The Quantum Decoding Problem

André Chailloux
Inria de Paris, France

Jean-Pierre Tillich
Inria de Paris, France

Abstract

One of the founding results of lattice based cryptography is a quantum reduction from the Short Integer Solution (SIS) problem to the Learning with Errors (LWE) problem introduced by Regev. It has recently been pointed out by Chen, Liu and Zhandry [12] that this reduction can be made more powerful by replacing the LWE problem with a quantum equivalent, where the errors are given in quantum superposition. In parallel, Regev's reduction has recently been adapted in the context of code-based cryptography by Debris, Remaud and Tillich [14], who showed a reduction between the Short Codeword Problem and the Decoding Problem (the DRT reduction). This motivates the study of the Quantum Decoding Problem (QDP), which is the Decoding Problem but with errors in quantum superposition and see how it behaves in the DRT reduction.

The purpose of this paper is to introduce and to lay a firm foundation for QDP. We first show QDP is likely to be easier than classical decoding, by proving that it can be solved in *quantum polynomial time* in a large regime of noise whereas no non-exponential quantum algorithm is known for the classical decoding problem. Then, we show that QDP can even be solved (albeit not necessarily efficiently) beyond the information theoretic Shannon limit for classical decoding. We give precisely the largest noise level where we can solve QDP giving in a sense the information theoretic limit for this new problem. Finally, we study how QDP can be used in the DRT reduction. First, we show that our algorithms can be properly used in the DRT reduction showing that our quantum algorithms for QDP beyond Shannon capacity can be used to find minimal weight codewords in a random code. On the negative side, we show that the DRT reduction cannot be, in all generality, a reduction between finding small codewords and QDP by exhibiting quantum algorithms for QDP where this reduction entirely fails. Our proof techniques include the use of specific quantum measurements, such as q-ary unambiguous state discrimination and pretty good measurements as well as strong concentration bounds on weight distribution of random shifted dual codes, which we relate using quantum Fourier analysis.

2012 ACM Subject Classification Theory of computation → Quantum information theory; Theory of computation → Error-correcting codes; Security and privacy → Cryptanalysis and other attacks

Keywords and phrases quantum information theory, code-based cryptography, quantum algorithms

Digital Object Identifier 10.4230/LIPIcs.TQC.2024.6

Related Version *Full Version*: `https://arxiv.org/pdf/2310.20651`

1 General cryptographic context

Error correcting codes appeared first as the fundamental tool to transmit information reliably through a noisy channel [25] and has found numerous applications in information theory and complexity. The hardness - even for quantum computers - of decoding random linear codes is also the core of code-based cryptography. In the cryptographic context, the decoding problem corresponds to decoding the k-dimensional vector space $\mathscr{C}$ (*i.e.*, the code) generated by the rows of a randomly generated $\mathbf{G} \in \mathbb{F}_q^{k\times n}$ (which is called a *generating matrix* of the code):

$$\mathscr{C} \stackrel{\triangle}{=} \left\{ \boldsymbol{u}\mathbf{G} \colon \boldsymbol{u} \in \mathbb{F}_q^k \right\}. \tag{1}$$

Here $\mathbb{F}_q$ denotes the finite field with q elements. In the decoding problem, we are given the noisy codeword $\boldsymbol{c}+\boldsymbol{e}$ where $\boldsymbol{c}$ belongs to $\mathscr{C}$ and we are asked to find the original codeword $\boldsymbol{c}$.

▶ **Problem 1** ($\mathsf{DP}(q,n,k,f)$)**.** *The decoding problem with positive integer parameters q,n,k and a probability distribution f on $\mathbb{F}_q^n$ is defined as:*

- *Input:* $(\mathbf{G},\boldsymbol{c}+\boldsymbol{e})$ *where* $\mathbf{G}\in\mathbb{F}_q^{k\times n}$ *and* $\boldsymbol{u}\in\mathbb{F}_q^k$ *are sampled uniformly at random over their domain - which generates a random codeword* $\boldsymbol{c}=\boldsymbol{u}\mathbf{G}$ *- and* $\boldsymbol{e}$ *is sampled from the distribution* f*.*
- *Goal: from* $(\mathbf{G},\boldsymbol{c}+\boldsymbol{e})$*, find* $\boldsymbol{c}$*.*

This problem for random codes has been studied for a long time and despite many efforts on this issue, the best (even quantum) algorithms are exponential in the codelength n for natural noise distributions f in the regime where k is linear in n and the rate $R\triangleq\frac{k}{n}$ bounded away from 0 and 1 [23, 27, 15, 19, 5, 20, 18, 9, 8]. This remains true even if we consider quantum algorithms

The most common noise distribution studied in this context is the uniform distribution over the errors of fixed Hamming weight t, but there are also other distributions, like in the binary case ($q=2$) the i.i.d Bernoulli distribution model which is frequently found in the Learning Parity with Noise problem (LPN) [16]. In code-based cryptography, the regime which is almost always relevant is a fixed number of samples n (or codelength) in the linear regime *i.e.* $k=\Theta(n)$. We will focus on this case here. While the security of many code-based cryptosystems relies on the hardness of the decoding problem, it can also be based on finding a "short" codeword (as in [21] or in [2, 7, 29] to build collision resistant hash functions), a problem which is stated as follows.

▶ **Problem 2** ($\mathsf{SCP}(q,n,k,w)$)**.** *The short codeword problem with parameters $q,n,k,w\in\mathbb{N}$ is defined as:*

- *Given:* $\mathbf{H}\in\mathbb{F}_q^{(n-k)\times n}$ *which is sampled uniformly at random,*
- *Find:* $\boldsymbol{c}\in\mathbb{F}_q^n\setminus\{\mathbf{0}\}$ *such that* $\mathbf{H}\boldsymbol{c}^\intercal=\mathbf{0}$ *and the weight* $|\boldsymbol{c}|$ *of* $\boldsymbol{c}$ *satisfies* $|\boldsymbol{c}|\leq w$*.*

Here we are looking for a non-zero codeword $\boldsymbol{c}$ of weight $\leq w$ in the k-dimensional code $\mathscr{C}$ defined by the so-called parity-check matrix $\mathbf{H}$, namely[1] :

$$\mathscr{C}\triangleq\left\{\boldsymbol{c}\in\mathbb{F}_q^n\colon \mathbf{H}\boldsymbol{c}^\intercal=\vec{0}\right\}.$$

The weight function which is generally used here is the Hamming weight, *i.e.* for a vector $\boldsymbol{x}=(x_1,\cdots,x_n)\in\mathbb{F}_q^n$, its Hamming weight is defined as

$$|\boldsymbol{x}|\triangleq\#\{i\in[\![1,n]\!]:x_i\neq 0\}.$$

We will only deal with this weight here. Decoding and looking for short codewords are problems that have been conjectured to be extremely close. They have been studied for a long time, and the best algorithms for solving these two problems are the same, namely Information Set Decoding algorithms [23, 27, 15, 5, 20, 6].

Recently, Debris-Alazard, Remaud and Tillich showed a quantum reduction from SCP to DP adapting Regev's reduction from the Short Integer Solution (SIS) problem to the Learning With Errors problem (LWE). It has recently been pointed out by Chen, Liu and

[1] The short codeword problem is usually defined by picking a random parity-check matrix $\mathbf{H}\in\mathbb{F}_q^{(n-k)\times n}$ and not a random generating matrix $\mathbf{G}\in\mathbb{F}_q^{k\times n}$ but the differences are minor (see for example [13]) and one could also define this problem via the generating matrix of a code as we did for the decoding problem.

Zhandry [12] that Regev's reduction can be made more powerful by replacing the LWE problem with a quantum equivalent, where the errors are given in quantum superposition. It is therefore natural to ask whether having errors in quantum superposition in the decoding problem can be applied to the DRT reduction in order to improve it.

The purpose of this article is to introduce and to lay a firm foundation for the Quantum Decoding Problem, which is the decoding problem but with errors in quantum superposition. We first present the DRT reduction for codes, properly define QDP, and then present in detail our contributions.

2 Regev's quantum reduction and follow-up work

Regev's quantum reduction[24] is at the core of complexity reductions for these problems, which with [1] essentially started lattice-based cryptography. His approach when rephrased in the coding context is based on the following observation. Suppose that we were able to construct a quantum superposition $\sqrt{\frac{1}{Z}}\sum_{\boldsymbol{c}\in\mathcal{C}}\sum_{\boldsymbol{e}\in\mathbb{F}_q^n}\sqrt{f(e)}|\boldsymbol{c}+\boldsymbol{e}\rangle$ of noisy codewords of some code $\mathcal{C}$ over $\mathbb{F}_q$, for a normalization factor Z. By applying the quantum Fourier transform on such a state, because of the periodicity property of such a state, we get a superposition concentrating solely on the codewords of the dual $\mathcal{C}^\perp$ of $\mathcal{C}$, that is $\frac{1}{\sqrt{Z}}\sum_{\boldsymbol{c}^\perp\in\mathcal{C}^\perp}\sqrt{\widehat{f}(\boldsymbol{c}^\perp)}|\boldsymbol{c}^\perp\rangle$. Here $\widehat{f}$ is the (classical) Fourier transform of f. Recall that the dual code is defined as

▶ **Definition 1** (dual code). *Let $\mathcal{C}$ be a k-dimensional linear code over $\mathbb{F}_q$ of length n. Let $\boldsymbol{x}\cdot\boldsymbol{y}$ be the inner product of vectors $\boldsymbol{x}$, $\boldsymbol{y}$ in $\mathbb{F}_q^n$ defined as $\boldsymbol{x}\cdot\boldsymbol{y}=\sum_i x_iy_i$. The dual code $\mathcal{C}^\perp$ is an $(n-k)$ dimensional subspace of $\mathbb{F}_q^n$ defined by $\mathcal{C}^\perp\stackrel{\triangle}{=}\{\boldsymbol{d}\in\mathbb{F}_q^n:\boldsymbol{d}\cdot\boldsymbol{c}=0,\ \forall\boldsymbol{c}\in\mathcal{C}\}$*

We can expect that if f concentrates on fairly small weights, then $\widehat{f}$ also concentrates on small weights. This gives a way of sampling low weight (dual) codewords and solve SCP for the dual code. The point is now that $\sqrt{\frac{1}{Z}}\sum_{\boldsymbol{c}\in\mathcal{C}}\sum_{\boldsymbol{e}\in\mathbb{F}_q^n}\sqrt{f(e)}|\boldsymbol{c}+\boldsymbol{e}\rangle$ can be obtained by solving DP on states that are easy to construct. This is the main idea of the DRT reduction. More precisely, the whole algorithm works as

Step 1. Creation of a superposition of noise tensored with a uniform superposition of codewords

$$|\phi_1\rangle=\sqrt{\frac{1}{q^k}}\sum_{\boldsymbol{e}\in\mathbb{F}_q^n}\sqrt{f(e)}|\boldsymbol{e}\rangle\sum_{\boldsymbol{c}\in\mathcal{C}}|\boldsymbol{c}\rangle.$$

Step 2. Entangling the codeword with the noise by adding the second register to the first one

$$|\phi_2\rangle=\sqrt{\frac{1}{q^k}}\sum_{\boldsymbol{c}\in\mathcal{C}}\sum_{\boldsymbol{e}\in\mathbb{F}_q^n}\sqrt{f(e)}|\boldsymbol{c}+\boldsymbol{e}\rangle|\boldsymbol{c}\rangle.$$

Step 3. Disentangling the two registers by decoding $\boldsymbol{c}+\boldsymbol{e}$ and therefore finding $\boldsymbol{c}$ which allows to erase the second register (a different normalization Z arises when decoding is imperfect and we condition on measuring $\mathbf{0}$ in the last register).

$$|\phi_3\rangle=\sqrt{\frac{1}{Z}}\sum_{\boldsymbol{c}\in\mathcal{C}}\sum_{\boldsymbol{e}\in\mathbb{F}_q^n}\sqrt{f(e)}|\boldsymbol{c}+\boldsymbol{e}\rangle|\mathbf{0}\rangle.$$

Step 4. Applying the quantum Fourier transform on the first register and get

$$\frac{1}{\sqrt{Z}} \sum_{\boldsymbol{d} \in \mathcal{C}^\perp} \sqrt{\widehat{f(\boldsymbol{d})}} |\boldsymbol{d}\rangle |\mathbf{0}\rangle$$

Step 5. Measure the first register and get some $\boldsymbol{d}$ in $\mathcal{C}^\perp$.

This approach is at the heart of the quantum reductions obtained in [24, 26, 14]. It is also a crucial ingredient in the paper [28] proving verifiable quantum advantage by constructing - among other things - one-way functions that are even collision resistant against classical adversaries but are easily invertible quantumly. In [24, 26, 14], the crucial erasing/disentangling step is performed with the help of a *classical* decoding algorithm. Indeed any (classical or quantum) algorithm that can recover $\boldsymbol{c}$ from $\boldsymbol{c}+\boldsymbol{e}$ can be applied coherently to erase the last register in step 3[2] .

A key insight observed in [12] is that it is actually enough to recover $|\boldsymbol{c}\rangle$ from the state $\sum_{\boldsymbol{e} \in \mathbb{F}_q^n} \sqrt{f(\boldsymbol{e})} |\boldsymbol{c}+\boldsymbol{e}\rangle$ so we are given a superposition of all the noisy codewords $\boldsymbol{c}+\boldsymbol{e}$ and not a fixed one. This means we have to solve the following problem

▶ **Problem 3** (QDP(q,n,k,f))**.** *The quantum decoding problem with positive integer parameters q,n,k and a probability distribution f on $\mathbb{F}_q^n$ is defined as:*

- *Input: Take $\mathbf{G} \in \mathbb{F}_q^{k \times n}$ and $\boldsymbol{u} \in \mathbb{F}_q^k$ sampled uniformly at random over their domain. Let $\boldsymbol{c} = \boldsymbol{u}\mathbf{G}$ and $|\psi_{\boldsymbol{c}}\rangle \stackrel{\triangle}{=} \sum_{\boldsymbol{e} \in \mathbb{F}_q^n} \sqrt{f(e)} |\boldsymbol{c}+\boldsymbol{e}\rangle$. The (quantum) input to this problem is $(\mathbf{G}, |\psi_{\boldsymbol{c}}\rangle)$.*
- *Goal: given $(\mathbf{G}, |\psi_{\boldsymbol{c}}\rangle)$, find $\boldsymbol{c}$.*

It's not clear a priori whether this is helpful or not. If one measures the state $|\psi_{\boldsymbol{c}}\rangle$ then one recovers a noisy codeword and we are back to the classical decoding problem. However, in the context of lattices, [12] showed that this approach can lead to improvements. A final small remark on the motivation of the quantum decoding problem. We are not in the context of noise coming from a realistic quantum channel so we do not need our noise model to emulate real quantum noise (which would certainly not be a q-ary symmetric channel with the same phases). The motivation of this definition really comes from an algorithmic and complexity perspective, as well as a quantum information-theoretic perspective but not from a quantum error correcting perspective.

3 Contributions

Here we focus on the noise model which is relevant for the Hamming metric in SCP, namely the Bernoulli noise of parameter p. This means we consider the error function

$$f(\boldsymbol{e}) = (1-p)^{n-|\boldsymbol{e}|} \left(\frac{p}{q-1}\right)^{|\boldsymbol{e}|}.$$

which in turn means that for any $\boldsymbol{c} = (c_1, \ldots, c_n) \in \mathbb{F}_q^n$, we can rewrite

$$|\psi_{\boldsymbol{c}}\rangle \stackrel{\triangle}{=} \sum_{\boldsymbol{e} \in \mathbb{F}_q^n} \sqrt{f(\boldsymbol{e})} |\boldsymbol{c}+\boldsymbol{e}\rangle = \bigotimes_{i=1}^{n} \left(\sqrt{1-p} |c_i\rangle + \sum_{\alpha \in \mathbb{F}_q^*} \sqrt{\frac{p}{q-1}} |c_i + \alpha\rangle \right).$$

[2] Indeed, having such an algorithm means we can construct the unitary $U : |\boldsymbol{c}+\boldsymbol{e}\rangle|0\rangle \to |\boldsymbol{c}+\boldsymbol{e}\rangle|\boldsymbol{c}\rangle$. Applying the inverse of this unitary will give the erasure operation.

For this Bernoulli noise with parameter p, the associated quantum decoding problem is written $\mathrm{QDP}(q, n, k, p)$. We will lay out a firm foundation for this quantum decoding problem by focusing on the case which is generally relevant in code-based cryptography, namely in the fixed code rate $R \stackrel{\triangle}{=} k/n$ regime and will show that QDP *departs significantly* from the classical decoding problem because we show that

- **(i)** QDP is likely to be easier than classical decoding, by proving that it can be solved in *polynomial time* in a large regime of noise whereas no non-exponential algorithm is known for the classical decoding problem.
- **(ii)** the problem can even be solved (albeit not necessarily efficiently) beyond the information theoretic Shannon limit for classical decoding. We will give precisely the largest noise level where we can solve QDP giving in a sense the information theoretic limit for the new problem.
- **(iii)** We study how these QDP algorithms fit in the DRT reduction. We show using our quantum polynomial algorithm in this reduction in order to obtain quantum polynomial algorithms recovering Prange's bound. Even more interestingly, we show that our quantum algorithm for QDP of (ii) can be used in order to find small codewords of weight as small as the *minimal distance* of the code, which is the best we can hope for. On the negative side, we show that the DRT reduction cannot be, in all generality, a reduction between QDP and finding small codewords by exhibiting quantum algorithms for QDP that make this reduction entirely fail.

We now perform a detailed description of our contributions.

3.1 Polynomial time quantum algorithm for QDP in a large parameter regime

Our first result is the following

▶ **Theorem 2.** *Let $R \in [0,1]$. For any $p < \left(\frac{(q-1)R}{q}\right)^{\perp}$, there exists a quantum algorithm that solves $\mathrm{QDP}(q, n, \lfloor Rn \rfloor, p)$ wp. $1 - 2^{-\Omega(n)}$ in time $\mathrm{poly}(n, \log(q))$, where for a real number $x \in [0,1]$, $x^{\perp} = \frac{\left(\sqrt{(1-x)(q-1)} - \sqrt{x}\right)^2}{q}$.*

Let us dive in how we obtain this result. We start from an input of $\mathrm{QDP}(q, n, \lfloor Rn \rfloor, p)$, which, for an (unknown) codeword $\boldsymbol{c} = c_1, \ldots, c_n$ is the state

$$|\Psi_{\boldsymbol{c}}\rangle = \bigotimes_{i=1}^{n} |\psi_{c_i}\rangle \quad \text{with} \quad |\psi_{c_i}\rangle \stackrel{\triangle}{=} \sqrt{1-p}|c_i\rangle + \sum_{\alpha \in \mathbb{F}_q^*} \sqrt{\frac{p}{q-1}}|c_i + \alpha\rangle$$

and the goal is to recover $\boldsymbol{c}$. Our algorithm performs Unambiguous State Discrimination (USD) on each of the n registers. USD is a quantum measurement that, on input $|\psi_{c_i}\rangle$, will output or the correct value c_i, or an abort symbol $\perp$ but will never output a value $\alpha \in \mathbb{F}_q$ different from c_i. Then, if we have enough correct values of c_i (essentially more than $\lceil Rn \rceil$), then one can recover the whole $\boldsymbol{c}$ using the description of the code and basic linear algebra.

Optimal unambiguous state discrimination is very well understood in the binary case ($q = 2$) but is not known in general for more than 2 states. In certain situations where we have a symmetric set of states [10] we know how to perform optimal USD. This would apply in our case case where q is prime. We have generalized the approach of [10] to be able to apply it to any finite field size q, and prove the following.

▶ **Proposition 3.** *Let q be a prime power, and $f : \mathbb{F}_q \to \mathbb{C}$ st. $\|f\|_2 = 1$. For each $y \in \mathbb{F}_q$, we define $|\psi_y\rangle = \sum_{\alpha\in\mathbb{F}_q} f(\alpha)|y+\alpha\rangle$. There exists a quantum measurement that, when given $|\psi_y\rangle$, outputs y wp. p_{USD} and $\perp$ wp. $1-p_{USD}$ where $p_{USD} = q \cdot \min_{\alpha\in\mathbb{F}_q} |\widehat{f}(\alpha)|^2$, and this is optimal. Moreover, if f corresponds to a Bernoulli noise of parameter p, this measurement can be done in time* $\mathrm{polylog}(q)$ *and we have $p_{USD} = \frac{qp^\perp}{q-1}$, where $p^\perp = \frac{\left(\sqrt{(1-p)(q-1)}-\sqrt{p}\right)^2}{q}$.*

Notice that [12] proposed a USD measurement with $p_{USD} = \frac{1}{q}\min_{\alpha\in\mathbb{F}_q} |\widehat{f}(\alpha)|^2$. Their measurement does not scale well with q contrarily to our measurement which has basically the right scaling with q. For instance, [12] requires that $q = \mathrm{poly}(n)$. We have no such restriction in our case and our algorithms work in polynomial time even for $q = 2^{\Omega(n)}$.

3.1.1 Interpretation as changing the noise channel

A nice interpretation of the above algorithm is that when the error is in quantum superposition, one can use quantum measurements to change the noise model. For instance, when we are given $|\psi_{c_i}\rangle = \sqrt{1-p}|c_i\rangle + \sum_{\alpha\in\mathbb{F}_q^*}\sqrt{\frac{p}{q-1}}|c_i+\alpha\rangle$ then

- One can measure in the computational basis to obtain c_i that has been flipped wp. p to one of the other $q-1$ values at random.
- One can use unambiguous state discrimination in which case c_i has been erased wp. $1-\frac{qp^\perp}{q-1}$.

What we show in Theorem 2 is that the second strategy is actually much more powerful for recovering the codeword $\boldsymbol{c}$. A natural question to ask is whether this can further be generalized to other measurements. In the binary setting we actually generalize USD as follows: given $|\psi_{c_i}\rangle$, the measurement sometimes outputs $\perp$ but it can also fail with some small probability. We prove the following

▶ **Proposition 4** (Partial USD). *Let $p, s \in [0, \frac{1}{2})$ with $s \le p$ and let $u = \frac{p^\perp}{s^\perp}$. There exists a quantum measurement that when applied to $|\psi_{c_i}\rangle = \sqrt{1-p}|c_i\rangle + \sqrt{p}|1-c_i\rangle$ outputs c_i wp. $u(1-s)$, $(1-c_i)$ wp. us and $\perp$ wp. $1-u$.*

Notice that this generalizes both the computational basis measurement (by taking $s = p$) and unambiguous state discrimination (by taking $s = 0$ giving $u = 2p^\perp$). This seems a very natural way of generalizing USD but is not something we have found in the literature and could be of independent interest. We can use this measurement not to provide new polynomial time algorithms but rather to give a reduction between different Quantum Decoding problems, which we detail in the full text.

3.2 Determining exactly the tractability of the quantum decoding problem

We are now interested in the tractability of $\mathrm{QDP}(q, n, k, p)$ meaning when is it possible from an information theoretic perspective to solve this problem. A fundamental quantity is relavant here, namely the Gilbert-Varshamov distance $\delta_{\min}(R)$ defined below

▶ **Notation 1.** *Let $R \in [0,1]$. We define $\delta_{\min}(R) \triangleq h_q^{-1}(1-R)$, where $h_q(x) \triangleq -(1-x)\log_q(1-x) - x\log_q\left(\frac{x}{q-1}\right)$. h_q is a bijection from $x \in \left[0, \frac{q-1}{q}\right]$ to $[0,1]$ and $h_q^{-1} : [0,1] \to \left[0, \frac{q-1}{q}\right]$ is the inverse of h_q.*

For the classical setting, it is well understood that $\mathrm{DP}(q,n,k,p)$ is not tractable when $p > \delta_{\min}(\frac{k}{n})$, meaning that even an unbounded algorithm will solve the problem wp. $o(1)$. We would like now to understand what happens in the quantum setting. Techniques based on (partial) USD will not work in the regime $p > \delta_{\min}(R)$. Since we are only interested in the tractability of the problem, we can consider optimal quantum algorithms for discriminating between the states $|\Psi_{\boldsymbol{c}}\rangle = \sum_{\boldsymbol{e}\in\mathbb{F}_q^n} \sqrt{f(\boldsymbol{e})}|\boldsymbol{c}+\boldsymbol{e}\rangle$ where f accounts for the Bernoulli noise of parameter p. This problem can be addressed by using the *Pretty Good Measurement* (PGM) which has turned out to be a very useful tool in quantum information. If we define $\mathrm{P}_{\mathrm{PGM}}$ as the probability that the pretty good measurement succeeds in solving our problem and define $\mathrm{P}_{\mathrm{OPT}}$ as the maximal probability that any measurement succeeds, we have [4, 22]

$$\mathrm{P}_{\mathrm{OPT}}^2 \leq \mathrm{P}_{\mathrm{PGM}} \leq \mathrm{P}_{\mathrm{OPT}}.$$

This means that in order to study the tractability of the quantum decoding problem, it is enough to look at the PGM associated with the problem of distinguishing the states $\{|\Psi_{\boldsymbol{c}}\rangle\}$. We show that

▶ **Theorem 5.** *Let $R \in (0,1)$.*

- *For $p < (\delta_{\min}(1-R))^{\perp}$, $\mathrm{QDP}(q,n,\lfloor Rn\rfloor,p)$ can be solved using the PGM wp. $P_{PGM} = 1-o(1)$ hence the problem is tractable.*
- *For $p > (\delta_{\min}(1-R))^{\perp}$, the probability that the PGM solves this problem is $P_{PGM} = o(1)$ hence $P_{OPT} = o(1)$ and the problem is intractable.*

In order to prove this theorem, we study the Pretty Good Measurement associated to the states $|\Psi_{\boldsymbol{c}}\rangle$ which are the possible inputs of QDP. In order to study our PGM, we define the shifted dual codes of $\mathcal{C}$

$$\mathcal{C}_{\boldsymbol{s}}^{\perp} \stackrel{\triangle}{=} \{\boldsymbol{x} \in \mathbb{F}_q^n : \mathbf{G}\boldsymbol{x}^{\intercal} = \boldsymbol{s}\}$$

We show the following

▶ **Proposition 6.** *We consider the Pretty Good Measurement associated to the states $\{|\Psi_{\boldsymbol{c}}\rangle\}_{\boldsymbol{c}\in\mathcal{C}}$ with $|\Psi_{\boldsymbol{c}}\rangle = \sum_{\boldsymbol{e}} f(\boldsymbol{e})|\boldsymbol{c}+\boldsymbol{e}\rangle$. This measurement outputs $\boldsymbol{c}$ given $|\Psi_{\boldsymbol{c}}\rangle$ wp.*

$$p_{PGM} = \frac{1}{q^k}\left(\sum_{\boldsymbol{s}\in\mathbb{F}_q^k} n_{\boldsymbol{s}}\right)^2 \quad \text{where} \quad n_{\boldsymbol{s}} = \sqrt{\sum_{\boldsymbol{y}\in\mathcal{C}_{\boldsymbol{s}}^{\perp}} |\hat{f}(\boldsymbol{y})|^2}.$$

This shows an interesting and unexpected link between the Pretty Good Measurement associated to the states $\{|\Psi_{\boldsymbol{c}}\rangle\}$ and the shifted dual codes $\mathcal{C}_{\boldsymbol{s}}^{\perp}$. In the particular case of a Bernoulli noise of parameter p, the value of $n_{\boldsymbol{s}}$ will be dominated by a quantity related to the number of words of weight close to $p^{\perp}$ in the shifted dual code $\mathcal{C}_{\boldsymbol{s}}^{\perp}$. In order to conclude, we use strong concentration bounds on the weight distribution of shifted dual codes.

3.2.1 Comparing the complexity of DP and QDP

With this full characterization, we compare the hardness, and tractability of the classical and quantum decoding problems. For $p = 0$, we have of course a polynomial time algorithm to solve $\mathrm{DP}(q,n,\lfloor Rn\rfloor,0)$. For $0 < p \leq \delta_{\min}(R)$, the problem is tractable and the best known classical or quantum algorithms run in time $2^{\Omega(n)}$. For $p > \delta_{\min}(R)$, we know the problem is intractable. For the Quantum Decoding Problem, we obtain a very different picture. A comparison of these results is presented in Figures 1 and 2 where we use the following terminology

- Easy: there exists an algorithm that runs in time $\text{poly}(n)$.
- Hard: the best known (classical or quantum) algorithm runs in time $2^{\Omega(n)}$, but there could potentially be more efficient algorithms.
- Intractable: we know that any (even unbounded) algorithm can solve the problem wp. at most $o(1)$.

Figure 1 Hardness and tractability of the decoding problem $\text{DP}(q, n, \lfloor Rn \rfloor, p)$, for any fixed $R \in [0, 1]$, as a function of p.

Figure 2 Hardness and tractability of the quantum decoding problem $\text{QDP}(q, n, \lfloor Rn \rfloor, p)$, for any fixed $R \in [0, 1]$, as a function of p.

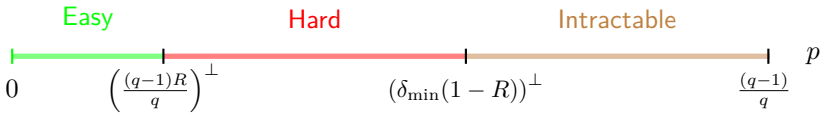

This gives a proper characterization of the difficulty of QDP. In our next contribution, we will apply them in the DRT quantum reduction in order to derive some results for the short codeword problem. As we will show, the results from Figure 2 will match exactly our knowledge for the short codeword problem.

3.3 Using our algorithms in Regev's reduction

We are now interested in solving the short codeword problem using Regev's reduction and the algorithms we described in the previous section. The known (classical and quantum) hardness of the short codeword problem is summarized in Fig. 3. In our coding context, the

Figure 3 Hardness and tractability of the short codeword problem $\text{SCP}(q, n, \lfloor Rn \rfloor, p)$ for a fixed $R \in (0, 1)$, as a function of p.

only known reduction is the following

▶ **Proposition 7** ([14], informal). *Fix integers $n, q \geq 2$ as well as parameters $R, p \in (0, 1)$ st. $p \leq \delta_{\min}(R)$. From any quantum algorithm that solves $DP(q, n, \lceil (1-R)n \rceil, p)$ with high probability, there exists a quantum algorithm that solves $SCP(q, n, \lfloor Rn \rfloor, p^{\perp})$ with high probability where recall that $p^{\perp} = \frac{\left(\sqrt{(1-p)(q-1)} - \sqrt{p}\right)^2}{q}$.*

In some sense, this reduction is far from tight. Indeed, if we take the best known algorithms for DP (see Figure 1) and we apply the above proposition, we get quantum algorithms much worst than the best known ones from Figure 3. On the other hand, if we could plug in our algorithms for QDP in this reduction, we would obtain quantum algorithms

for SCP that have the same complexities as the one in Figure 3. From our discussion in Section 2, it seems that one could perform the same reduction as above but replacing DP with QDP. The reality is quite more tricky. Indeed, if the quantum algorithm for QDP succeeds wp. 1 then the reduction works. However, even a small error in the quantum algorithm for QDP can lead to large error in the corresponding algorithms for SCP.

We first show that this is not an issue when we use the quantum algorithms for QDP we described in the previous section in the DRT reduction.

▶ **Theorem 8.** *For any* $p < \left(\frac{(q-1)R}{q}\right)^{\perp}$*, if we plug the quantum polynomial time algorithm of Theorem 2 for* $\mathrm{QDP}(q, n, \lfloor (1-R)n \rfloor, p^{\perp})$ *in Regev's reduction, we obtain a quantum polynomial time algorithm for* $SCP(q, n, \lfloor R \rfloor, p)$*.*

▶ **Theorem 9.** *For any* $p < \delta_{min}(1-R)$*, if we plug (a slight variant of) the quantum algorithm of Theorem 5 for* $\mathrm{QDP}(q, n, \lfloor (1-R)n \rfloor, p^{\perp})$ *in Regev's reduction, we obtain a quantum polynomial time algorithm for* $SCP(q, n, \lfloor R \rfloor, p)$*.*

3.3.1 Efficiency of the reduction

We graphically compare here the DRT reduction using DP and using our algoirthms for QDP.

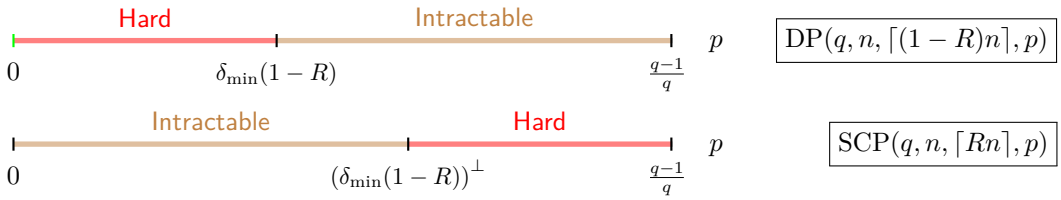

Figure 4 On the top, best known (classical or quantum) algorithms for $\mathrm{DP}(q, n, \lceil (1-R)n \rceil, p)$. On the bottom, complexity of a quantum algorithm for $\mathrm{SCP}(q, n, \lceil Rn \rceil, p)$ that uses the best algorithm for $\mathrm{DP}(q, n, \lceil (1-R)n \rceil, p)$ and then uses Proposition 7.

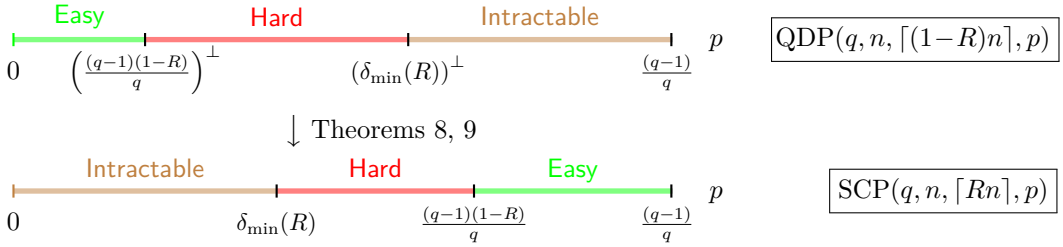

Figure 5 On the top, our quantum algorithms for $\mathrm{QDP}(q, n, \lceil (1-R)n \rceil, p)$. On the bottom, complexity of a quantum algorithm for $\mathrm{SCP}(q, n, \lceil Rn \rceil, p)$ that would use our algorithms $\mathrm{QDP}(q, n, \lceil (1-R)n \rceil, p)$ and then Theorems 8 and 9.

What we find quite remarkable is that our algorithm for QDP used at the limit of the tractability bound can be used to recover minimal weight codewords of weight $\delta_{\min}(R)$ in the dual. A natural question is whether we can have generic reductions between SCP and QDP. We show that the DRT reduction fails for this task and the increase in error in the reduction can be drastic.

▶ **Theorem 10.** *For any $p < \delta_{min}(1-R)$, there exists a quantum algorithm that solves* $\mathrm{QDP}(q, n, \lfloor (1-R)n \rfloor, p^{\perp})$ *wp.* $1 - o(1)$ *st. if we plug it in the DRT reduction, the resulting algorithm for* $SCP(q, n, \lfloor Rn \rfloor, p)$ *never succeeds.*

These results show that, while it is impossible to have a generic reduction from SCP to QDP with this method, it is (at least in our examples) possible to find algorithms for QDP that give results according to Fig. 5, and recover the areas where the problem is easy or tractable. This can be seen as quite a surprise since our bounds on QDP come from information theory and best known bounds on SCP comes from classical coding theory and seem unrelated at first.

4 Related work

Our main starting point is [12] so it natural to compare our contributions with this work. In [12], they introduce the S-$|\mathrm{LWE}\rangle$ problem, the lattice equivalent of QDP. They construct from it a quantum algorithm for SIS_∞ via Regev's reduction while our work focuses on QDP mainly for its own sake. Regarding their quantum polynomial time algorithm for S-$|\mathrm{LWE}\rangle$, it is obtained by performing a variant of Unambiguous State discrimination where we only rule out certain values for the code-symbols, and then they use the Arora-Ge algorithm [3] for recovering completely the codeword by solving an algebraic system which for the parameters that are considered there, is of polynomial complexity. Our quantum polynomial time algorithm is inspired by this approach but since we work with the q-ary Bernoulli noise, we can directly use unambiguous state discrimination. Also, we perform a more efficient q-ary USD which allows us to work even when $q = 2^{\Omega(n)}$ while the work of [12] works only when $q = \mathrm{poly}(n)$. Our other results on the (in)tractability results, as well as the discussion on the DRT reduction are entirely novel and do not have an equivalent version in [12].

Another quantum variant has been presented [17] but where they consider a quantum superposition of samples *i.e.* superpositions of generating matrices. They show that in this case the problem can be solved in quantum polynomial time. Their setting is very different from ours as we fix a code and do not have codes in superposition. Moreover, their techniques are not applicable to our setting.

A recent result [11] also studies the S-$|\mathrm{LWE}\rangle$ problem. They perform a quantum reduction between LWE and S-$|\mathrm{LWE}\rangle$ with extra unknown phases. The parameters of this S-$|\mathrm{LWE}\rangle$ are such that if they did not have these unknown phases, they could solve it with a subexponential quantum algorithm using a subexponential number of samples. This is very far from our parameter range and hence not comparable but gives an interesting use of Kuperberg's sieve for this kind of problem.

5 Discussion

5.1 A problem which is interesting in its own

Our work lays firm theoretical foundations for the Quantum Decoding Problem, where we find an interesting parameter range where the problem can be solved in quantum polynomial time. Moreover, we precisely characterize up to what level of noise the problem is tractable from an information theoretic point of view. Finally we show how our algorithms can be used in Regev's reduction for finding short codewords.

Beyond this, it seems to us that the quantum decoding problem is a natural and important problem in its own. We did not study the quantum decoding problem to relate it to the classical decoding problem so the aim of our results is not to say something about the

complexity of the classical decoding problem (even though it is linked to the related Short Codeword Problem via Regev's reduction). We are actually more interested in the differences between these two problems. Having errors in quantum superposition can be used to change the noise model from a q-ary symmetric channel to an erasure channel. We even have a complete characterization in the binary setting via our notion of partial unambiguous state discrimination. Also quite surprisingly, we can decode beyond classical information theoretic limits. Moreover, the optimal bound $(\delta_{\min}(1-R))^{\perp}$ that we exhibit in the binary setting is exactly the first linear programming which bounds the minimal distance of a code depending on its size so this bound, which comes from the study of the Pretty Good Measurement has links, again, with fundamental quantities from classical coding theory.

But as we said, we have another strong motivation for studying the quantum decoding problem. Indeed, it is the problem which directly appears when performing Regev's reduction. This means studying the quantum decoding problem also gives us a better understanding of this reduction both from a complexity point of view and from a quantum algorithmic point of view. From a complexity point of view, our results explain why this reduction between the Decoding Problem and the Short Codeword problem (or equivalently between LWE and SIS) gives far from tight results. It is because this reduction is genuinely a reduction between the Quantum Decoding Problem and the Short Codeword and our analysis shows the former is much simpler than its classical counterpart. Also, Figure 5 shows that with this reduction, we recover the polynomial zone and the tractability zone of the short codeword problem which shows in some sense the tightness of this reduction. The fact that bounds on optimal Unambiguous State Discrimination and on the Pretty Good Measurement leads via Regev's reduction to Prange's bound (for the polynomial case) and to the Gilbert-Varshamov bound (for the tractability bound) shows interesting and, in our opinion, quite aesthetic links between quantum and classical information theory.

From an algorithmic point of view, these results can be seen as a new class of quantum algorithms for the short codeword problem. One might say that we do not improve on existing algorithms. For example, our quantum polynomial algorithm finds short codeword for some weights t that can also be found by the classical Prange's algorithm. Let us just observe here that if we could find a polynomial quantum algorithm that would beat Prange's bound then that would significantly change our understanding of the quantum hardness of these problems and the post-quantum security claims of code-based cryptography (and may be even lattice-based cryptography) would be affected. In the exponential regime, we propose a new family of quantum algorithms for the short codeword problem and there are many directions (changing the noise function f, determining the complexity of measurements required in the QDP, strict analysis in Regev's reduction ...) that look very promising for future work.

5.2 Technical takeaways

In the first part of the paper, we use binary Unambiguous State Discrimination for constructing our quantum polynomial algorithm. Once the idea is found, the techniques used are known and simple. We then extend this to partial unambiguous state discrimination - where we still allow some probability of failure much less than in Helstrom's measurement. This is a technique that we have not seen previously in the literature and could be of independent interest. In the q-ary setting, we extend optimal bounds for unambiguous state discrimination of symmetric states to the case q is a prime power and also show how to construct this measurement in time $\log(q)$ for Bernoulli noise. The second part of the paper, which deals with (in)tractability bounds is arguably the most technical part of the paper. A first

interesting technical result was to precisely characterize the Pretty Good Measurement used in the quantum decoding problem of a code $\mathcal{C}$ as a projective measurement that involves the shifted dual codes of $\mathcal{C}$. Then, the most technical part was to actually compute the success probability of the Pretty Good Measurement as a function of the noise rate which requires precise and well used concentration and anti-concentration bounds on the weight distribution of shifted dual codes of a random code. In the third part of the paper, where we apply our algorithms to Regev's reduction, we mainly use our analysis of the Pretty Good Measurement developed in the previous section. Regarding the reduction using Unambiguous State Discrimination, the analysis is fairly simple. One interesting fact though is that we *do not* construct the state $\sum_{\boldsymbol{c},\boldsymbol{e}} f(\boldsymbol{e})|\boldsymbol{c}+\boldsymbol{e}\rangle$ but another state related to a punctured code $\mathcal{C}_J$ which allows us to find small dual codewords. This circumvents many issues arising when one wants to go from solving QDP (aka S-|LWE⟩) to construct the state $\sum_{\boldsymbol{c},\boldsymbol{e}} f(\boldsymbol{e})|\boldsymbol{c}+\boldsymbol{e}\rangle$ (aka C-|LWE⟩).

6 Proofs

The proofs of this article are presented in the full version of this paper (`https://arxiv.org/pdf/2310.20651`), which we omit here due to space restrictions.

References

1 Miklós Ajtai. Generating hard instances of lattice problems (extended abstract). In Gary L. Miller, editor, *Proceedings of the Twenty-Eighth Annual ACM Symposium on the Theory of Computing, Philadelphia, Pennsylvania, USA, May 22-24, 1996*, pages 99–108. ACM, 1996. doi:10.1145/237814.237838.

2 Benny Applebaum, Naama Haramaty, Yuval Ishai, Eyal Kushilevitz, and Vinod Vaikuntanathan. Low-complexity cryptographic hash functions. In *ITCS*, volume 67 of *LIPIcs*, pages 7:1–7:31. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2017.

3 Sanjeev Arora and Rong Ge. New algorithms for learning in presence of errors. In Luca Aceto, Monika Henzinger, and Jiří Sgall, editors, *Automata, Languages and Programming*, volume 6755 of *LNCS*, pages 403–415. Springer Berlin Heidelberg, 2011. doi:10.1007/978-3-642-22006-7_34.

4 H. Barnum and E. Knill. Reversing quantum dynamics with near-optimal quantum and classical fidelity. *Journal of Mathematical Physics*, 43(5):2097–2106, April 2002. doi:10.1063/1.1459754.

5 Anja Becker, Antoine Joux, Alexander May, and Alexander Meurer. Decoding random binary linear codes in $2^{n/20}$: How $1+1=0$ improves information set decoding. In *Advances in Cryptology – EUROCRYPT 2012*, LNCS. Springer, 2012.

6 Leif Both and Alexander May. Optimizing BJMM with Nearest Neighbors: Full Decoding in $2^{2/21n}$ and McEliece Security. In *WCC Workshop on Coding and Cryptography*, September 2017. URL: http://wcc2017.suai.ru/Proceedings{_}WCC2017.zip.

7 Zvika Brakerski, Vadim Lyubashevsky, Vinod Vaikuntanathan, and Daniel Wichs. Worst-case hardness for LPN and cryptographic hashing via code smoothing. In Yuval Ishai and Vincent Rijmen, editors, *Advances in Cryptology - EUROCRYPT 2019 - 38th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Darmstadt, Germany, May 19-23, 2019, Proceedings, Part III*, volume 11478 of *LNCS*, pages 619–635. Springer, 2019. doi:10.1007/978-3-030-17659-4_21.

8 Kevin Carrier, Thomas Debris-Alazard, Charles Meyer-Hilfiger, and Jean-Pierre Tillich. Statistical decoding 2.0: Reducing decoding to LPN. In *Advances in Cryptology - ASIACRYPT 2022*, LNCS. Springer, 2022. URL: https://eprint.iacr.org/2022/1000.

9 André Chailloux, Thomas Debris-Alazard, and Simona Etinski. Classical and quantum algorithms for generic syndrome decoding problems and applications to the lee metric. In Jung Hee Cheon and Jean-Pierre Tillich, editors, *Post-Quantum Cryptography*, pages 44–62, Cham, 2021. Springer International Publishing.

10 Anthony Chefles and Stephen M. Barnett. Optimum unambiguous discrimination between linearly independent symmetric states. *Physics Letters A*, 250(4):223–229, 1998. doi:10.1016/S0375-9601(98)00827-5.

11 Yilei Chen, Zihan Hu, Qipeng Liu, Han Luo, and Yaxin Tu. On the hardness of $\mathsf{S}|\mathsf{LWE}\rangle$ with gaussian and other amplitudes, 2023. arXiv:2310.00644.

12 Yilei Chen, Qipeng Liu, and Mark Zhandry. Quantum algorithms for variants of average-case lattice problems via filtering. In Orr Dunkelman and Stefan Dziembowski, editors, *Advances in Cryptology - EUROCRYPT 2022 - 41st Annual International Conference on the Theory and Applications of Cryptographic Techniques, Trondheim, Norway, May 30 - June 3, 2022, Proceedings, Part III*, volume 13277 of *LNCS*, pages 372–401. Springer, 2022. doi:10.1007/978-3-031-07082-2_14.

13 Thomas Debris-Alazard. Code-based cryptography: Lecture notes, arxiv cs.cr 2304.03541, 2023.

14 Thomas Debris-Alazard, Maxime Remaud, and Jean-Pierre Tillich. Quantum reduction of finding short code vectors to the decoding problem. *IEEE Trans. Inform. Theory*, November 2023. in press, see also arXiv:2106.02747 (v2). doi:10.1109/TIT.2023.3327759.

15 Il'ya Dumer. Two decoding algorithms for linear codes. *Probl. Inf. Transm.*, 25(1):17–23, 1989.

16 Oded Goldreich, Shafi Goldwasser, and Dana Ron. Property testing and its connection to learning and approximation. *J. ACM*, 45(4):653–750, 1998. doi:10.1145/285055.285060.

17 Alex B. Grilo, Iordanis Kerenidis, and Timo Zijlstra. Learning-with-errors problem is easy with quantum samples. *Phys. Rev. A*, 99:032314, March 2019. doi:10.1103/PhysRevA.99.032314.

18 Ghazal Kachigar and Jean-Pierre Tillich. Quantum information set decoding algorithms. In *Post-Quantum Cryptography 2017*, volume 10346 of *LNCS*, pages 69–89, Utrecht, The Netherlands, June 2017. Springer.

19 Alexander May, Alexander Meurer, and Enrico Thomae. Decoding random linear codes in $O(2^{0.054n})$. In Dong Hoon Lee and Xiaoyun Wang, editors, *Advances in Cryptology - ASIACRYPT 2011*, volume 7073 of *LNCS*, pages 107–124. Springer, 2011.

20 Alexander May and Ilya Ozerov. On computing nearest neighbors with applications to decoding of binary linear codes. In E. Oswald and M. Fischlin, editors, *Advances in Cryptology - EUROCRYPT 2015*, volume 9056 of *LNCS*, pages 203–228. Springer, 2015.

21 Rafael Misoczki, Jean-Pierre Tillich, Nicolas Sendrier, and Paulo S. L. M. Barreto. MDPC-McEliece: New McEliece variants from moderate density parity-check codes, 2012. doi:10.1109/ISIT.2013.6620590.

22 Ashley Montanaro. On the distinguishability of random quantum states. *Communications in Mathematical Physics*, 273, July 2006. doi:10.1007/s00220-007-0221-7.

23 Eugene Prange. The use of information sets in decoding cyclic codes. *IRE Transactions on Information Theory*, 8(5):5–9, 1962. doi:10.1109/TIT.1962.1057777.

24 Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing, Baltimore, MD, USA, May 22-24, 2005*, pages 84–93, 2005. doi:10.1145/1060590.1060603.

25 Claude E. Shannon. A mathematical theory of communication. *Bell System Technical Journal*, 27(3):379–423, 1948. doi:10.1002/j.1538-7305.1948.tb01338.x.

26 Damien Stehlé, Ron Steinfeld, Keisuke Tanaka, and Keita Xagawa. Efficient public key encryption based on ideal lattices. In Mitsuru Matsui, editor, *Advances in Cryptology - ASIACRYPT 2009, 15th International Conference on the Theory and Application of Cryptology and Information Security, Tokyo, Japan, December 6-10, 2009. Proceedings*, volume 5912 of *LNCS*, pages 617–635. Springer, 2009. doi:10.1007/978-3-642-10366-7_36.

27 Jacques Stern. A method for finding codewords of small weight. In G. D. Cohen and J. Wolfmann, editors, *Coding Theory and Applications*, volume 388 of *LNCS*, pages 106–113. Springer, 1988.

28 Takahashi Yamakawa and Mark Zhandry. Verifiable quantum advantage without structure. In *63rd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2022, Denver, CO, USA, October 31 – November 3, 2022*, pages 69–74. IEEE, 2022. doi:10.1109/FOCS54457.2022.00014.

29 Yu Yu, Jiang Zhang, Jian Weng, Chun Guo, and Xiangxue Li. Collision resistant hashing from sub-exponential learning parity with noise. In *ASIACRYPT (2)*, volume 11922 of *Lecture Notes in Computer Science*, pages 3–24. Springer, 2019.