Improved Codes and Decoders for HQC

Sebastian Bitzer[1], Bharath Purtipli[1*], Antonia Wachter-Zeh[1]

[1]Technical University of Munich, Munich, Germany.

*Corresponding author(s). E-mail(s): bharath.purtipli@tum.de;
Contributing authors: sebastian.bitzer@tum.de;
antonia.wachter-zeh@tum.de;

Abstract

The Hamming Quasi-Cyclic (HQC) cryptosystem, recently selected by NIST for standardization, uses a concatenation of Reed–Muller (RM) and Reed–Solomon (RS) codes to ensure correct decryption. This work optimizes the underlying code and decoder with the goal of reducing ciphertext and public-key sizes. First, we propose a two-level generalized concatenated code (GCC), exploiting the observation that the inner RM code contains a repetition subcode. The repetition layer incurs a smaller inner failure probability and, therefore, admits a higher code rate than the RS code. Second, we develop a reliability-based decoding framework for the outer code. The inner decoder provides a reliability measure whose distribution, jointly with the decoding outcome, is bounded analytically. These bounds allow deriving conservative DFR guarantees for two erasure-assignment strategies (threshold-based and partition-based). Combining both contributions, we observe that public-key and ciphertext sizes can be reduced by up to 4.34 %.

Keywords: HQC, Code-based cryptography, Generalized Code Concatenation, Errors-and-Erasures Decoding

1 Introduction

Hamming Quasi-Cyclic (HQC) [1] is a code-based key-encapsulation mechanism whose security relies on the hardness of decoding quasi-cyclic codes in the Hamming metric. Following several rounds of evaluation, HQC was selected by NIST as the fifth post-quantum algorithm [2]. To encrypt a message $\mathbf{m} \in \mathbb{F}_2^{\lambda}$, it is first encoded via an efficient error-correcting code $\mathcal{C}_{\text{pub}}$ before applying a random-looking mask. During decryption, the mask is removed using the secret key, but a residual error $\mathbf{e}$ of moderate Hamming weight remains. To achieve IND-CCA2 security and thwart attacks based on

decryption failures [3, 4], the decoder of $\mathcal{C}_{\text{pub}}$ must correct $\mathbf{e}$ with probability $1 - 2^{-\lambda}$, where λ denotes the bit security level (e.g., $\lambda = 128$ for NIST security level 1).

The Decryption Failure Rate (DFR) analysis of [1] models the HQC error $\mathbf{e}$ as a vector sampled from a Bernoulli distribution with an error probability p, which is decreasing in the code length n. The public-key size is n bits, and the ciphertext size is approximately $2n$ bits, so reducing n directly reduces both. Under this model, an analytical DFR bound for a hard-decision decoder of a concatenated code was derived in [1]. Initially, HQC used the concatenation of an inner repetition and an outer BCH code [5]. In the current version, this was replaced by an inner (duplicated) Reed–Muller and outer Reed–Solomon code, resulting in shorter codes and better runtime.

1.1 Contributions

We propose two complementary techniques to reduce HQC's public-key and ciphertext sizes, achieving a reduction of up to 4.34 %.

Generalized concatenated codes. We observe that the inner RM code currently used in HQC contains the repetition code as a subcode, forming a natural partition chain (Section 3). This motivates a 2-level generalized concatenated code (GCC) [6, 7], which can be interpreted as combining the round-one (repetition-BCH) and round-two (RM-RS) constructions. The repetition layer has twice the minimum distance of the RM layer and therefore incurs a much smaller inner failure probability, admitting a higher-rate outer code. The additional information bits carried at this level allow shortening the overall code while maintaining the required DFR and message length. Section 3.1 provides a general overview of GCCs and their DFR analysis [8]; Section 3.2 instantiates the framework for HQC with a concrete 2-level construction.

Reliability-based outer decoding. The current HQC specification [1] uses a hard-decision outer decoder, which does not exploit the confidence information available from the inner decoding. We develop a framework for errors-and-erasures outer decoding, in which positions with low inner-decoding confidence are declared as erasures (Section 4). Since an errors-and-erasures decoder can correct e errors and s erasures whenever $2e + s < d$, converting suspected errors into erasures effectively doubles their correction efficiency.

The framework is based on a reliability measure derived from the inner code's geometry. We bound its distribution jointly with the decoding outcome (Section 4.1), and use these bounds to derive rigorous DFR guarantees for two erasure-assignment strategies: a *threshold-based* scheme, where all symbols with reliability below a fixed threshold are erased (Section 4.2), and a *partition-based* scheme, where a fixed-size subset of the least reliable symbols is erased (Section 4.3). The DFR analysis uses stochastic dominance and order-statistic techniques, respectively. Additionally, we extend the DFR analysis for the threshold-based scheme to use any error weight distribution.

Reproducibility. Scripts for reproducing computations and simulations are available at https://gitlab.com/bharath.purtipli/hqc_gcc.

1.2 Related Work

Aguilar-Melchor et al. [9] introduced the current RM-RS concatenated coding scheme for HQC, replacing the earlier BCH-repetition construction. They provide a detailed DFR analysis of the duplicated first-order RM code under hard-decision RS decoding. Our work builds on their construction by extending it to a GCC and by adding reliability-based outer decoding.
HARE [10] is a recently proposed adaptation of the HQC standard that achieves significant reductions in ciphertext and public-key sizes through ciphertext compression, unbalanced encryption weights, and an errors-and-erasures RS decoder. Their errors-and-erasures decoder uses the same reliability measure and threshold-based erasure assignment as Section 4.2 of this work, with a DFR analysis specific to that setting. In contrast, we develop a general analytical framework for the reliability measure (Section 4.1) that supports multiple erasure strategies (Sections 4.2 and 4.3) and additionally propose a GCC structure (Section 3).
GMD decoding of HQC's RS outer code was proposed in [11]. While their approach is well-motivated, the DFR bounds rely on extrapolation from simulation data, which may not yield valid guarantees for DFRs of $2^{-\lambda}$. In contrast, our DFR analysis is fully analytical and yields provable upper bounds.

2 Preliminaries

2.1 Notation

We write $[n] \triangleq \{1, \ldots, n\}$. Integers and field elements are denoted by lower-case Roman letters, vectors by bold lower-case, and matrices by bold upper-case. Denote by $\mathbb{F}_q$ the finite field with q elements. For $\mathbf{u}, \mathbf{v} \in \mathbb{F}_q^n$, the Hamming distance is $\mathrm{d_H}(\mathbf{u}, \mathbf{v}) \triangleq |\{i \in [n] : u_i \neq v_i\}|$ and the Hamming weight is $|\mathbf{v}|_{\mathrm{H}} \triangleq \mathrm{d_H}(\mathbf{v}, \mathbf{0})$. Let $\mathcal{R}_n \triangleq \mathbb{F}_2[x]/(x^n - 1)$; for $\mathbf{h}, \mathbf{y} \in \mathcal{R}_n$, $\mathbf{h} \cdot \mathbf{y}$ denotes the product in $\mathcal{R}_n$.

2.2 Codes and Decoders

A *linear code* $\mathcal{C}$ with parameters $[n, k, d]_q$ is a k-dimensional subspace of $\mathbb{F}_q^n$ with minimum Hamming distance d. For given $\mathcal{C}$, we also write $k(\mathcal{C})$ to refer to its dimension and $d(\mathcal{C})$ to refer to its minimum distance. An *encoder* $\mathsf{enc}(\cdot)$ for $\mathcal{C}$ is a bijective linear map $\mathbb{F}_q^k \to \mathcal{C}$, characterized by a generator matrix $\mathbf{G} \in \mathbb{F}_q^{k \times n}$ with $\mathsf{enc}(\mathbf{m}) = \mathbf{mG}$.
Decoders. A *decoder* $\mathsf{dec}(\cdot)$ maps a received vector $\mathbf{v} \in \mathbb{F}_q^n$ to a codeword $\hat{\mathbf{c}} \in \mathcal{C}$. A *maximum-likelihood (ML)* decoder returns $\hat{\mathbf{c}} = \arg\min_{\mathbf{c} \in \mathcal{C}} \mathrm{d_H}(\mathbf{v}, \mathbf{c})$, with ties broken by a fixed deterministic rule. A *bounded-minimum-distance (BMD)* decoder successfully decodes any error pattern with $|\mathbf{e}|_{\mathrm{H}} \leq \lfloor (d-1)/2 \rfloor$.

An *errors-and-erasures decoder* accepts input $\mathbf{v} \in (\mathbb{F}_q \cup \{\Delta\})^n$, where Δ denotes an erasure (unknown symbol value with known position). A BMD errors-and-erasures decoder succeeds whenever $2e + s < d$, where s is the number of erasures and e the number of errors among the non-erased positions.
Reed–Solomon (RS) codes. For any prime power $q \geq n$, there exists an RS code with parameters $[n, k, n - k + 1]_q$. RS codes are maximum-distance separable (MDS) and admit efficient BMD errors-and-erasures decoding [12].

Reed–Muller (RM) codes. The r-th order binary Reed–Muller code $\mathrm{RM}(r,m)$ has length 2^m, dimension $\sum_{i=0}^{r}\binom{m}{i}$, and minimum distance 2^{m-r}. For $r=0$, this is the binary repetition code $\mathrm{Rep} \triangleq \{(0,\ldots,0),(1,\ldots,1)\}$ of length 2^m. For $r=1$, ML decoding can be performed efficiently via the Fast Walsh–Hadamard Transform (FWHT) [13].

Duplicated RM codes. Following [1, 9], we use the tensor product $\mathrm{Rep}\otimes\mathrm{RM}(r,m)$, which repeats every bit of an $\mathrm{RM}(r,m)$ codeword δ times. For duplication factor δ, the resulting code has parameters $[\delta\cdot 2^m,\ \sum_{i=0}^{r}\binom{m}{i},\ \delta\cdot 2^{m-r}]_2$. ML decoding first accumulates the δ copies of each bit, then applies the FWHT to decode the RM component. A bound on the ML failure probability is provided in [1, 9].

2.3 HQC in a Nutshell

We provide a brief overview of the HQC public-key encryption scheme, focusing on the error-correction mechanism that forms the center of this work. For cryptographic details, the interested reader is referred to [1].

HQC uses a linear $[n_{\mathrm{out}}n_{\mathrm{in}}, k=\lambda, d_{\mathrm{pub}}]_2$ code $\mathcal{C}_{\mathrm{pub}}$ with $n_{\mathrm{out}}n_{\mathrm{in}} \le n$. The secret key is a random pair of moderately sparse vectors $\mathbf{x},\mathbf{y}\in\mathcal{R}_n$ each of Hamming weight $w_{\mathtt{sk}}$. The corresponding public key is the pair of vectors $\mathbf{h}$ and $\mathbf{s}=\mathbf{x}+\mathbf{h}\cdot\mathbf{y}$.

Encryption. To encrypt a message $\mathbf{m}\in\mathbb{F}_2^{\lambda}$, one creates a pair of vectors $\mathbf{u}=\mathbf{r}_1+\mathbf{h}\cdot\mathbf{r}_2$ and $\mathbf{v}=\mathcal{C}_{\mathrm{pub}}.\mathsf{enc}(\mathbf{m})+(\mathbf{s}\cdot\mathbf{r}_2+\mathbf{r}_3)_{[n_{\mathrm{out}}n_{\mathrm{in}}]}$ where $\mathbf{r}_1,\mathbf{r}_2,\mathbf{r}_3\in\mathcal{R}_n$ are random vectors of weight $w_{\mathtt{ct}}$.

Decryption. Decryption applies an efficient decoder of $\mathcal{C}_{\mathrm{pub}}$ on $\mathbf{v}+\mathbf{u}\cdot\mathbf{y} = \mathcal{C}_{\mathrm{pub}}.\mathsf{enc}(\mathbf{m})+\mathbf{e}_{[n_{\mathrm{out}}n_{\mathrm{in}}]}$ where the residual error is the truncation of the HQC error $\mathbf{e}=\mathbf{x}\cdot\mathbf{r}_2+\mathbf{y}\cdot\mathbf{r}_1+\mathbf{r}_3$. The message $\mathbf{m}$ is recovered by applying the decoder of $\mathcal{C}_{\mathrm{pub}}$. Since $\mathbf{e}$ is random, decryption may fail; the code parameters are chosen so that this *decryption failure rate* (DFR) is at most $2^{-\lambda}$.

Error model. The error $\mathbf{e}$ can be modeled as having i.i.d. $\mathrm{Bern}(p)$ entries, i.e., errors occur independently with probability p. As shown in [1], p depends on the system parameters via

$$p=\Pr\left[\mathbf{e}_i=1\right]=2(1-\tilde{p})\tilde{p}\left(1-\tfrac{w_{\mathtt{ct}}}{n}\right)+\left((1-\tilde{p})^2+\tilde{p}^2\right)\frac{w_{\mathtt{ct}}}{n},$$

$$\text{where } \tilde{p}=\binom{n}{w_{\mathtt{ct}}}^{-1}\cdot\sum_{\substack{1\le\ell\le\min\{w_{\mathtt{sk}},w_{\mathtt{ct}}\}\\ \ell\text{ odd}}}\binom{w_{\mathtt{sk}}}{\ell}\binom{n-w_{\mathtt{sk}}}{w_{\mathtt{ct}}-\ell}.$$

The crossover probability p is inversely decreasing in n: longer codes operate at a lower noise level.

Alternative error models that account for dependencies in the error $\mathbf{e}$ have been studied in [14–16]. While most of this work relies on the independence assumption described above, we also use the correlated-hypergeometric model of [14, Model 2].

Definition 1 (Model 2 [14]) Let $\mathbf{t'_1},\mathbf{t'_2}\sim\mathrm{Hypergeometric}(n,w_{\mathtt{sk}},w_{\mathtt{ct}})^n$, conditioned on $\sum_{i=1}^{n}t'_{1,i}=\sum_{i=1}^{n}t'_{2,i}=w_{\mathtt{sk}}\cdot w_{\mathtt{ct}}$. Let $\mathbf{t_1}=\mathbf{t'_1} \bmod 2$, $\mathbf{t_2}=\mathbf{t'_2} \bmod 2$, and $|\mathbf{r}_3|_{\mathrm{H}}=w_{\mathtt{ct}}$.

Model 2, the correlated-hypergeometric assumption, models the HQC error $\mathbf{e}$ as

$$\mathbf{t_1} + \mathbf{t_2} + \mathbf{r}_3 \,.$$

3 More Efficient Error-Correcting Code

This section provides an overview on generalized concatenated codes (Section 3.1) and describes the code construction tailored to HQC (Section 3.2).

3.1 An Overview on Generalized Concatenated Codes

Zinoviev [6] and Blokh and Zyablov [7] provided a generalization of concatenated codes to several levels: whereas standard concatenation pairs a single outer code with a single inner code, a GCC associates a separate outer code with each level of an inner code partition chain. For a thorough introduction to generalized code concatenation, also known as multi-level concatenation, see [17] or [12], and references therein.

To construct a binary s-level GCC, select a nested sequence of inner codes

$$\{\mathbf{0}\} = \mathcal{B}^{(s+1)} \subseteq \mathcal{B}^{(s)} \subseteq \cdots \subseteq \mathcal{B}^{(1)} \subseteq \mathbb{F}_2^{n_{\text{in}}}.$$

Note that the indexing convention has $\mathcal{B}^{(1)}$ as the *largest* (and weakest) inner code and $\mathcal{B}^{(s)}$ as the smallest (and strongest). Correspondingly, the inner minimum distance $d(\mathcal{B}^{(j)})$ is non-decreasing in j. By abuse of notation, $\mathcal{B}^{(j)}/\mathcal{B}^{(j+1)} \subset \mathbb{F}_q^{n_{\text{in}}}$ denotes the (not necessarily unique) linear code of dimension $k(\mathcal{B}^{(j)}) - k(\mathcal{B}^{(j+1)})$ such that

$$\mathcal{B}^{(j)} = \{\mathbf{c}_1 + \mathbf{c}_2 \mid \mathbf{c}_1 \in \mathcal{B}^{(j+1)}, \mathbf{c}_2 \in \mathcal{B}^{(j)}/\mathcal{B}^{(j+1)}\}.$$

Before introducing GCCs using this notation, we illustrate it by partitioning the inner code used in HQC's code concatenation [1].

Example 1 Let $\mathcal{B}^{(1)}$ be a duplicated first-order RM code of dimension k_{RM} as used in [1]. $\mathcal{B}^{(1)}$ is generated by $\langle \mathbf{g}_1, \ldots, \mathbf{g}_{k_{\text{RM}}} \rangle$ with $\mathbf{g}_1 = (1, \ldots, 1)$. Hence, $\mathcal{B}^{(1)}$ contains the length-n_{in} repetition code as a subcode $\mathcal{B}^{(2)}$. Then, we may select $\mathcal{B}^{(1)}/\mathcal{B}^{(2)} = \langle \mathbf{g}_2, \ldots, \mathbf{g}_{k_{\text{RM}}} \rangle$.

Definition 2 Let a nested sequence of inner codes $\{\mathbf{0}\} = \mathcal{B}^{(s+1)} \subseteq \mathcal{B}^{(s)} \subseteq \cdots \subseteq \mathcal{B}^{(1)} \subseteq \mathbb{F}_2^{n_{\text{in}}}$ be given. Let outer codes $\mathcal{A}_j \subseteq \mathbb{F}_{2^{m_j}}^{n_{\text{out}}}$ with $m_j = k(\mathcal{B}^{(j)}) - k(\mathcal{B}^{(j+1)})$ be given for all $j \in [s]$. The *generalized concatenated code* GCC is constructed as

$$\text{GCC} = \left\{ \left(\sum_{j \in [s]} (\mathcal{B}^{(j)}/\mathcal{B}^{(j+1)}).\mathsf{enc}(a_{j,\ell}) \right)_{\ell \in [n_{\text{out}}]} \,\middle|\, \mathbf{a}_j \in \mathcal{A}_j \ \forall j \in [s] \right\}.$$

Then, GCC has length $n_{\text{out}} \cdot n_{\text{in}}$ and $\mathbb{F}_2$-dimension $\sum_{j \in [s]} m_j \cdot k(\mathcal{A}_j)$.

The following example highlights that Definition 2 specializes to a standard concatenated code (tensor product) for $s = 1$, as currently deployed in HQC [1].

Example 2 To construct the concatenated code proposed in [1] for NIST security category 1, we set $s = 1$. As $\mathcal{B}^{(1)}$, pick a first-order duplicated RM code with parameters $[384, 8, 192]_2$. As $\mathcal{A}_1$, a RS code with parameters $[46, 16, 31]_{2^8}$ is used, with alphabet $\mathbb{F}_{2^8}$ matching the inner code dimension. To construct the error-correcting codes at NIST security categories 3 and 5, the length of the inner RM code and the parameters of the outer RS code are increased.

The minimum distance of a GCC is lower-bounded as

$$d(\mathrm{GCC}) \geq \min_{j \in [s]} d(\mathcal{A}_j) \cdot d(\mathcal{B}^{(j)}), \tag{1}$$

As the inner minimum distance $d(\mathcal{B}^{(j)})$ is increasing in j, it is natural to have $d(\mathcal{A}_j)$ decrease with j (allowing higher-dimensional outer codes).

In the following, we recall the standard decoding algorithm for GCCs. Instantiated with suitable decoders for the component codes, it is guaranteed to correct all errors up to half the bound in Equation (1).

GCC decoder. Let $\mathbf{r} = (\mathbf{r}_1, \ldots, \mathbf{r}_{n_{\mathrm{out}}}) = \mathbf{c} + \mathbf{e}$ be given with codeword $\mathbf{c} \in \mathrm{GCC}$ and error $\mathbf{e}$ (cf. Section 2.3). The decoder proceeds level by level in a multistage fashion. Instantiate $\mathbf{r}^{(1)} = \mathbf{r}$. Then, for each level $j \in [s]$ proceed as follows:

1. **Inner decoding at level j.** For each $i \in [n_{\mathrm{out}}]$, a decoder of $\mathcal{B}^{(j)}$ is applied to each $\mathbf{r}_i^{(j)}$ to obtain an estimated inner codeword $\hat{\mathbf{c}}_i \in \mathcal{B}^{(j)}$. Optionally (see Section 4), a reliability measure α_i indicating the confidence of the inner decoding may be derived from the geometry of $\mathbf{r}_i^{(j)}$ and $\mathcal{B}^{(j)}$.
2. **Outer decoding at level j.** Each $\hat{\mathbf{c}}_i \in \mathcal{B}^{(j)}$ corresponds to an estimated symbol of the outer code $\mathcal{A}_j$. To correct errors in this sequence of estimated symbols, a decoder of $\mathcal{A}_j$ is applied, which may utilize the reliabilities α_i. This yields an estimated outer codeword $\hat{\mathbf{a}}_j$ (or possibly a decoding failure).
3. **Cancellation.** The estimated codeword $\hat{\mathbf{a}}_j$ is re-encoded as

$$\hat{\mathbf{c}}^{(j)} = \left((\mathcal{B}^{(j)} / \mathcal{B}^{(j+1)}).\mathsf{enc}(\hat{a}_{j,\ell}) \right)_{\ell \in [n_{\mathrm{out}}]}.$$

Before continuing with level $j+1$, $\hat{\mathbf{c}}^{(j)}$ is canceled from $\mathbf{r}^{(j)}$ via $\mathbf{r}^{(j+1)} = \mathbf{r}^{(j)} - \hat{\mathbf{c}}^{(j)}$.

DFR analysis The described decoder fails if and only if the outer decoding fails on at least one level. Denote by $E^{(j)} = \{\hat{\mathbf{a}}_j \neq \mathbf{a}_j\}$ the decoding failure event at level j. Then, using the union bound [8], an upper bound for the DFR is obtained as

$$\mathrm{DFR} = \Pr\left[\bigcup_{j \in [s]} E^{(j)} \right] \leq \sum_{j \in [s]} \Pr\left[E^{(j)} \right]. \tag{2}$$

This allows analyzing the failure probability at each level separately. Parameters are chosen such that $\sum_{j \in [s]} \Pr\left[E^{(j)}\right] \leq 2^{-\lambda}$. This simple possibility of analyzing the DFR makes GCCs a suitable candidate for error correction in HQC. In the following section, we propose a concrete instantiation of Definition 2 with two levels, which builds on the construction in Example 2.

3.2 A Two-Level GCC for HQC

As illustrated in Example 2, the error correcting code of HQC is currently instantiated as a concatenation of an inner RM and an outer RS code. Previously, the concatenation of an inner repetition code with an outer binary BCH code was used [5], which was, however, less efficient.

Table 1: Proposed 2-level GCC structure.

Level j	Outer Code $\mathcal{A}_j$		Inner Code $\mathcal{B}_j$	
1	RS	$[n_{\text{out}}, k_{\text{RS}}, d_{\text{RS}}]_{2^m}$	RM	$[n_{\text{in}}, k_{\text{RM}}, \frac{n_{\text{in}}}{2}]_2$
2	$\mathcal{C}$	$[n_{\text{out}}, k_{\mathcal{C}}, d_{\mathcal{C}}]_2$	Rep	$[n_{\text{in}}, 1, n_{\text{in}}]_2$

In Example 1, we have observed that the repetition code and the duplicated first-order RM code form a partition chain. This motivates the 2-level GCC construction for $\mathcal{C}_{\text{pub}}$ described in Table 1. As in [1], we pair the RM code with an outer RS code at level 1. The repetition code at level 2 requires a binary outer code; as in [5], we use a BCH code to achieve efficient hard-decision or errors-and-erasure decoding (cf. Section 4). Both the repetition and the RM code can be efficiently maximum-likelihood decoded.

Performance improvement. The key advantage of the 2-level construction is that the repetition code (level 2) has minimum distance n_{in}, which is twice the RM code's minimum distance. Consequently, the inner decoding error probability at level 2 is significantly smaller than at level 1. The corresponding binary outer code $\mathcal{C}$ therefore requires a smaller minimum distance, and can operate at a higher rate. Since the overall message length is fixed to λ, this improvement translates to a smaller RS dimension or, equivalently, a shorter overall code length. This effect reduces the size of the public key and the ciphertext.

The achievable improvement depends on the outer decoder; the following example illustrates the interplay between the GCC structure and the decoder for NIST 1 parameters.

Example 3 The current HQC specification [1] uses a single-level concatenation with an outer $[46, 16, 31]_{2^8}$ RS code and an inner $[384, 8, 192]_2$ RM code, giving $n = 17669$. At $p = 0.340$, a DFR of $2^{-132.9}$ is achieved using inner ML and outer hard-decision decoding.

Using the improved errors-and-erasures decoder of Section 4, we can reduce the outer length to $n_{\text{out}} = 44$ (and correspondingly $n = 16901$, reducing the public-key size by 4.34 %). At this shorter length, the DFR analysis requires an RS minimum distance of $d_{\text{RS}} = 30$, leaving $k_{\text{RS}} = n_{\text{out}} - d_{\text{RS}} + 1 = 15$ information symbols. For a single concatenation level, this yields $15 \cdot 8 = 120$ information bits, which is insufficient for $\lambda = 128$.

The 2-level GCC resolves this shortfall. At level $j = 2$, the repetition code has minimum distance $n_{\text{in}} = 384$, twice that of the RM code. This results in a much smaller inner failure probability, so that an outer binary BCH code with minimum distance $d_{\mathcal{C}} = 8$ suffices. A shortened $[n_{\text{out}}, 25, 8]_2$ BCH code contributes $k_{\mathcal{C}} = 25$ information bits, for a total of $15 \cdot 7 + 25 = 130 \geq 128 = \lambda$.

Refer to Section 5 for further parameter suggestions, also relying on the improved decoding presented in Section 4, which is fully compatible with the 2-level construction.

On adding a third level via RM(2,m). Higher-order RM codes are supercodes of the duplicated first-order RM code. It is therefore a natural question whether it makes sense to extend the code presented in Table 1 to $s = 3$ levels via the partition chain

$$\mathrm{Rep} \subset \mathrm{RM}(1, m) \subset \mathrm{RM}(2, m),$$

where the minimum distance of $\mathrm{RM}(2, m)$ is half the minimum distance of $\mathrm{RM}(1, m)$. To estimate the failure probability of the duplicated $\mathrm{RM}(2, m)$ under maximum-likelihood decoding, we use the union bound

$$\bar{P}_{\mathrm{in}} \triangleq \sum_{w=1}^{n_{\mathrm{in}}} A_w \sum_{i=\lceil w/2 \rceil}^{w} \binom{w}{i} p^i (1-p)^{w-i},$$

where $A_0, \ldots A_{n_{\mathrm{in}}}$ denotes the weight distribution of the duplicated code, which is known [18]. For parameters suggested in [1], we find $\bar{P}_{\mathrm{in}} > 1$. While this does not directly imply that the true error probability equals 1 (the union bound may be loose), it indicates that the inner code operates far beyond its error-correction capability at this noise level. Consequently, the error-correction capability of the corresponding outer code would need to be enormous, limiting the number of information bits that can be allocated to this level. Due to this limitation, and the computational cost of actually decoding $\mathrm{RM}(2, m)$, we restrict ourselves to the 2-level construction shown in Table 1.

4 Errors-and-Erasure Decoder for HQC

This section analyzes errors-and-erasures decoding of the outer code. The reliability measure that we derive from the inner decoding is inspired by [19], and we conservatively bound its distribution in Section 4.1. Section 4.2 considers an erasure assignment scheme in which all symbols with the reliabilities below a predetermined threshold are erased. In Section 4.3, we consider the approach of erasing the ϵ least reliable symbols. All results are stated for a single concatenation level and assuming a generic inner code $\mathcal{B}$; the extension to multiple levels via Equation (2) is immediate. Without loss of generality, we assume $\mathbf{c}_i = \mathbf{0}$ (so $\mathbf{r}_i = \mathbf{e}_i$) throughout the section.

4.1 Analysis of Proposed Reliability Measure

In addition to the estimated symbol, the inner decoding inherently provides reliability information. For example, when $\hat{\mathbf{c}}_i = \mathbf{r}_i$, the inner decoding is correct with high probability. On the other hand, if there is another codeword $\hat{\mathbf{c}}'_i$ with $\mathrm{d_H}(\hat{\mathbf{c}}_i, \mathbf{r}_i) = \mathrm{d_H}(\hat{\mathbf{c}}'_i, \mathbf{r}_i)$, the inner decoding is correct with probability at most $\frac{1}{2}$. More generally,

the posterior probability of the decoded codeword given the received word equals

$$\Pr[\hat{\mathbf{c}}_i \,|\, \mathbf{r}_i] = \frac{\Pr[\mathbf{r}_i \,|\, \hat{\mathbf{c}}_i]}{\sum_{\mathbf{c}\in\mathcal{C}} \Pr[\mathbf{r}_i \,|\, \mathbf{c}]} = \frac{(\frac{p}{1-p})^{\mathrm{d_H}(\hat{\mathbf{c}}_i,\mathbf{r}_i)}}{\sum_{\mathbf{c}\in\mathcal{C}} (\frac{p}{1-p})^{\mathrm{d_H}(\mathbf{c},\mathbf{r}_i)}}$$

and constitutes the optimal reliability measure.

Example 4 For a binary repetition code of length n_{in}, we have

$$\Pr[\hat{\mathbf{c}}_i \,|\, \mathbf{r}_i] = \frac{(\frac{p}{1-p})^{\mathrm{d_H}(\hat{\mathbf{c}}_i,\mathbf{r}_i)}}{\sum_{\mathbf{c}\in\mathcal{C}} (\frac{p}{1-p})^{\mathrm{d_H}(\mathbf{c},\mathbf{r}_i)}} = \left(1 + (\frac{p}{1-p})^{d_i'-d_i}\right)^{-1},$$

where $d_i = \mathrm{d_H}(\hat{\mathbf{c}}_i, \mathbf{r}_i)$ denotes the distance to the estimated codeword $\hat{\mathbf{c}}_i$ and $d_i' = \mathrm{d_H}(\hat{\mathbf{c}}_i', \mathbf{r}_i)$ the distance to the second-best one. As $(\frac{p}{1-p}) < 1$, $\Pr[\hat{\mathbf{c}}_i \,|\, \mathbf{r}_i]$ is strictly increasing in $\alpha_i = d_i' - d_i$, which is therefore a suitable choice of reliability measure, i.e., it captures all information contained in $\Pr[\hat{\mathbf{c}}_i \,|\, \mathbf{r}_i]$. Moreover, $d_i' = n_{\mathrm{in}} - d_i$.

Example 4 illustrates that precisely characterizing the distribution of the reliability information (also conditioning on erroneous or successful decoding) is straightforward for binary repetition codes, since there are only two codewords. For general inner codes, however, such a characterization requires accounting for the complete code geometry, which quickly becomes infeasible as $|\mathcal{B}|$ grows. To overcome this issue, observe that typically the best estimate $\hat{\mathbf{c}}_i$ and the second-best $\hat{\mathbf{c}}_i'$ are closer to $\mathbf{r}_i$ than any other codeword by a significant margin. Therefore, we may approximate $\sum_{\mathbf{c}\in\mathcal{B}} \Pr[\mathbf{r}_i \,|\, \mathbf{c}] \approx \Pr[\mathbf{r}_i \,|\, \hat{\mathbf{c}}_i] + \Pr[\mathbf{r}_i \,|\, \hat{\mathbf{c}}_i']$, which yields

$$\Pr[\hat{\mathbf{c}}_i \,|\, \mathbf{r}_i] \approx \frac{1}{1 + (\frac{p}{1-p})^{\alpha_i}} \text{ with } \alpha_i \triangleq d_i' - d_i = \mathrm{d_H}(\hat{\mathbf{c}}_i', \mathbf{r}_i) - \mathrm{d_H}(\hat{\mathbf{c}}_i, \mathbf{r}_i),$$

which is monotonically increasing in α_i, justifying our choice of α_i as the reliability measure.

Remark 1 This approximation is exact for repetition codes and tight whenever the two nearest codewords dominate the likelihood sum; the formal analysis below via the surrogate $\tilde{\alpha}_i$ and union bounds does not rely on this approximation for correctness.

Analysis of distribution. Looking ahead, the DFR analyses in Sections 4.2 and 4.3 require the distribution of α_i jointly with the decoding outcome, i.e., the joint probabilities $\Pr\big[\alpha_i \geq x, \hat{\mathbf{c}}_i = \mathbf{c}_i\big]$ and $\Pr\big[\alpha_i \geq x, \hat{\mathbf{c}}_i \neq \mathbf{c}_i\big]$. Since the joint distribution of α_i and the decoding outcome depends on the full distance profile of $\mathcal{B}$, we define an alternative quantity $\tilde{\alpha}_i$ that admits a tractable bound.

Lemma 1 *Define* $\hat{\mathbf{c}}_i^\star \triangleq \arg\min_{\mathbf{v}\in\mathcal{B}\setminus\{\mathbf{c}_i\}} \mathrm{d_H}(\mathbf{v},\mathbf{r}_i)$ *as the closest codeword in* $\mathcal{B}\setminus\{\mathbf{c}_i\}$, $w_i = |\mathbf{e}_i|_\mathrm{H}$, *and* $w_i^\star = \mathrm{d_H}(\hat{\mathbf{c}}_i^\star,\mathbf{r}_i)$. *Then*

$$\tilde{\alpha}_i \triangleq \left|w_i^\star - w_i\right| \begin{cases} = \alpha_i & \text{if } \hat{\mathbf{c}}_i = \mathbf{c}_i \\ \geq \alpha_i & \text{otherwise.} \end{cases}$$

Proof For $\hat{\mathbf{c}}_i = \mathbf{c}_i$, we have $w_i = \mathrm{d_H}(\mathbf{r}_i,\hat{\mathbf{c}}_i) = d_i$ and $w_i^\star = \mathrm{d_H}(\mathbf{r}_i,\hat{\mathbf{c}}_i') = d_i'$. For $\hat{\mathbf{c}}_i \neq \mathbf{c}_i$, $\hat{\mathbf{c}}_i^\star = \hat{\mathbf{c}}_i$ (since $\hat{\mathbf{c}}_i \in \mathcal{B}\setminus\{\mathbf{c}_i\}$ minimizes $\mathrm{d_H}(\mathbf{v},\mathbf{r}_i)$ over all $\mathbf{v}\in\mathcal{B}$, hence also over $\mathcal{B}\setminus\{\mathbf{c}_i\}$), so $w_i^\star = \mathrm{d_H}(\mathbf{r}_i,\hat{\mathbf{c}}_i^\star) = \mathrm{d_H}(\mathbf{r}_i,\hat{\mathbf{c}}_i) = d_i$ and $w_i = \mathrm{d_H}(\mathbf{r}_i,\mathbf{c}_i) \geq \mathrm{d_H}(\mathbf{r}_i,\hat{\mathbf{c}}_i') = d_i'$. □

The key property of $\tilde{\alpha}_i$ is that it coincides with α_i upon correct decoding and may only *overestimate* the reliability upon decoding failure. As a consequence, any upper bound on $\Pr\big[\tilde{\alpha}_i \geq x, \hat{\mathbf{c}}_i \neq \mathbf{c}_i\big]$ is also a valid upper bound on $\Pr\big[\alpha_i \geq x, \hat{\mathbf{c}}_i \neq \mathbf{c}_i\big]$, and any lower bound on $\Pr\big[\tilde{\alpha}_i \geq x, \hat{\mathbf{c}}_i = \mathbf{c}_i\big]$ is a valid lower bound on $\Pr\big[\alpha_i \geq x, \hat{\mathbf{c}}_i = \mathbf{c}_i\big]$. Both directions are conservative for DFR estimation.

Theorem 2 *Let* $\bar{P}_\mathrm{in}$ *be any upper bound on the inner failure probability, i.e.,* $\Pr\big[\hat{\mathbf{c}}_i \neq \mathbf{c}_i\big] \leq \bar{P}_\mathrm{in}$. *Then, for all* $i \in [n_\mathrm{out}]$ *and all* x, *with* $\tilde{\alpha}_i$ *as in Lemma 1,*

$$\Pr\big[\alpha_i \geq x, \hat{\mathbf{c}}_i \neq \mathbf{c}_i\big] \leq \Pr\big[\tilde{\alpha}_i \geq x, \hat{\mathbf{c}}_i \neq \mathbf{c}_i\big] \leq \Pr\big[\alpha_i' \geq x, E_i'\big],$$
$$\Pr\big[\alpha_i \geq x, \hat{\mathbf{c}}_i = \mathbf{c}_i\big] = \Pr\big[\tilde{\alpha}_i \geq x, \hat{\mathbf{c}}_i = \mathbf{c}_i\big] \geq \Pr\big[\alpha_i' \geq x, C_i'\big],$$

where the surrogate quantities are defined, for $x > 0$, *by*

$$\Pr\big[\alpha_i' \geq x, E_i'\big] \triangleq \sum_{w_i=0}^{n_\mathrm{in}} \Pr\big[|\mathbf{e}_i|_\mathrm{H} = w_i\big]\, \tilde{\mathrm{P}}(\leq w_i - x \,|\, w_i),$$
$$\Pr\big[\alpha_i' \geq x, C_i'\big] \triangleq \sum_{w_i=0}^{n_\mathrm{in}} \Pr\big[|\mathbf{e}_i|_\mathrm{H} = w_i\big]\, \Big(1 - \tilde{\mathrm{P}}(< w_i + x \,|\, w_i)\Big),$$

and, for $x = 0$, *by* $\Pr\big[\alpha_i' \geq x, E_i'\big] \triangleq \bar{P}_\mathrm{in}$ *and* $\Pr\big[\alpha_i' \geq x, C_i'\big] \triangleq 1 - \bar{P}_\mathrm{in}$. *Here*

$$\tilde{\mathrm{P}}\big(\leq w^\star \,\big|\, w_i\big) \triangleq \min\left\{ \sum_{\mathbf{v}\in\mathcal{B}\setminus\{\mathbf{0}\}} \Pr\big[\mathrm{d_H}(\mathbf{e}_i,\mathbf{v}) \leq w^\star \,\big|\, |\mathbf{e}_i|_\mathrm{H} = w_i\big],\ 1 \right\}$$

is a union bound on the probability that some nonzero codeword lies within distance $w^\star$ *of the error pattern.*

Proof As $\tilde{\alpha}_i$ is by definition non-negative,

$$\Pr\big[\tilde{\alpha}_i \geq 0, \hat{\mathbf{c}}_i \neq \mathbf{c}_i\big] = \Pr\big[\hat{\mathbf{c}}_i \neq \mathbf{c}_i\big] \leq \Pr\big[\alpha_i' \geq 0, E_i'\big] = \bar{P}_\mathrm{in}$$
$$\Pr\big[\tilde{\alpha}_i \geq 0, \hat{\mathbf{c}}_i = \mathbf{c}_i\big] = \Pr\big[\hat{\mathbf{c}}_i = \mathbf{c}_i\big] \geq \Pr\big[\alpha_i' \geq 0, C_i'\big] = 1 - \bar{P}_\mathrm{in}.$$

For $\tilde{\alpha}_i \geq x > 0$, we show the case of $\hat{\mathbf{c}}_i \neq \mathbf{c}_i$, the case $\hat{\mathbf{c}}_i = \mathbf{c}_i$ is analogous. Denote the random variables $W_i = |\mathbf{e}_i|_\mathrm{H}$ and $W_i^\star = \mathrm{d_H}(\hat{\mathbf{c}}_i^\star,\mathbf{e}_i)$. Since $\tilde{\alpha}_i = |W_i^\star - W_i|$, the event $\{\tilde{\alpha}_i \geq x,\ \hat{\mathbf{c}}_i \neq \mathbf{c}_i\}$ decomposes as

$$\begin{aligned}\Pr\big[\tilde{\alpha}_i \geq x, \hat{\mathbf{c}}_i \neq \mathbf{c}_i\big] &= \Pr\big[(W_i^\star - W_i \geq x \cup W_i - W_i^\star \geq x), \hat{\mathbf{c}}_i \neq \mathbf{c}_i\big] \\ &= \Pr\big[W_i - W_i^\star \geq x, \hat{\mathbf{c}}_i \neq \mathbf{c}_i\big],\end{aligned}$$

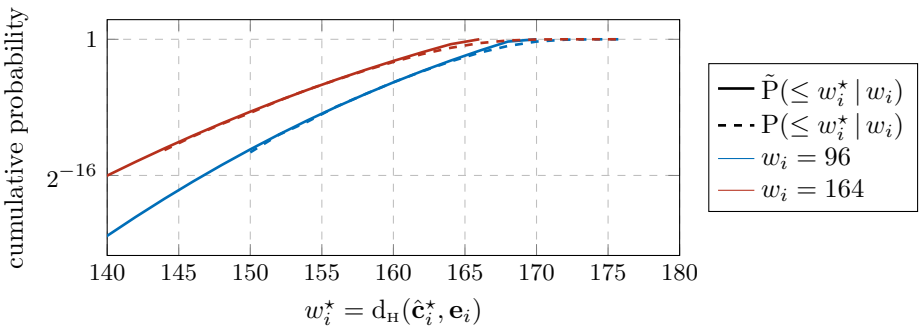

Fig. 1: Cumulative distribution of $w_i^\star$ for inner $\mathcal{B} = \mathrm{RM}$ with $n_{\mathrm{in}} = 384$; $p = 0.34$.

where the second equality holds because $W_i^\star - W_i \geq x > 0$ would imply $\hat{\mathbf{c}}_i = \mathbf{c}_i$ (the transmitted codeword is closer than all others). On the other hand, $W_i - W_i^\star \geq x > 0$ already implies $\hat{\mathbf{c}}_i \neq \mathbf{c}_i$, so the intersection is redundant and it follows that

$$\begin{aligned}\Pr\big[W_i - W_i^\star \geq x, \hat{\mathbf{c}}_i \neq \mathbf{c}_i\big] &= \Pr\big[W_i^\star \leq W_i - x\big] \\ &= \sum_{w_i} \Pr\big[W_i^\star \leq w_i - x \,\big|\, W_i = w_i\big] \Pr\big[W_i = w_i\big]\end{aligned}$$

The result follows using the union bound:

$$\Pr\big[W_i^\star \leq w_i - x \,\big|\, W_i = w_i\big] \leq \sum_{\mathbf{v} \in \mathcal{B} \setminus \{\mathbf{0}\}} \Pr[\mathrm{d_H}(\mathbf{v}, \mathbf{e}_i) \leq w_i - x \,|\, W_i = w_i].$$

□

The quantity $\tilde{\mathrm{P}}(\leq w^\star \,|\, w)$ is a simple union bound on the probability that at least one nonzero codeword lies within distance $w^\star$ of the error pattern and can be evaluated in closed form for arbitrary inner codes. Therefore, a valid choice for $\bar{P}_{\mathrm{in}}$ is

$$\bar{P}_{\mathrm{in}} = \sum_{w_i} \Pr\big[|\mathbf{e}_i|_{\mathrm{H}} = w_i\big] \, \tilde{\mathrm{P}}(\leq w_i - 1 \,|\, w_i);$$

A tighter bound for (duplicated) first-order RM codes is given in [1].

Simulation results. Figure 1 compares the bound $\tilde{\mathrm{P}}(\leq w_i^\star \,|\, w_i)$ from Theorem 2 with the empirically estimated CDF $\mathrm{P}(\leq w_i^\star \,|\, w_i) \triangleq \Pr[\mathrm{d_H}(\hat{\mathbf{c}}_i^\star, \mathbf{e}_i) \leq w_i^\star \,|\, w_i]$ for fixed error weights $w_i \in \{96, 164\}$. For both values (as well as other tested error weights), the union bound $\tilde{\mathrm{P}}(\leq w_i^\star \,|\, w_i)$ provides a tight upper bound on the empirical distribution, shifting mass toward smaller values of $w_i^\star$ as expected.

Figure 2 compares the bounds on $\Pr\big[\tilde{\alpha}_i \geq x, \hat{\mathbf{c}}_i = \mathbf{c}_i\big]$ and $\Pr\big[\tilde{\alpha}_i \geq x, \hat{\mathbf{c}}_i \neq \mathbf{c}_i\big]$ from Theorem 2 with the empirical distributions of $\Pr\big[\alpha_i \geq x, \hat{\mathbf{c}}_i = \mathbf{c}_i\big]$ and $\Pr\big[\alpha_i \geq x, \hat{\mathbf{c}}_i \neq \mathbf{c}_i\big]$. For the erroneous-decoding event, the analytical curve is a conservative upper bound on the empirical one for all x, confirming the validity of the bound from Theorem 2. For the correct-decoding event, the analytical and empirical curves are visually indistinguishable on the logarithmic scale; their difference is below 10^{-3} throughout. For α_i such that the Monte Carlo estimates are accurate to within 10^{-3}, the empirical values are consistently above the analytical lower bound, as expected.

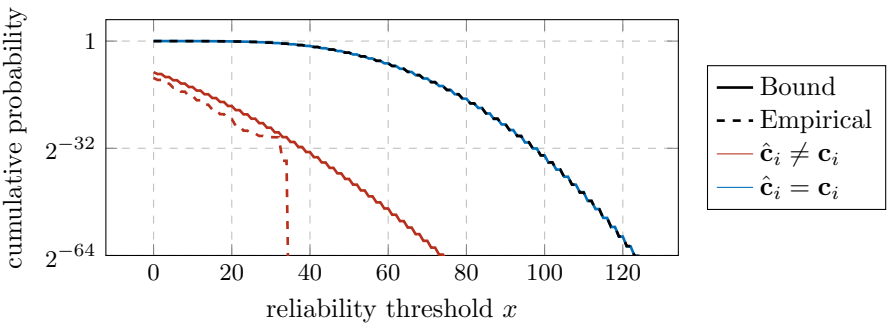

Fig. 2: Distribution of α_i for the inner code $\mathcal{B} = \mathrm{RM}$ with $n_{\mathrm{in}} = 384$ and $p = 0.34$.

4.2 Threshold-based Errors-and-Erasures Decoding

This section analyzes the classical decoding strategy in which outer symbols whose inner decoding is deemed unreliable are *erased* before outer decoding [20]. This reduces the DFR by converting potential errors (which cost the outer code two units of minimum distance each) into erasures (which cost only one).

Erasure assignment scheme. Fix a threshold $T \in \mathbb{Z}$. After inner decoding, each outer symbol $i \in [n_{\mathrm{out}}]$ is classified as follows:

- If $\alpha_i \geq T$, the estimate $\hat{\mathbf{c}}_i$ is *accepted* and forwarded to the outer decoder.
- If $\alpha_i < T$, the position i is declared an *erasure*.

The outer decoder then operates in errors-and-erasures mode.

Error and erasure probabilities. Assuming $\mathbf{e} \sim \mathrm{Bern}(p)^n$, the inner blocks are independent and identically distributed. The decoding outcome at each position i is determined by three mutually exclusive events:

1. *Correct acceptance*, with probability $p_{\mathrm{c}}(T) \triangleq \Pr\big[\alpha_i \geq T,\, \hat{\mathbf{c}}_i = \mathbf{c}_i\big]$.
2. *Undetected error*, with probability $p_{\mathrm{e}}(T) \triangleq \Pr\big[\alpha_i \geq T,\, \hat{\mathbf{c}}_i \neq \mathbf{c}_i\big]$.
3. *Erasure*, with probability $p_{\mathrm{s}}(T) \triangleq \Pr\big[\alpha_i < T\big] = 1 - p_{\mathrm{c}}(T) - p_{\mathrm{e}}(T)$.

Using Theorem 2, we bound the probabilities of Case 1 and 2 as

$$p_{\mathrm{e}}(T) = \Pr\big[\alpha_i \geq T,\, \hat{\mathbf{c}}_i \neq \mathbf{c}_i\big] \leq \Pr\big[\alpha'_i \geq T,\, E'_i\big] \triangleq p'_{\mathrm{e}}(T), \tag{3}$$

$$p_{\mathrm{c}}(T) = \Pr\big[\alpha_i \geq T,\, \hat{\mathbf{c}}_i = \mathbf{c}_i\big] \geq \Pr\big[\alpha'_i \geq T,\, C'_i\big] \triangleq p'_{\mathrm{c}}(T). \tag{4}$$

To cover Case 3, we further define

$$p'_{\mathrm{s}}(T) \triangleq 1 - p'_{\mathrm{e}}(T) - p'_{\mathrm{c}}(T). \tag{5}$$

In general, $p'_{\mathrm{s}}(T)$ is not guaranteed to be a lower or upper bound on $p_{\mathrm{s}}(T)$. However, for the DFR bound, $p'_{\mathrm{c}}(T) \leq p_{\mathrm{c}}(T)$ and $p'_{\mathrm{e}}(T) \geq p_{\mathrm{e}}(T)$ is sufficient.

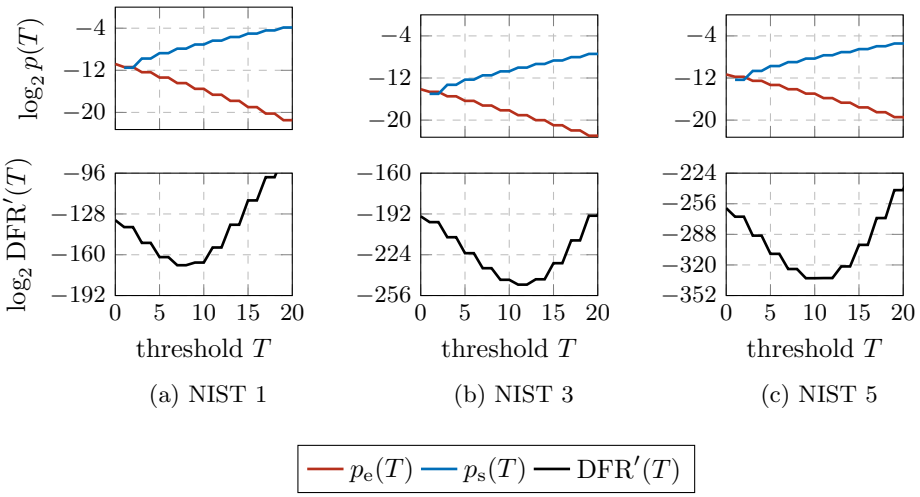

(a) NIST 1 (b) NIST 3 (c) NIST 5

Fig. 3: Probability of inner error, inner erasure, and decryption failure as trade-off via threshold T. Concatenation of RM and RS code as in [1], see Table 2.

Proposition 3 *For any threshold $T \geq 0$, the DFR is bounded as*

$$\mathrm{DFR}(T) \leq \mathrm{DFR}'(T) \triangleq \sum_{\substack{e,s \geq 0 \\ 2e+s \geq d_\mathrm{out}}} \binom{n_\mathrm{out}}{e,s} p'_\mathrm{e}(T)^e \, p'_\mathrm{s}(T)^s \, p'_\mathrm{c}(T)^{n_\mathrm{out}-e-s},$$

where the multinomial coefficient is $\binom{n_\mathrm{out}}{e,s} = \frac{n_\mathrm{out}!}{e!\,s!\,(n_\mathrm{out}-e-s)!}$ for $e+s \leq n_\mathrm{out}$.

Proof Denote the number of undetected errors (Case 2) as e and the number of erasures (Case 3) as s, i.e., there are $n_\mathrm{out}-e-s$ correct acceptances (Case 1). The outer decoding fails if and only if $2e+s \geq d_\mathrm{out}$. Since each position corresponds to an independent multinomial trial with respect to the three possible outcomes, the DFR is computed as

$$\mathrm{DFR}(T) = \sum_{\substack{e,s \geq 0 \\ 2e+s \geq d_\mathrm{out}}} \binom{n_\mathrm{out}}{e,s} p_\mathrm{e}(T)^e \, (1 - p_\mathrm{c}(T) - p_\mathrm{e}(T))^s \, p_\mathrm{c}(T)^{n_\mathrm{out}-e-s},$$

Replacing $p_\mathrm{e}(T)$ by $p'_\mathrm{e}(T)$ reduces the erasure probability in favor of the error probability, and replacing $p_\mathrm{c}(T)$ by $p'_\mathrm{c}(T)$ reduces the correct acceptance probability in favor of the erasure probability. Errors cost more error-correction budget than erasures, and erasures cost more than correct decoding. Therefore, these substitutions can only increase the probability of incorrect outer decoding. □

Choice of threshold. As illustrated in Figure 3 for code parameters suggested in [1], the threshold T controls the trade-off between errors and erasures. For $T = 0$, no erasures are declared, and the hard-decision decoder employed in HQC is recovered. As T increases, inner decoding failures with low reliability are increasingly caught as

erasures, at the cost of increasing the erasure probability. We pick the threshold $T^\star$ that minimizes the DFR bound, i.e.,

$$T^\star = \arg\min_{T \geq 0} \ \mathrm{DFR}'(T).$$

For NIST security category 1, $T^\star = 8$ guarantees a DFR of at most 2^{-168}. Typically, a slightly higher or lower threshold offers a similar DFR performance.

Practical considerations. The threshold-based erasure assignment scheme proposed in this section (and also used in HARE [10]) inherently results in a varying number of erasures. The number of erasures depends on the secret key and may leak information about it. Therefore, a practical implementation has to ensure that the number of erasures is not leaked through side channels. A sequence of side-channel attacks targeting the HQC decoder (see [21] and references therein) testifies that this is non-trivial. The following section proposes an assignment scheme with a fixed number of erasures that achieves comparable performance (cf. Section 5).

4.3 Partition-based Errors-and-Erasures Decoding

An alternative to the threshold-based scheme of Section 4.2 is to erase a *fixed number* of outer symbols—specifically, those with the lowest reliability [20]. By fixing the number of erasures rather than the threshold, the erasure count becomes deterministic, which avoids worst-case fluctuations and simplifies the interaction with the outer code's error-correction budget.

Erasure assignment. Fix a parameter $\epsilon \in \{0, 1, \ldots, n_{\mathrm{out}}\}$. After inner decoding, partition $[n_{\mathrm{out}}] = \mathcal{J} \cup \mathcal{I}$ with $|\mathcal{I}| = \epsilon$, $|\mathcal{J}| = n_{\mathrm{out}} - \epsilon$, such that

$$\alpha_i \leq \alpha_j \quad \forall\, i \in \mathcal{I},\ j \in \mathcal{J},$$

with ties broken arbitrarily (e.g., by block index). The positions indexed by $\mathcal{I}$ are declared erasures, and the remaining $n_{\mathrm{out}} - \epsilon$ symbols are forwarded to the outer errors-and-erasures decoder, which corrects up to $t_\epsilon \triangleq \lfloor (d_{\mathrm{out}} - 1 - \epsilon)/2 \rfloor$ additional errors.

DFR analysis. Agrawal and Vardy [19] (see also [22]) analyze the failure probability of the described decoder. However, in their case, the exact joint probability law of decoding outcome and reliability measure is available. In what follows, we show that their bound still holds when we work with the surrogates of Theorem 2.

We describe each position by its state s_i, where $s_i = (\hat{\mathbf{c}}_i = \mathbf{c}_i, \alpha_i)$ indicates correct inner decoding of block i with reliability α_i, while $s_i = (\hat{\mathbf{c}}_i \neq \mathbf{c}_i, \alpha_i)$ corresponds to erroneous inner decoding. We order states from *worst* to *best* as

$$(\hat{\mathbf{c}}_i \neq \mathbf{c}_i, n_{\mathrm{in}}) \succ \cdots \succ (\hat{\mathbf{c}}_i \neq \mathbf{c}_i, 0) \succ (\hat{\mathbf{c}}_i = \mathbf{c}_i, 0) \succ \cdots \succ (\hat{\mathbf{c}}_i = \mathbf{c}_i, n_{\mathrm{in}}). \tag{6}$$

That is, with respect to the outcome of the errors-and-erasures decoder, any error is worse than a correct symbol; among errors, higher reliability is worse, while, among correct symbols, lower reliability is worse. This is formally shown in the following lemma.

Lemma 4 *Let $f_\epsilon(s_1, \ldots, s_{n_{\text{out}}}) \in \{0, 1\}$ indicate whether the decoder fails for input states $s_1, \ldots, s_{n_{\text{out}}}$. Then f_ϵ is non-decreasing in the product order induced by (6): moving positions to a worse state can turn a success into a failure, the reverse is impossible.*

Proof The order is total and generated by the covering moves $(\hat{\mathbf{c}}_i \neq \mathbf{c}_i, \alpha) \succ (\hat{\mathbf{c}}_i = \mathbf{c}_i, \alpha)$, $(\hat{\mathbf{c}}_i \neq \mathbf{c}_i, \alpha + 1) \succ (\hat{\mathbf{c}}_i \neq \mathbf{c}_i, \alpha)$, and $(\hat{\mathbf{c}}_i = \mathbf{c}_i, \alpha - 1) \succ (\hat{\mathbf{c}}_i = \mathbf{c}_i, \alpha)$. It suffices to check each covering move on a single coordinate i, holding the others fixed. The erasure set depends only on the reliabilities, and every move keeps exactly ϵ positions erased, so we track the one position i that may cross the cutoff.

$(\hat{\mathbf{c}}_i = \mathbf{c}_i, \alpha) \to (\hat{\mathbf{c}}_i \neq \mathbf{c}_i, \alpha)$*:* The reliability is unchanged, so the erasure set is unchanged. If i is not erased, the number of kept errors rises by one; otherwise, nothing changes.

$(\hat{\mathbf{c}}_i \neq \mathbf{c}_i, \alpha) \to (\hat{\mathbf{c}}_i \neq \mathbf{c}_i, \alpha + 1)$*:* If the erasure set is unchanged, nothing changes. If position i is moved out of the erasure set, the number of kept errors may increase or remain the same.

$(\hat{\mathbf{c}}_i = \mathbf{c}_i, \alpha) \to (\hat{\mathbf{c}}_i = \mathbf{c}_i, \alpha - 1)$*:* Symmetrically, if i moves from kept to erased, some previously erased position is retained instead, which may increase the error count.

For any move, the number of kept errors is either constant or increases. Hence, f_ϵ remains either constant or increases. A general single-coordinate increase is a chain of moves, and the product order is handled one coordinate at a time. □

Having established an order on the possible states with respect to their impact on decoding failure events, we can formally prove that the surrogate distribution of Theorem 2 favors bad states compared to the actual state distribution.

Lemma 5 *Denote by S the actual state distribution that the decoder sees, and by S' the surrogate from Theorem 2, characterized by $\Pr[\alpha'_i \geq x, E'_i]$ and $\Pr[\alpha'_i \geq x, C'_i]$. Then, S' favors worse states compared to S: Let $s \sim S$ and $s' \sim S'$. Then, for any state σ,*

$$\Pr[s \succeq \sigma] \leq \Pr[s' \succeq \sigma].$$

Proof We distinguish whether σ corresponds to an inner decoding failure or not.

Error event $\sigma = (\hat{\mathbf{c}}_i \neq \mathbf{c}_i, x)$: The set of states at least as bad as σ is exactly the errors of reliability at least x, with probability $\Pr[\alpha_i \geq x, \hat{\mathbf{c}}_i \neq \mathbf{c}_i]$ under S and $\Pr[\alpha'_i \geq x, E'_i]$ under S'. By the first inequality of Theorem 2, it holds that

$$\Pr[s' \succeq \sigma] = \Pr[\alpha'_i \geq x, E'_i] \geq \Pr[\alpha_i \geq x, \hat{\mathbf{c}}_i \neq \mathbf{c}_i] = \Pr[s \succeq \sigma] .$$

Correct event $\sigma = (\hat{\mathbf{c}}_i = \mathbf{c}_i, x)$: The set of states worse than σ contains all errors together with the correct symbols of reliability at most x; equivalently, everything except the correct symbols with reliability at least $x + 1$. Thus, the second inequality of Theorem 2 implies

$$\Pr[s' \succeq \sigma] = 1 - \Pr[\alpha'_i \geq x + 1, C'_i] \geq 1 - \Pr[\alpha_i \geq x + 1, \hat{\mathbf{c}}_i = \mathbf{c}_i] = \Pr[s \succeq \sigma].$$ □

Together, Lemmas 4 and 5 lead to a DFR bound based on Theorem 2.

Proposition 6 *For $i \in [n_{\text{out}}]$, let $s_i \sim S$ and $s'_i \sim S'$. Then*

$$\text{DFR}(\epsilon) = \Pr[f_\epsilon(s_1, \ldots, s_{n_{\text{out}}}) = 1] \leq \Pr[f_\epsilon(s'_1, \ldots, s'_{n_{\text{out}}}) = 1] = \text{DFR}'(\epsilon),$$

i.e., the actual DFR is bounded from above by the surrogate of Theorem 2.

Proof The proof proceeds by a sequence of hybrids. For the hybrid experiment ℓ, the first ℓ states are sampled from S' and the last $n_{\text{out}} - \ell$ states from S. The corresponding DFR is

$$F_\ell \triangleq \Pr\left[f_\epsilon(s'_1, \ldots, s'_\ell, s_{\ell+1}, \ldots, s_{n_{\text{out}}}) = 1\right] = \mathbb{E}[f_\epsilon(s'_1, \ldots, s'_\ell, s_{\ell+1}, \ldots, s_{n_{\text{out}}})].$$

Experiment $\ell - 1$ and ℓ differ only in position ℓ, which switches from distribution S to distribution S'; all other positions are identical and independent. $F_{\ell-1} \leq F_\ell$ follows from a standard stochastic dominance argument: We write $\tilde{\mathbf{s}} = (s'_1, \ldots, s'_{\ell-1}, s_{\ell+1}, \ldots, s_{n_{\text{out}}})$ and

$$g_{\tilde{\mathbf{s}}}(s) = f_\epsilon(s'_1, \ldots, s'_{\ell-1}, s, s_{\ell+1}, \ldots, s_{n_{\text{out}}}).$$

Then, $F_{\ell-1} = \sum_{\tilde{\mathbf{s}}} \Pr[\tilde{\mathbf{s}}]\, \mathbb{E}_{s\sim S}[g_{\tilde{\mathbf{s}}}(s)]$ and $F_\ell = \sum_{\tilde{\mathbf{s}}} \Pr[\tilde{\mathbf{s}}]\, \mathbb{E}_{s'\sim S'}[g_{\tilde{\mathbf{s}}}(s')]$, where $\Pr[\tilde{\mathbf{s}}]$ is identical in both experiments. By Lemma 4, moving to a worse state in position ℓ can only turn a success into a failure, so $g_{\tilde{\mathbf{s}}}(s)$ is non-decreasing in the order (6). By Lemma 5, S' puts more mass on bad states compared to S, i.e., S' stochastically dominates S in that order. Together, Lemmas 4 and 5 imply $\mathbb{E}_{s'\sim S'}[g_{\tilde{\mathbf{s}}}(s')] \geq \mathbb{E}_{s\sim S}[g_{\tilde{\mathbf{s}}}(s)]$ for any $\tilde{\mathbf{s}}$; since $g_{\tilde{\mathbf{s}}}(s)$ is non-decreasing, a stochastically larger law yields no smaller expectation. It follows that $F_{\ell-1} \leq F_\ell$ for any $\ell \in [n_{\text{out}}]$, and we conclude

$$\text{DFR}(\epsilon) = F_0 \leq F_1 \leq \ldots \leq F_{n_{\text{out}}} = \text{DFR}'(\epsilon). \qquad \square$$

Computing the DFR bound. Proposition 6 establishes that the DFR can be computed by replacing the true distribution with the bound of Theorem 2 and following the blueprint of Agrawal and Vardy [19]. We decompose the DFR bound as

$$\text{DFR}'(\epsilon) = \sum_{w=0}^{n_{\text{out}}} \binom{n_{\text{out}}}{w} (P'_{\text{in}})^w (1 - P'_{\text{in}})^{n_{\text{out}}-w} \cdot \Pr[f_\epsilon = 1 \,|\, w]. \tag{7}$$

Here, $P'_{\text{in}} \triangleq \Pr\left[\alpha'_i \geq 0, E'_i\right]$ denotes the inner failure probability according to Theorem 2 and $\Pr[f_\epsilon = 1 \,|\, w]$ denotes the outer failure probability conditioned on w erroneous inner decodings. The following lemma, a direct adaptation from [19], provides the bound on $\Pr[f_\epsilon = 1 \,|\, w]$ that we use.

Lemma 7 *Let $\beta_1(w) \geq \cdots \geq \beta_w(w)$ denote the ordered reliabilities of the w erroneous positions and $\gamma_1(n_{\text{out}} - w) \geq \cdots \geq \gamma_{n_{\text{out}}-w}(n_{\text{out}} - w)$ those of the $n_{\text{out}} - w$ correct positions. Then, the outer failure probability is upper-bounded by*

$$\Pr[f_\epsilon = 1 \,|\, w] \leq \begin{cases} 0, & \text{if } w \leq t_\epsilon, \\ \Pr\left[\beta_{t_\epsilon+1}(w) \geq \gamma_{n_{\text{out}}-\epsilon-t_\epsilon}(n_{\text{out}} - w)\right], & \text{if } t_\epsilon < w \leq \epsilon + t_\epsilon, \\ 1, & \text{if } w > \epsilon + t_\epsilon. \end{cases} \tag{8}$$

Choice of erasure number. Figure 4 illustrates how the parameter ϵ controls the trade-off between errors and erasures. For $\epsilon = 0$, no erasures are used, and the classical hard-decision decoder of HQC is recovered. As ϵ increases, more potentially erroneous symbols are erased, which reduces $\Pr[f_\epsilon = 1 \,|\, w]$ for moderate ϵ. For larger ϵ the remaining error-correction radius $t_\epsilon = \lfloor (d_{\text{out}} - 1 - \epsilon)/2 \rfloor$ is too low. If $\epsilon \equiv d_{\text{out}} \pmod 2$, $\epsilon \pm 1$ can correct strictly more error patterns than ϵ; this leads to the ripple of the curve in Figure 4. Typically, several choices for ϵ offer a similar DFR performance. Therefore, among all ϵ values that achieve a sufficiently low DFR, the value that is

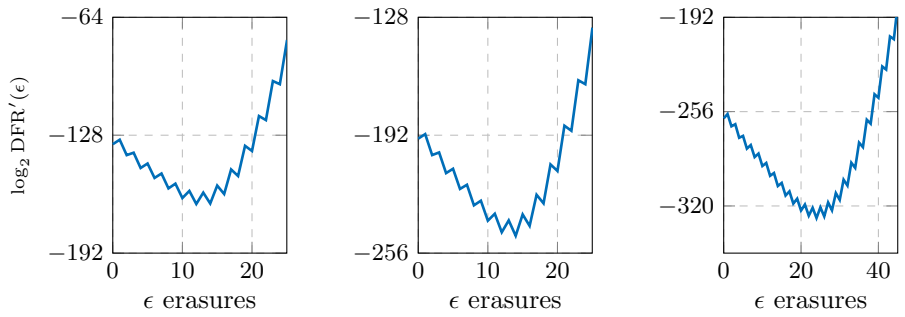

(a) NIST 1 (b) NIST 3 (c) NIST 5

Fig. 4: DFR bound as a function of the number of erased positions ϵ. Concatenation of RM and RS code as in [1], see Table 2.

most convenient implementation-wise may be chosen. When proposing parameters, we pick $\epsilon^\star$, that minimizes the DFR bound, i.e.,

$$\epsilon^\star = \arg\min_{\epsilon} \mathrm{DFR}'(\epsilon).$$

4.4 Extending to Any Error Weight Distribution

Model 1: Bernoulli error. In Section 4.2 and Section 4.3, the DFR of the decoders was bounded under the assumption that the HQC error $\mathbf{e}$ is a vector of i.i.d. Bernoulli random variables. This assumption, to which we are going to refer as Model 1, is justified by the observation that it overestimates the DFR compared to the true behavior of the HQC error [1].

Model 2: Alternative weight distribution. Recently, [14–16] proposed alternative models for the HQC error that account for dependencies in $\mathbf{e}$. These models offer probability mass functions $\Pr\left[|\mathbf{e}|_{\mathrm{H}} = w\right]$ that differ from Model 1, while still assuming that all errors with $|\mathbf{e}|_{\mathrm{H}} = w$ are equally likely. Here, we assume $\Pr\left[|\mathbf{e}|_{\mathrm{H}} = w\right]$ as given by the Model 2 we introduced in [14] (see Definition 1), but the models of [15, 16] are expected to yield only a minor difference.

DFR under Model 2. To bound the DFR under Model 2, we rewrite Equation (2) as

$$\mathrm{DFR} \leq \sum_{w=0}^{n_{\mathrm{out}} n_{\mathrm{in}}} \Pr\left[|\mathbf{e}|_{\mathrm{H}} = w\right] \sum_{j \in [s]} \Pr\left[E^{(j)} \,\middle|\, |\mathbf{e}|_{\mathrm{H}} = w\right]. \tag{9}$$

That is, we require the failure probability of GCC level j under a fixed error weight.

This computation is very challenging for the partition-based decoder of Section 4.3: The fixed error weight implies that the reliabilities of the inner decoding are dependent. Hence, the analysis would require computing the order statistics of *correlated* random variables, cf. Lemma 7.

Unlike the partition-based decoder, the threshold-based decoder of Section 4.2 assigns erasures based on the individual reliability. In the following, we show that this simpler assignment rule allows analyzing the DFR under Model 2.

Analysis of threshold-based decoding. For ease of notation, we represent by $(w_1, \ldots, w_{n_{\text{out}}})$ the joint event that for each i, block i of $\mathbf{e}$ has weight w_i. Then

$$\Pr\left[E^{(j)} \,\middle|\, |\mathbf{e}|_{\text{H}} = w\right] = \sum_{\substack{w_1,\ldots,w_{n_{\text{out}}} \\ \sum_i w_i = w}} \Pr\left[E^{(j)} \,\middle|\, (w_1, \ldots, w_{n_{\text{out}}})\right] \Pr\left[(w_1, \ldots, w_{n_{\text{out}}})\right]. \quad (10)$$

To compute the outer failure probability given the weight split $(w_1, \ldots, w_{n_{\text{out}}})$, recall that inner decoding assigned each block one of three events. Similar to Section 4.3, the probabilities given error weight w_i can be bounded via Theorem 2:

1. *Correct acceptance*, with probability $p'_{\text{c}}(T \,|\, w_i) \triangleq 1 - \tilde{\text{P}}(< w_i + T \,|\, w_i)$
2. *Undetected error*, with probability $p'_{\text{e}}(T \,|\, w_i) \triangleq \tilde{\text{P}}(\leq w_i - T \,|\, w_i)$.
3. *Erasure*, with probability $p'_{\text{s}}(T \,|\, w_i) = 1 - p'_{\text{c}}(T \,|\, w_i) - p'_{\text{e}}(T \,|\, w_i)$.

Denote the number of undetected errors as e and the number of erasures as s. As the outer decoder fails if and only if $2e + s \geq d_{\text{out}}$, we can write

$$\Pr\left[E^{(j)} \,\middle|\, (w_1, \ldots, w_{n_{\text{out}}})\right] = 1 - \sum_{\delta < d_{\text{out}}} \Pr[2e + s = \delta \,|\, (w_1, \ldots, w_{n_{\text{out}}})]. \quad (11)$$

Dynamic programming allows us to efficiently evaluate $\Pr[2e + s = \delta \,|\, (w_1, \ldots, w_{n_{\text{out}}})]$ via the recursion

$$\begin{aligned}\Pr[2e + s = \delta \,|\, (w_1, \ldots, w_{n_{\text{out}}})] = {} & p'_{\text{c}}(T \,|\, w_{n_{\text{out}}}) \Pr[2e + s = \delta \,|\, (w_1, \ldots, w_{n_{\text{out}}-1})] \\ & + p'_{\text{s}}(T \,|\, w_{n_{\text{out}}}) \Pr[2e + s = \delta - 1 \,|\, (w_1, \ldots, w_{n_{\text{out}}-1})] \\ & + p'_{\text{e}}(T \,|\, w_{n_{\text{out}}}) \Pr[2e + s = \delta - 2 \,|\, (w_1, \ldots, w_{n_{\text{out}}-1})].\end{aligned}$$

Note that this recursion allows us to reuse computations for $(w_1, \ldots, w_{n_{\text{out}}})$ for another weight partition. In fact, this allows computing $\Pr\left[E^{(j)} \,\middle|\, |\mathbf{e}|_{\text{H}} = w\right]$ for *all* weights w via a single pass over a trellis. With each of the n_{out} layers an additional block is incorporated into the computation. For each layer, the nodes correspond to the current total weight and keep track of the distribution of δ. Overall this technique allows bounding the DFR in time $O({n_{\text{out}}}^2 \cdot n_{\text{in}} \cdot d_{\text{out}})$.

As can be seen from the parameters presented in the following section, the results obtained via Model 2 confirm the conservative nature of the DFR bound for the threshold-based decoder under Model 1.

5 Parameters

For various NIST security levels, parameters relying on the proposed GCC (Section 3) and the reliability-based outer decoding (Section 4) are suggested in Table 2. Concretely, we provide results for both the threshold- and partition-based erasure

Table 2: Suggested parameters for NIST 1, 3, and 5. GCCs are constructed according to Table 1. Outer code and decoder parameters are chosen to guarantee a DFR of at most $2^{-\lambda}$ via Equation (2) and the bounds of Section 4.

Security category	n	n_{in}	n_{out}	Outer RS			$-\log_2(\text{DFR})$		Outer $\mathcal{C}$ (BCH)			$-\log_2(\text{DFR})$	
				k_{RS}	d_{RS}	$\mathsf{dec}(\cdot)$	M. 1	M. 2	$k_{\mathcal{C}}$	$d_{\mathcal{C}}$	$\mathsf{dec}(\cdot)$	M. 1	M. 2
NIST 1	17 669	384	46	16	31	HDD	132.9	145.1	none [1]				
	16 901	384	44	15	30	$T^\star=8$	132.9	149.3	25	8	$T^\star=20$	142.4	146.9
	16 901	384	44	15	30	$\epsilon^\star=13$	130.1	-	26	7	$\epsilon^\star=4$	129.4	-
NIST 3	35 851	640	56	24	33	HDD	193.9	205.2	none [1]				
	34 589	640	54	24	31	$T^\star=10$	194.3	207.7	30	9	$T^\star=30$	200.0	205.0
	34 589	640	54	24	31	$\epsilon^\star=12$	194.2	-	30	9	$\epsilon^\star=4$	218.1	-
NIST 5	57 637	640	90	32	59	HDD	260.6	276.5	none [1]				
	55 541	640	86	31	56	$T^\star=10$	259.4	-	43	14	$T^\star=28$	256.6	-
	55 603	768	72	32	41	$\epsilon^\star=16$	258.7	-	37	11	$\epsilon^\star=4$	257.4	-

assignment schemes. For a fixed RM and n_{out} value, we check the DFR of several RS codes of length n_{out}. Then, we perform a similar search for the second level to find a code $\mathcal{C}$ which meets the DFR requirement. By reducing n_{out}, we can reduce both the public-key length n bits (we choose a primitive prime larger than $n_{\text{out}}n_{\text{in}}$) and the ciphertext length $n+n_{\text{out}}n_{\text{in}}$ bits. For NIST security level 1, we found that public-key and ciphertext lengths can be reduced by 4.34 %, which is on par with the improvements achieved in [14]. Finally, we confirmed the DFR of the threshold-based decoder using Model 2 as described in Section 4.4.

6 Conclusion

In this work, we proposed a two-level generalized concatenated code that exploits the repetition subcode within HQC's inner Reed–Muller code, together with a reliability-based errors-and-erasures decoding framework for the outer codes. We derived rigorous upper bounds on the DFR for both threshold-based and partition-based erasure assignment, and identified parameters that reduce public-key and ciphertext sizes by up to 4.34 % for NIST security level 1.

Several directions for further improvement are possible. The unbalanced encryption weights proposed in [10] are orthogonal to both our GCC structure and our decoder; combining them with the techniques of this work would yield additional size reductions. Similarly, ciphertext compression [14] is compatible with our construction and could further reduce bandwidth usage. Finally, analyzing the DFR of the partition-based errors-and-erasures decoder under an advanced error model of [14–16] remains an open question.

References

[1] Gaborit, P., Aguilar Melchor, C., Aragon, N., Bettaieb, S., Bidoux, L., Blazy, O.,

Deneuville, J.-C., Persichetti, E., Zémor, G., Bos, J., Dion, A., Lacan, J., Robert, J.-M., Veron, P., Barreto, P., Gueron, S., Güneysu, T., Misoczki, R., Sendrier, N., Tillich, J.-P., Vasseur, V., Ghosh, S., Richter-Brockmann, J.: HQC. Technical report, National Institute of Standards and Technology (2025). version of August 2025, available at https://pqc-hqc.org/doc/hqc_specifications_2025_08_22.pdf

[2] Alagic, G., Bros, M., Ciadoux, P., Cooper, D., Dang, Q., Dang, T., Kelsey, J., Lichtinger, J., Liu, Y.-K., Miller, C., Moody, D., Peralta, R., Perlner, R., Robinson, A., Silberg, H., Smith-Tone, D., Waller, N.: Status report on the fourth round of the NIST post-quantum cryptography standardization process. NIST Interagency/Internal Report (NISTIR) NIST IR 8545, National Institute of Standards and Technology, Gaithersburg, MD (2025). https://doi.org/10.6028/NIST.IR.8545

[3] Guo, Q., Johansson, T., Stankovski, P.: A key recovery attack on MDPC with CCA security using decoding errors. In: International Conference on the Theory and Application of Cryptology and Information Security, pp. 789–815 (2016). Springer

[4] Guo, Q., Johansson, T.: A new decryption failure attack against hqc. In: International Conference on the Theory and Application of Cryptology and Information Security, pp. 353–382 (2020). Springer

[5] Aguilar-Melchor, C., Blazy, O., Deneuville, J.-C., Gaborit, P., Zémor, G.: Efficient encryption from random quasi-cyclic codes. IEEE Transactions on Information Theory **64**(5), 3927–3943 (2018)

[6] Zinov'ev, V.A.: Generalized cascade codes. Problemy Peredachi Informatsii **12**(1), 5–15 (1976)

[7] Blokh, È.L., Zyablov, V.V.: Coding of generalized concatenated codes. Problemy Peredachi Informatsii **10**(3), 45–50 (1974)

[8] Puchinger, S., Müelich, S., Bossert, M., Hiller, M., Sigl, G.: On error correction for physical unclonable functions. In: SCC 2015; 10th International ITG Conference on Systems, Communications and Coding, pp. 1–6 (2015). VDE

[9] Aguilar-Melchor, C., Aragon, N., Deneuville, J.-C., Gaborit, P., Lacan, J., Zémor, G.: Efficient error-correcting codes for the HQC post-quantum cryptosystem. Designs, Codes and Cryptography **92**(12), 4511–4530 (2024)

[10] Wang, T., Teng, Q., Wang, A., Zhang, J., Pang, B., Zhao, C., Hu, S., Wang, X.: HARE: Compact HQC via Distance-Informed Erasure Decoding. Cryptology ePrint Archive, Paper 2026/544 (2026). https://eprint.iacr.org/2026/544

[11] Cai, J., Zhang, X.: HQC Post-Quantum Cryptography Decryption with Generalized Minimum-Distance Reed-Solomon Decoder (2026). https://arxiv.org/abs/

2603.20156

[12] Bossert, M.: Channel Coding for Telecommunications. John Wiley & Sons, Chichester, England (1999)

[13] MacWilliams, F.J., Sloane, N.J.A.: The Theory of Error-correcting Codes vol. 16. North-Holland, Amsterdam, The Netherlands (1977)

[14] Bitzer, S., Deneuville, J.-C., Munisamy, E., Purtipli, B., Ritterhoff, S., Wachter-Zeh, A.: HQC beyond the standard: Ciphertext compression and refined dfr analysis. Annual International Conference on the Theory and Applications of Cryptographic Techniques (2026). Springer

[15] Annechini, A., Barenghi, A., Pelosi, G.: On the decoding failure rate of HQC. In: 2026 IEEE International Symposium on Information Theory, pp. 1–6 (2026)

[16] Mesnard, A., Tillich, J.-P., Vasseur, V.: The syndrome weight distribution in quasi-cyclic codes, applications to BIKE and HQC. In: International Conference on Post-Quantum Cryptography, pp. 156–187 (2026). Springer

[17] Dumer, I.I.: Concatenated codes and their multilevel generalizations. Handbook of coding theory (1998)

[18] Kasami, T., Tokura, N.: On the weight structure of Reed-Muller codes. IEEE Transactions on Information Theory **16**(6), 752–759 (1970)

[19] Agrawal, D., Vardy, A.: Generalized minimum distance decoding in Euclidean space: Performance analysis. IEEE Transactions on Information Theory **46**(1), 60–83 (2000)

[20] Weber, J.H., Abdel-Ghaffar, K.A.: Reduced GMD decoding. IEEE Transactions on Information Theory **49**(4), 1013–1027 (2003)

[21] Maillet, N., Nugier, C., Migliore, V.: Key recovery from side-channel power analysis attacks on non-SIMD HQC. In: CRYPTO 2025, vol. 16004, p. 70 (2025). Springer Nature

[22] Fossorier, M., Liu, S.: Error performance analysis for reliability-based decoding algorithms. In: 2000 IEEE International Symposium on Information Theory (2000)