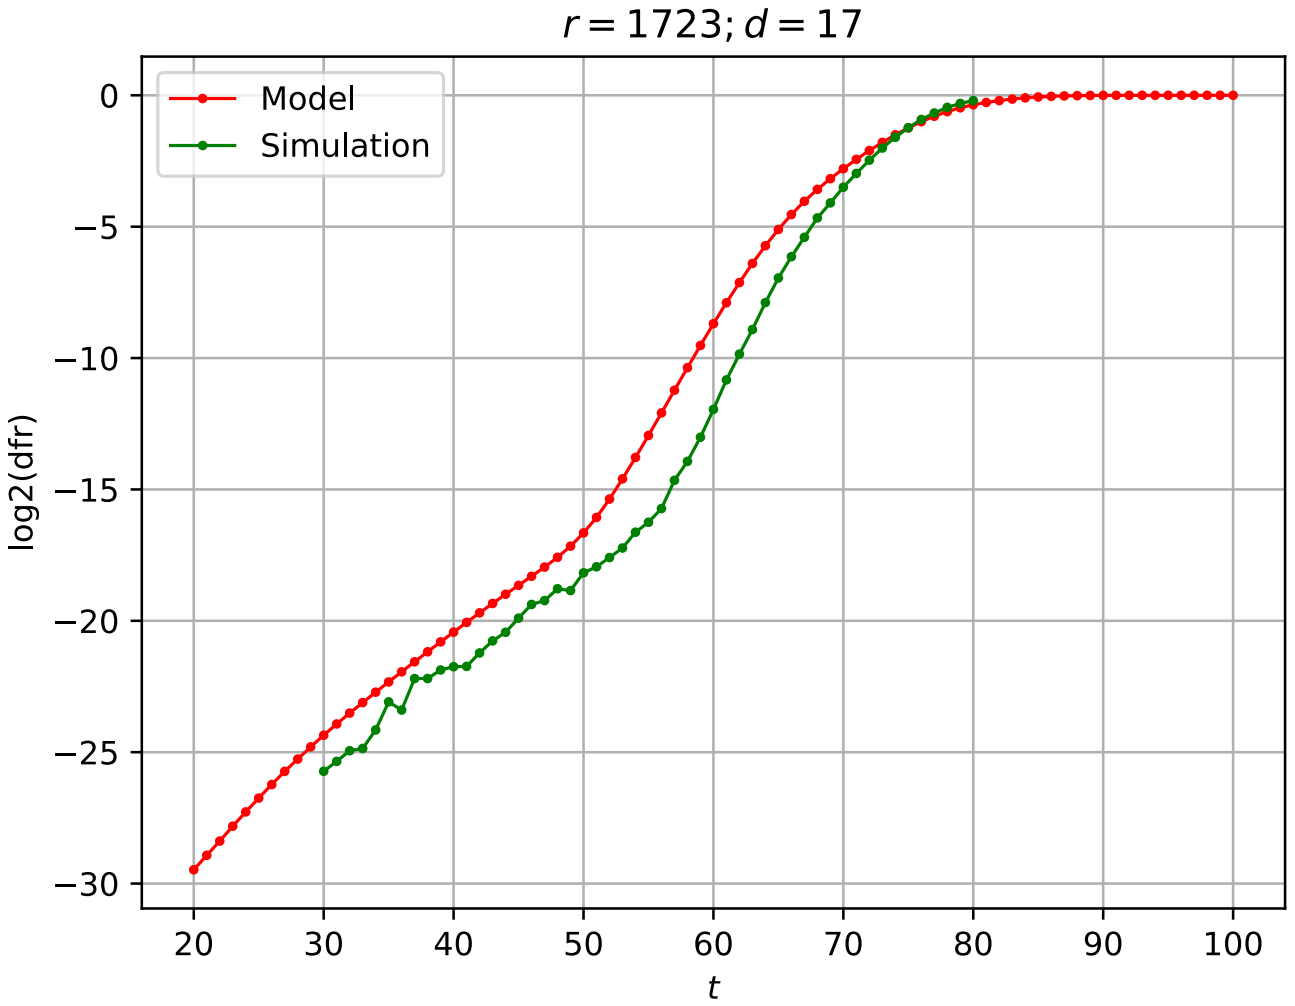

Mitigation of Weak Key Impact

Observation: known weak keys are characterized by bad distance spectrum multiplicities.

1. Our favor goes to the ongoing effort towards error floor modeling

   Weak keys behavior is hopefully captured by those models

2. Other factors mitigate the impact of weak key in BIKE
   - Ephemeral key setting unaffected by failures
   - New threshold schedule makes weak key attacks less effective
   - Successful key attack scenarios require $> 2^{64}$ decapsulation queries
   - If it comes to that, filtering out weak keys is an easy task (distance spectrum with multiplicity is easy to compute)