Advances in Mathematics of Communications
Volume X, No. 0X, 20xx, X–XX

doi:10.3934/amc.xx.xx.xx

ON THE COMPUTATIONAL HARDNESS OF THE CODE EQUIVALENCE PROBLEM IN CRYPTOGRAPHY

Alessandro Barenghi

Department of Electronics and Information
Politecnico di Milano

Jean-François Biasse

Department of Mathematics and Statistics
University of South Florida

Edoardo Persichetti

Department of Mathematical Sciences
Florida Atlantic University

Paolo Santini

Department of Information Engineering
Università Politecnica delle Marche

Abstract. Code equivalence is a well-known concept in coding theory. Recently, literature saw an increased interest in this notion, due to the introduction of protocols based on the hardness of finding the equivalence between two linear codes. In this paper, we analyze the security of code equivalence, with a special focus on the hardest instances, in the interest of cryptographic usage. Our work stems from a thorough review of existing literature, identifies the various types of solvers for the problem, and provides a precise complexity analysis, where previously absent. Furthermore, we are able to improve on the state of the art, providing more efficient algorithm variations, for which we include numerical simulation data. Our results include also a dedicated method for solving code equivalence with a quantum algorithm, as well as a refinement of quantum Information-Set Decoding (ISD) algorithms. In the end, the goal of this paper is to provide a complete, single point of access, which can be used as a tool for designing schemes that rely on the code equivalence problem.

1. Introduction

Code-based cryptography is one of the main areas of research aiming to provide security in a post-quantum scenario. The area is largely based on the well-known and understood Syndrome Decoding Problem (SDP), which leads to very good solutions for key establishment [1, 2, 21], but has shown to be far from optimal when designing signature schemes. With this in mind, a recent approach was presented in 2020, leveraging the code equivalence problem as the main hardness assumption; the result is the scheme known as LESS [11], a zero-knowledge protocol that can be converted to signature scheme via the Fiat-Shamir transformation [15].

2010 *Mathematics Subject Classification:* 11T71, 94A60.
Key words and phrases: Code Equivalence, Linear Codes.

The publication of LESS stirred the community into giving a deeper look at the hardness of code equivalence in practice; for example, shortly after the appearance of LESS, Beullens introduced an improved algorithm to solve the code equivalence problem [10] for certain specific instances. This had an immediate effect on LESS; new parameter choices were published in [6], alongside a variety of computational optimizations aimed at improving the protocol's efficiency. Such optimizations are possible, in the first place, as code equivalence can be seen as a particular type of cryptographic group action, thus drawing another line in the sand when compared with previous solutions from code-based cryptography. In the end, it is clear that the practical hardness of solving code equivalence is worthy of further investigation.

Our Contribution. In this work, we analyze and improve on the computational methods to solve instances of the Code Equivalence problem (CE) for which a solution exists, with a particular emphasis on the instances that are relevant to cryptography. This contribution is of fundamental interest, and it will have an important impact on future developments regarding schemes based on code equivalence, such as LESS. More specifically, we focus on the most efficient methods to solve the Linear Equivalence Problem (LEP) for codes over fields of cardinality $q \geq 5$, i.e. *the overwhelming majority* of the LEP instances. Note that CE efficiently reduces to LEP, which means that the techniques described in this paper apply to the resolution of almost all instances of CE. For the instances of LEP that we focus on, no efficient method applies, and the best known technique is due to Beullens [10].

- In Section 5 we study the costs of Leon's and Beullens' algorithms for LEP. We provide new arguments to measure the performance, that were not available previously, and that are essential to assess any further improvement.
- In Section 6, we describe a new technique for solving LEP, and we demonstrate that it is an improvement over the state of the art (Leon and Beullens)[1].

In Section 4, we also analyze Leon's and Beullens' algorithms for computationally hard instances of the Permutation Equivalence Problem (PEP). Instances of PEP are a narrow subset of the set of instances of LEP (almost all LEP instances are not a PEP instance), but there are efficient methods for generating hard PEP instances (i.e. when codes have large hull) which can be used in cryptography. Therefore, it is essential to have a precise analysis of the best computational methods for solving PEP in these special instances. Additionally, we describe in Section 7 quantum algorithms for solving the hard instances of PEP and LEP presented above.

- In Appendix A, we present an adaptation of the quantum Information Set Decoding (ISD) algorithm of Kachigar and Tillich [18] to q-ary codes. While this adaptation is straightforward, to the best of our knowledge, it does not appear in any prior works.
- In Section 7.1, we present a quantum adaptation Beullens' algorithm for PEP that offers an asymptotic speed-up over its classical counterpart.
- In Section 7.2, we present quantum adaptations of both Beullens' algorithm and of the new algorithm presented in Section 6 for LEP that offer an asymptotic speed-up over their classical counterparts.
- In Section 7.3, we present ideas for further improvements on these new quantum methods.

[1] When possible, we validate our analysis with numerical simulations; the employed Sage scripts are available at https://github.com/paolo-santini/LESS_project

As for the classical algorithms presented in this paper, the quantum methods we introduce offer the best performance on *almost all* instances of the Code Equivalence problem, including all instances that are used in cryptography. However, they are not competitive on special cases for which classical (quasi-)polynomial time algorithms exist (i.e. PEP on codes of small hull, or LEP on codes over fields of size $q = 2, 3, 4$).

Organization of the paper. We begin in Section 2 by recalling some background notions about coding theory, as well as quantum search algorithms. In Section 3, we describe the code equivalence problem and give a high level overview of its hardness, and what are the main approaches for solvers. The permutation equivalence case is treated first, in Section 4; we then describe solvers for linear equivalence separately, in Section 5, including our improved technique. Quantum solvers are discussed, for both cases together, in Section 7. Finally, we conclude in Section 8.

2. Background

We will use the conventions of Table 1 throughout the rest of the paper.

a	a scalar
A	a set
$\boldsymbol{a}$	a vector
$\boldsymbol{A}$	a matrix
a	a function or relation
$\mathcal{A}$	an algorithm
$\boldsymbol{I}_n$	the $n \times n$ identity matrix
$[a; b]$	the set of integers $\{a, a+1, \ldots, b\}$
$\mathbb{U}(A)$	the uniform distribution over the set A
$\xleftarrow{\$} A$	sampling uniformly at random from A

Table 1. Notation used in this document.

We denote with $\mathbb{Z}_q$ the ring of integers modulo q, and with $\mathbb{F}_q$ the finite field of order q, as is customary; obviously, we have $\mathbb{Z}_q = \mathbb{F}_q$ when q is a prime. The multiplicative group of $\mathbb{F}_q$ is indicated as $\mathbb{F}_q^*$. Given a vector $\boldsymbol{a} \in \mathbb{F}_q^n$, we denote by $\mathsf{Values}(\boldsymbol{a})$ the ordered multiset formed by its entries. We denote with $\mathsf{Aut}(\mathbb{F}_q)$ the group of automorphisms of the field $\mathbb{F}_q$. The sets of vectors and matrices with elements in $\mathbb{Z}_q$ (resp. $\mathbb{F}_q$) are denoted by $\mathbb{Z}_q^n$ and $\mathbb{Z}_q^{m \times n}$ (resp. $\mathbb{F}_q^n$ and $\mathbb{F}_q^{m \times n}$). We write $\mathsf{GL}_k(q)$ for the set of invertible $k \times k$ matrices with elements in $\mathbb{F}_q$, or simply GL_k when the finite field is implicit. Let S_n be the set of permutations over n elements. Given a vector $\boldsymbol{x} = (x_1, \ldots, x_n) \in \mathbb{F}_q^n$ and a permutation $\pi \in \mathsf{S}_n$, we write the action of π on $\boldsymbol{x}$ as $\pi(\boldsymbol{x}) = (x_{\pi(1)}, \ldots, x_{\pi(n)})$. Note that a permutation can equivalently be described as an $n \times n$ matrix with exactly one 1 per row and column. Analogously, for *linear isometries*, i.e. transformations $\tau = (\boldsymbol{v}; \pi) \in \mathbb{F}_q^{*n} \rtimes \mathsf{S}_n$, we write the action on a vector $\boldsymbol{x}$ as $\tau(\boldsymbol{x}) = (v_1 x_{\pi(1)}, \ldots, v_n x_{\pi(n)})$. Then, we can also describe these in matrix form as a product $\boldsymbol{Q} = \boldsymbol{D}\boldsymbol{P}$ where $\boldsymbol{P}$ is an $n \times n$ permutation matrix and $\boldsymbol{D} = \{d_{ij}\}$ is an $n \times n$ diagonal matrix with entries in $\mathbb{F}_q^*$. We denote with M_n the set of such matrices, usually known as *monomial* matrices.

2.1. Coding Theory. An $[n,k]$-*linear code* $\mathfrak{C}$ of length n and dimension $k \leq n$ over $\mathbb{F}_q$ is a k-dimensional vector subspace of $\mathbb{F}_q^n$. It can be represented by a full-rank matrix $\boldsymbol{G} \in \mathbb{F}_q^{k \times n}$ with rank k, called *generator matrix*, whose rows form a basis for the vector space, i.e. $\mathfrak{C} = \{\boldsymbol{uG}, \ \ \boldsymbol{u} \in \mathbb{F}_q^k\}$. Alternatively, a linear code can be represented as the kernel of a full-rank matrix $\boldsymbol{H} \in \mathbb{F}_q^{(n-k)\times n}$, known as *parity-check matrix*, i.e. $\mathfrak{C} = \{\boldsymbol{x} \in \mathbb{F}_q^n : \boldsymbol{Hx}^T = 0\}$. For both representations, there may exist a standard choice, called *systematic form*, which corresponds, respectively, to $\boldsymbol{G} = (\boldsymbol{I}_k \mid \boldsymbol{M})$ and $\boldsymbol{H} = (-\boldsymbol{M}^T \mid \boldsymbol{I}_{n-k})$. Generator (resp. parity-check) matrices in systematic form can be obtained very simply by calculating the row-reduced echelon formstarting from any other generator (resp. parity-check) matrix. We denote such a procedure by sf. The parity-check matrix is important also as it is a generator for the *dual code*, defined as the set of words that are orthogonal to the code, i.e. $\mathfrak{C}^\perp = \{\boldsymbol{y} \in \mathbb{F}_q^n \ : \ \forall \boldsymbol{x} \in \mathfrak{C}, \ \ \boldsymbol{x} \cdot \boldsymbol{y}^T = 0\}$. Codes that are contained in their dual, i.e. $\mathfrak{C} \subseteq \mathfrak{C}^\perp$, are called *self-orthogonal* or *weakly self-dual*, and codes that are equal to their dual, i.e. $\mathfrak{C} = \mathfrak{C}^\perp$, are called simply *self-dual*.

We now proceed by recalling some well known definitions and results, which we will frequently use in the rest of the paper.

Definition 2.1. Let $\mathfrak{C} \subseteq \mathbb{F}_q^n$ be a code with dimension k. We define the permutation automorphism group of $\mathfrak{C}$ as

$$\mathsf{Aut}_{\mathsf{S}_n}(\mathfrak{C}) = \{\pi \in \mathsf{S}_n \mid \pi(\mathfrak{C}) = \mathfrak{C}\}.$$

Analogously, we define the monomial automorphism group of $\mathfrak{C}$ as

$$\mathsf{Aut}_{\mathsf{M}_n}(\mathfrak{C}) = \{\mu \in \mathsf{M}_n \mid \mu(\mathfrak{C}) = \mathfrak{C}\}.$$

Note that, if $\pi \in \mathsf{Aut}_{\mathsf{S}_n}$, then for any $\boldsymbol{G}$ that generates $\mathfrak{C}$, there must exist $\boldsymbol{S} \in \mathsf{GL}_k$ such that $\boldsymbol{G} = \boldsymbol{S}\pi(\mathfrak{C})$. Clearly, analogous relation applies to the monomial automophism group.

Definition 2.2 (Code support)**.** For a linear code $\mathfrak{C} \subseteq \mathbb{F}_q^n$, we define the support $\mathsf{Supp}(\mathfrak{C}) \subset \{1, \dots, n\}$ as the set of indexes i for which there is at least one codeword $\boldsymbol{c} \in \mathfrak{C}$ such that $c_i \neq 0$.

We now introduce another concept which will be fundamental for the analysis we develop in this paper.

Definition 2.3. Let $\mathfrak{C} \subseteq \mathbb{F}_q^n$ be a linear code with dimension k. A k'-dimensional subcode $\mathfrak{C}'$ of $\mathfrak{C}$ is a k'-dimensional vector space that can be generated by k' codewords of $\mathfrak{C}$. The set of all k'-dimensional subcodes with support size w is indicated as $A_w^{(k')}(\mathfrak{C})$. We refer to such a set as the *k'-dimensional Hamming sphere with radius w.*

The concept of code support can be deemed as a direct generalization of the notion of support for a vector (i.e. the set of indexes pointing at non null coordinates). In particular, for a vector, the cardinality of its support is referred to as *Hamming weight*:

$$\mathsf{wt}(\boldsymbol{a}) : \mathbb{F}_q^n \mapsto \mathbb{N} \ \ := \ \ \mathsf{wt}(\boldsymbol{a}) = |\mathsf{Supp}(\boldsymbol{a})|.$$

Remark 1. For $k' = 1$, the set $A_w^{(k')}(\mathfrak{C})$ contains all the codewords that have Hamming weight w and are distinct, even when considering multiple scalars. To ease the notation, we will refer to such a set as $A_w(\mathfrak{C})$.

We now continue with some properties of linear codes, in terms of subcodes having a desired support size.

Lemma 1. Let $\mathfrak{C} \subseteq \mathbb{F}_q^n$ be a k-dimensional linear code. Then, the number of subcodes of $\mathfrak{C}$ with dimension $k' \leq k$ is given by

$$\begin{bmatrix} k \\ k' \end{bmatrix}_q = \frac{(q^k - 1)\dots(q^k - q^{k'-1})}{(q^{k'} - 1)\dots(q^{k'} - q^{k'-1})} = \prod_{i=0}^{k'-1} \frac{q^k - q^i}{q^{k'} - q^i}.$$

When a code is picked at random, it is safe to assume that the contained k'-dimensional subcodes are random as well, that is, uniformly distributed over the set of all possible k'-dimensional vector subspaces of $\mathbb{F}_q^n$. Starting from this consideration (which is a standard assumption in coding theory), we can count the number of subcodes having a certain support size.

Proposition 1. Let $\mathfrak{C} \subseteq \mathbb{F}_q^n$ be a random linear code with dimension k. Then, the average number of subcodes with dimension k' and support size w is bounded from above by

$$\frac{(q^{k'} - 1)^w \binom{n}{w}}{\prod_{i=0}^{k'-1} (q^{k'} - q^i)} \frac{\left[{k \atop k'} \right]_q}{\left[{n \atop k'} \right]_q}.$$

Proof. Let $J \subseteq \{1, \cdots, n\}$ of size w, and consider the codes with dimension k' and whose support is exactly J. We can upper bound the number of such codes by $\frac{(q^{k'}-1)^w}{\prod_{i=0}^{k'-1} q^{k'} - q^i}$. Indeed, $(q^{k'} - 1)^w$ counts the number of matrices that have no null column among those indexed by J, while all the other ones are null; we divide this number by $\prod_{i=0}^{k'-1} q^{k'} - q^i$ to take into account all possible bases. Note that this is an upper bound since not all the considered matrices will have full rank k'. Since we have $\binom{n}{w}$ choices for J, we obtain $\frac{\binom{n}{w}(q^{k'}-1)^w}{\prod_{i=0}^{k'-1}(q^{k'}-q^i)}$ as an upper bound for the number of k'-dimensional subcodes of $\mathbb{F}_q^n$ with support size w. We now assume that all of the $\left[{k \atop k'} \right]_q$ subcodes of $\mathfrak{C}$ with dimension k' are randomly and uniformly picked among the set of $\left[{n \atop k'} \right]_q$ subspaces of $\mathbb{F}_q^n$ with dimension k'. Then, the probability that a specific k'-dimensional subcode has support size w is $\frac{(q^{k'}-1)^w \binom{n}{w}}{\prod_{i=0}^{k'-1}(q^{k'}-q^i)} \frac{1}{\left[{n \atop k'} \right]_q}$. Multiplying the above probability by $\left[{k \atop k'} \right]_q$ (that is, the number of subcodes in $\mathfrak{C}$ with dimension k') we obtain the estimate. □

Remark 2. When $k' = 1$, the number of subcodes is equal to that of codewords with Hamming weight w (without counting scalar multiples). For simplicity, we will denote this quantity as N_w, and have

$$N_w = N_w^{(1)} = \binom{n}{w} (q-1)^{w-1} \frac{q^k - 1}{q^n - 1}.$$

Remark 3. When $k' = 2$, we can improve upon the upper bound of Proposition 1 and obtain the average number of subcodes with support size w. To do this, it is enough to subtract from $(q^2 - 1)^w$ (that is, the number of matrices with w non null columns) the number of matrices that generate a one-dimensional space. Notice that these matrices are such that both rows have weight w and are identical up to a scalar multiplication; hence, they can be counted as $(q-1)^w(q-1) = (q-1)^{w+1}$. Consequently, we can set $\frac{(q^2-1)^w - (q-1)^{w+1}}{(q^2-1)(q^2-1)}$ as the number of subcodes of $\mathbb{F}_q^n$ with dimension 2 and support J of size w. Plugging this estimate into the proof of the above Proposition, we can estimate the average number of support size w subcodes of a random code as

$$N_w^{(2)} = \binom{n}{w} \frac{(q^2-1)^w - (q-1)^{w+1}}{(q^2-1)(q^2-q)} \frac{\left[{k \atop 2}\right]_q}{\left[{n \atop 2}\right]_q}.$$

2.2. ISD ALGORITHMS. Information Set Decoding (ISD) is the best technique to produce low weight codewords in a given code. There is a vast literature on ISD algorithms, most of which apply to the binary case; an extensive review can be found for example in [4]. For the more general, q-ary case (which is of interest to us), the work of Peters [22] is usually considered the go-to reference. In this paper, we will denote as $C_{ISD}(q, n, k, w)$ the cost of finding a specific codeword with weight w, in a code with length n and dimension k, defined over $\mathbb{F}_q$. In other words, if $\boldsymbol{c}$ is a codeword with weight w, then $C_{ISD}(q, n, k, w)$ is the cost to have an ISD routine return exactly $\boldsymbol{c}$. To assess $C_{ISD}(q, n, k, w)$, we rely on the analysis in [22]. Note that ISD is a randomized algorithm and, in case a code contains $N_w > 1$ codewords with weight w, then ISD will randomly return one of these codewords. In such a case, the complexity to find a codeword with weight w can be assessed as $\frac{C_{ISD}(q,n,k,w)}{N_w}$.

2.3. QUANTUM SEARCH ALGORITHMS. Our quantum algorithms for solving code equivalence problems rely on two different building blocks related to search problems, namely Grover's search algorithms and Quantum walks. Note that the former is a special case of the latter. Grover's algorithm assumes we know a set S and a function $f : S \to \{0, 1\}$ that is implemented by a quantum algorithm $\mathcal{O}_f$. Grover's algorithm returns (with constant probability) a marked element, that is $x \in S$ such that $f(x) = 1$. If we denote by $f^{-1}(\{1\}) = M \subseteq S$ and $\varepsilon = |M|/|S|$, the cost of Grover's algorithm is in

$$O\left(\frac{\text{Cost}(\mathcal{O}_f)}{\sqrt{\varepsilon}}\right).$$

Grover's search algorithm is generalized by the notion of random walk on a graph. We assume that a graph G is given by a set of vertices V and edges E, and we assume that we are looking for a marked element in $M = f^{-1}(\{1\})$ for some $f : V \to \{0, 1\}$. The general strategy of a random walk is to start from a vertex $x \in V$, check if $f(x) = 1$, and if not, then walk in the graph by sampling neighboors uniformly at random long enough to ensure the new vertex x' attained is distributed almost uniformly at random in V, then test if $f(x') = 1$. This is repeated until a marked element is found. In addition to running $\mathcal{O}_f$, there are two main steps in a quantum walk that contribute to the overall cost:

- Setup: sampling the first vector and initializing the data structure.
- Update: sampling a neighboor and updating the data structure (we need to update the current node and its neighbors).

Each of the aforementioned steps have a cost that depend on the data structure that is chosen to navigate the graph (note that depending on the model of computation chosen, memory-intensive data structures can penalize the cost). Moreover, the cost is impacted by the shape of the transition matrix $\boldsymbol{M}$. In the case of a d-regular graph (which is relevant to our problem), $\boldsymbol{M} = \frac{1}{d}\boldsymbol{A}$ where $\boldsymbol{A}$ is the adjacency matrix of the graph. The number of update steps required to reach a node almost uniformly distributed is $\tilde{O}\left(\frac{1}{\delta}\right)$ where δ is the *spectral gap* of $\boldsymbol{M}$, i.e. $\delta := 1 - \max_{i>1} |\lambda_i|$ where

$(\lambda_i)_{i>1}$ are the eigenvalues of $\boldsymbol{M}$ not equal to 1. The cost of a quantum walk is given by

$$\mathrm{Cost(Setup)} + \frac{1}{\sqrt{\varepsilon}}\left(\mathrm{Cost}(\mathcal{O}_f) + \frac{1}{\sqrt{\delta}}\mathrm{Cost(Update)}\right).$$

The search for solutions of the 4-sum problem reduces to a walk in the product of Johnson graphs of the V_i's. A Johnson graph $J(n,r)$ is an undirected graph whose vertices are the subsets of size r of a given set of size n. There is an edge between vertices S and S' if and only if $|S \cap S'| = r-1$ (i.e. they differ by only 1 element). The Johnson graph $J(n,r)$ has $\binom{n}{r}$ elements, is $r(n-r)$-regular and its spectral gap is

$$\delta = \frac{n}{r(n-r)}.$$

The product $J^m(n,r)$ of m copies of $J(n,r)$ is the graph whose vertices are of the form $(v_1,\dots,v_m)$ where each v_i is a vertex of $J(n,r)$, and there is an edge between $(v_1,\dots,v_m)$ and $(v'_1,\dots,v'_m)$ if and only if there is an edge between v_i and v'_i for some i, and $v_j = v'_j$ for all $j \neq i$. As recalled in [18], $J^m(n,r)$ has $\binom{n}{r}^m$ elements, is $mr(n-r)$-regular, and its spectral gap satisfies

$$\delta(J^m(n,r)) \geq \frac{1}{m}\delta(J(n,r)).$$

3. The Code Equivalence Problem

The concept of *equivalence* between two codes, in its most general formulation, is defined as follows.

Definition 3.1 (Code Equivalence)**.** We say that two linear codes $\mathfrak{C}_1$ and $\mathfrak{C}_2$ are *equivalent*, and write $\mathfrak{C}_1 \sim \mathfrak{C}_2$, if there exist a field automorphism $\alpha \in \mathsf{Aut}(\mathbb{F}_q)$ and a linear isometry $\tau = (\boldsymbol{v};\pi) \in \mathbb{F}_q^{*n} \rtimes \mathsf{S}_n$ that map $\mathfrak{C}_1$ into $\mathfrak{C}_2$, i.e. such that $\mathfrak{C}_2 = \tau(\alpha(\mathfrak{C}_1)) = \{\boldsymbol{y} \in \mathbb{F}_q^n \ : \boldsymbol{y} = \tau(\alpha(\boldsymbol{x})),\ \boldsymbol{x} \in \mathfrak{C}_1\}$.

Clearly, if $\mathfrak{C}_1$ and $\mathfrak{C}_2$ are two codes with generator matrices $\boldsymbol{G}_1$ and $\boldsymbol{G}_2$, respectively, it holds that

$$\mathfrak{C}_1 \sim \mathfrak{C}_2 \iff \exists(\boldsymbol{S};(\alpha,\boldsymbol{Q})) \in \mathsf{GL}_k \rtimes (\mathsf{Aut}(\mathbb{F}_q) \times \mathsf{M}_n) \text{ s.t. } \boldsymbol{G}' = \boldsymbol{S}\alpha(\boldsymbol{G}\boldsymbol{Q}).$$

The notion we just presented is usually known as *semilinear equivalence* and it is the most generic. If the field automorphism is the trivial one (i.e. $\alpha = id$), then the notion is simply known as *linear equivalence*. If, furthermore, the monomial matrix is a permutation (i.e. $\boldsymbol{Q} = \boldsymbol{D}\boldsymbol{P}$ with $\boldsymbol{D} = \boldsymbol{I}_n$), then the notion is known as *permutation equivalence*. Note that, in cryptographic applications (e.g. [11, 6]), the fields considered are always prime, and therefore the last two notions are the only ones of interest to us. Finally, we state the following computational[2] problem.

Problem 1 (Code Equivalence)**.** *Let $\boldsymbol{G}_1, \boldsymbol{G}_2 \in \mathbb{F}_q^{k\times n}$ be two generator matrices for two linearly equivalent codes $\mathfrak{C}_1$ and $\mathfrak{C}_2$. Find two matrices $\boldsymbol{S} \in \mathsf{GL}_k$ and $\boldsymbol{Q} \in \mathsf{M}_n$ such that $\boldsymbol{G}_2 = \boldsymbol{S}\boldsymbol{G}_1\boldsymbol{Q}$.*

We normally refer, respectively, to *permutation equivalence problem (PEP)* or *linear equivalence problem (LEP)*, according to the notion of code equivalence considered, or simply to the *code equivalence problem* where such distinction is not important.

[2] Note that this problem is traditionally formulated as a decisional problem in literature, yet for our purposes it is more natural to present here the search version.

3.1. High Level Hardness Overview. As proven in [23], the permutation equivalence problem is unlikely to be NP-complete, since this property would imply a collapse of the polynomial hierarchy. Yet, while the problem can be efficiently solved for some families of codes, there are many instances that, after almost 40 years of study, are still intractable. In the remainder of the paper, we analyze the best known solvers for the code equivalence problem. We first deal with the case of permutation equivalence, and report the complexity of all techniques to solve this problem. Then, we show how these techniques adapt to the case of linear equivalences.

We begin by recalling a trivial property of code equivalence.

Proposition 2. Let $\mathfrak{C}_1, \mathfrak{C}_2 \subseteq \mathbb{F}_q^n$ be two linear codes with dimension k, and let $\mathfrak{C}_1^\perp$, $\mathfrak{C}_2^\perp$ be their duals. Then

i. if $\pi \in \mathsf{S}_n$ is such that $\pi(\mathfrak{C}_1^\perp) = \mathfrak{C}_2^\perp$, then also $\pi(\mathfrak{C}_1) = \mathfrak{C}_2$;

ii. if $\tau \in \mathsf{M}_n$ is such that $\tau(\mathfrak{C}_1^\perp) = \mathfrak{C}_2^\perp$, then also $\tau'(\mathfrak{C}_1) = \mathfrak{C}_2$, where τ' is derived from τ by taking the inverses of the scaling factors.

The above proposition is crucial to understand the hardness of solving the code equivalence problem. Indeed, the problem can equivalently be solved by looking at the given codes, or at their duals. For the sake of simplicity, in this work, we will describe all the algorithms and procedures by considering solely the codes initially given; to derive the corresponding complexity for the attack on the duals, it is enough to replace k with $n-k$ in all the provided formulas.

To avoid studying vacuously hard instances (i.e., those represented by codes that are not equivalent), we will always consider the case in which at least a solution is guaranteed to exist. Namely, we consider that:

- the code $\mathfrak{C}_1$ is chosen at random;
- for PEP, we have $\mathfrak{C}_2 \overset{\$}{\leftarrow} \{\pi(\mathfrak{C}_1) \mid \pi \in \mathsf{S}_n\}$;
- for LEP, we have $\mathfrak{C}_2 \overset{\$}{\leftarrow} \{\tau(\mathfrak{C}_1) \mid \tau \in \mathsf{M}_n\}$.

Note that the number of solutions to PEP is equal to the size of the automorphism group. Indeed, if π solves PEP and σ is such that $\sigma(\mathfrak{C}_1) = \mathfrak{C}_1$, then we have another solution to PEP by combining π and σ. Clearly, the same considerations hold for LEP. To the best of our knowledge, the behaviour of the autormorphism groups of random codes under this perspective has never been formally studied. However, it is essentially folklore that these groups is trivial. Consequently, in our study we are going to make use of the following structural assumption.

Assumption 1. *We assume that the permutation and monomial automorphism groups of the considered codes are trivial.*

As a result of the above assumption, all the code equivalence instances we consider admit only one solution.

3.1.1. *The easy cases.* We begin our analysis by discussing algorithms that treat special cases, leading to very efficient solvers. The first such algorithm is the Support Splitting Algorithm (SSA), introduced by Sendrier [26]. This solver is based on the idea of *signature function*, i.e. a function $\mathcal{S}$ that fixes the action of the permutation on each position in the code. A signature function is said to be *fully discriminant* if it returns a different value in each position, and this allows to reveal the permutation linking the two codes. The signature function proposed by Sendrier in [26] is based on the *hull space* of a code, that is, the intersection between a code and its dual, for which the *weight enumerator* is computed. In particular, to create a dependence between the signature value and the code positions, one can *puncture* the code, i.e. remove coordinates from the codewords. Putting these considerations together, in [26, Section 5.2] Sendrier proposes to build a signature as

$$\mathcal{S}(\mathfrak{C}_i) := \left\{ \mathsf{Wef}\Big(\mathfrak{H}(\mathfrak{C}_{\setminus i}) \Big) \;,\; \mathsf{Wef}\Big(\mathfrak{H}(\mathfrak{C}^{\perp}_{\setminus i}) \Big) \right\},$$

where $\mathfrak{C}_{\setminus i}$ is the code obtained from $\mathfrak{C}$ punctured in position i, $\mathfrak{H}$ denotes the hull and Wef denotes the Weight Enumerator Function. The hull computation requires simple linear algebra, and comes with a cost of $O(n^3)$ operations in the finite field. To compute the weight enumerator of a code, one usually needs to enumerate all of its codewords: assuming that the hull has dimension h, we can use $O(nq^h)$ as an estimate for the cost of each Wef computation. On the other hand, heuristically, we observe that using $\ln(n)$ refinements is enough to obtain a fully discriminant signature. In the end, the complexity of SSA can be estimated as $O\left(n^3 + n^2 q^h \ln(n)\right)$. Thus, the hull dimension plays a crucial role in the analysis of the performance of SSA. For random codes, this dimension is with high probability equal to a small constant [27], de facto making SSA a polynomial-time solver for PEP. On the other hand, SSA is very inefficient for codes that have a large hull. This is, for instance, the case of (weakly) self-dual codes, for which SSA can be made arbitrarily hard by choosing codes with a sufficiently large dimension. SSA can be extended to solve the linear equivalence problem as well; however, in this case, the algorithm is less efficient. In fact, such an adaptation requires applying SSA to the *closure* of the code, i.e. the linear code defined as $\{\boldsymbol{c} \otimes \boldsymbol{a}, \;\; \boldsymbol{c} \in \mathfrak{C}\}$, where $\boldsymbol{a} = (a_1, \cdots, a_{q-1})$ is any ordering of the non-zero elements of $\mathbb{F}_q$. A fundamental point is that, for $q \geq 5$, the closure of a code is always weakly-self dual, and thus has a hull of maximum dimension, leading to exactly the hardest instances for SSA to solve. These results are corroborated by the analysis in [25].

Note that SSA trivially fails in the case of codes with an empty hull. In this case, however, another approach is possible. In 2019, Bardet et al. [5] proposed a new method to solve the permutation equivalence problem, which fully exploits the connection between the permutation equivalence problem and the notion of *graph isomorphism*. The core idea of [5] is to reduce code equivalence to an instance of the *Weighted Graph Isomorphism (WGI) problem*. This is done by building matrices of the form $\boldsymbol{A}_{\mathfrak{C}_i} = \boldsymbol{G}_i^{\top} \left(\boldsymbol{G}_i \boldsymbol{G}_i\right)^{-1} \boldsymbol{G}_i$ from the codes considered, and observing that $\boldsymbol{A}_{\mathfrak{C}_1} = \boldsymbol{P}^{\top} \boldsymbol{A}_{\mathfrak{C}_1} \boldsymbol{P}$ allows to recover the permutation $\boldsymbol{P}$ that connects the two codes. Indeed, $\boldsymbol{A}_{\mathfrak{C}_1}$ and $\boldsymbol{A}_{\mathfrak{C}_2}$ are interpreted as the adjacency matrices of two graphs, and hence can be given as input to some routine which solves the WGI problem. Given that, to compute $\boldsymbol{A}_{\mathfrak{C}_1}$ and $\boldsymbol{A}_{\mathfrak{C}_2}$, only $O(n^{2.373})$ operations in the finite field are required (this is essentially the cost of matrix inversion), we have that this approach gives a complexity of

$$O\Big(n^{2.373} C_{WGI}(n) \Big),$$

where $C_{WGI}(n)$ denotes the complexity of a solver for the weighted graph isomorphism problem. Note that this problem can be solved, for many classes of graphs, with very efficient algorithms. Furthermore, Babai's recent breakthrough paper [3] shows that the problem can be solved, in the worst case, with quasi-polynomial complexity. Hence, even in the worst case scenario, this solver runs in a time that is quasi-polynomial in the code length, on codes that have a trivial hull. For the more general case of codes with a non-trivial hull, the reduction from graph isomorphism works in a different way. In this case the complexity scales heavily with the dimension of the hull and thus the solver is, in practice, much less efficient; a proof of this fact can be found in [5, Theorem 10].

Finally, an algebraic approach was investigated in [25], where the author shows how it is possible to solve permutation equivalence by modeling it as a quadratic system. When the hull is trivial, it is possible to add several linear equations (through a technique called *block linearization*), which makes the system very easy to solve. However, in the general case of a non-trivial hull, the methods proposed by the author (using shortened codes or searching for the closest vector in the code) always end up in exponential complexity; for example, the latter scales proportionally to q^k. It follows that, as mentioned by the author himself, this approach can be deemed efficient only for the case of trivial hulls, once again.

To conclude this first section, we clarify the main takeaway to the reader. All the methods described above provide efficient solvers for very specific cases (small or trivial hulls); however, for codes with large hulls, these methods become quickly impractical. More to the point: when considering code equivalence in cryptography, it is easy to avoid these attacks. In fact, for the linear equivalence problem, it is enough to consider random codes defined over a large enough alphabet ($q \geq 5$), and then the value $q^h = q^k$ is already large enough for any realistic choice of code parameters. On the other hand, if one wants to use permutation equivalence, choosing a weakly-self dual code is sufficient to guarantee maximum hull dimension. All these considerations are already taken into account in the original LESS work, and constitute essentially just a set of "best practices", to be considered when designing a cryptosystem based on code equivalence. We now move on to summarizing algorithms that are relevant to the analysis of such systems.

3.1.2. *Solvers for hard instances.* There are other algorithms that are able to solve the hard instances described above, for which the previous solvers are ineffective. This is because the complexity of such algorithms does not depends on the size of the hull. Instead, the algorithms are based on a different observation, namely, that both permutation and monomial transformations preserve the Hamming weight distribution of the codewords. In particular, if two codes $\mathfrak{C}_1$ and $\mathfrak{C}_2$ are linked by some permutation or monomial transformation, say τ, then we have that for any subset of weight-w codewords $A_1 \subseteq \mathfrak{C}_1$, there must exist some subset of weight-w codewords $A_2 \subseteq \mathfrak{C}_2$ such that $\tau(A_1) = A_2$. Starting from this basic reasoning, the goal becomes that of finding sets of codewords that i) can efficiently been computed, and ii) have enough structure to allow for the reconstruction of τ.

Leon's algorithm [19], which dates as the first technique to solve code equivalence, chooses A_1 and A_2 as the set of all codewords having some low Hamming weight w. The choice of w is crucial to determine the algorithm effectiveness. On the one hand, in fact, if w is too low then A_1 and A_2 may have not enough structure (i.e., they contain very few codewords) so that reconstructing τ may not be possible. Yet, low-weight codewords can be found with ISD algorithms (see Appendix A), with a cost that is significantly smaller than that of enumerating the whole code. On the other hand, if w is too high, the number of codewords in A_1 and A_2 may become too high, so that determining the sets becomes too time-consuming.

Recently, Beullens [10] proposed an algorithm which is able, in some cases, to improve over Leon's algorithm. For the permutation equivalence case (i.e., when $\tau \in \mathsf{S}_n$), one observes that the multisets formed by the entries of the codewords are preserved as well. Hence, one can construct the sets A_1 and A_2 by considering pairs of codewords (one in A_1, one in A_2) having the same multisets of entries. To avoid too many collisions (which would make the algorithm perform worse than Leon's), one can consider only the codewords having some moderately low weight. As a key observation, Beullens has shown how it is not necessary to find all of these matching codewords (differently from what one does in Leon's algorithm).

For the remainder of this work, we will focus our analysis on algorithms of this second type, as they constitute the most efficient attack avenue for cryptographic schemes based on code equivalence.

4. Solvers for Hard Permutation Equivalence Instances

In this section we recall the algorithms for the permutation equivalence problem, whose complexity does not depend on the hull size, that we anticipated in the previous section.

Leon's Algorithm. Chronologically, the first method capable of solving the code equivalence problem is due to Leon [19], and is based on the following reasoning. Let $\mathfrak{C}_1$ and $\mathfrak{C}_2$ be two linear codes with length n and dimension k, and $\pi \in \mathsf{S}_n$ such that $\pi(\mathfrak{C}_1) = \mathfrak{C}_2$. Let X be a set of codewords picked from $\mathfrak{C}_1$. Then, there must exist a set Y formed by codewords of $\mathfrak{C}_2$ and such that $\pi(X) = Y$: among all the maps from X to Y, there must necessarily also be those mapping $\mathfrak{C}_1$ into $\mathfrak{C}_2$. In [19], Leon proposes an algorithm that constructs the ensemble of permutations between two sets, with a running time that is polynomial in the cardinality of the sets. Starting from the observation that permutations preserve the Hamming weight, Leon suggests to form X and Y using the codewords with a properly low weight w. Let $A_w(\mathfrak{C}_1)$ and $A_w(\mathfrak{C}_2)$ denote such sets, and $\mathsf{Mor}_{\mathsf{S}_n}\big(A_w(\mathfrak{C}_1), A_w(\mathfrak{C}_2)\big)$ be the set of all permutations $\pi \in \mathsf{S}_n$ such that $\pi\big(A_w(\mathfrak{C}_1)\big) = A_w(\mathfrak{C}_2)$. In a nutshell, Leon's algorithm operates as follows:

1. compute $A_w(\mathfrak{C}_1)$ and $A_w(\mathfrak{C}_2)$;
2. construct $\mathsf{Mor}_{\mathsf{S}_n}\big(A_w(\mathfrak{C}_1), A_w(\mathfrak{C}_2)\big)$;
3. check if there exists $\pi \in \mathsf{Mor}_{\mathsf{S}_n}\big(A_w(\mathfrak{C}_1), A_w(\mathfrak{C}_2)\big)$ such that $\pi(\mathfrak{C}_1) = \mathfrak{C}_2$.

As Leon proves in the original paper, the complexity of the second and third steps is polynomial in the cardinality of $A_w(\mathfrak{C}_1)$ and $A_w(\mathfrak{C}_2)$, which we estimate with N_w as in Proposition 1. This also allows us to properly choose the value of w. Indeed, N_w grows (exponentially) with w: when w is too high, N_w may become so large that the first and second steps of the algorithm become too time-consuming. On the other hand, if w is too low, then the sets $A_w(\mathfrak{C}_1)$ and $A_w(\mathfrak{C}_2)$ are rather small and do not possess enough structure, in the sense that there may exist a very large number of maps from $A_w(\mathfrak{C}_1)$ to $A_w(\mathfrak{C}_2)$.

Heuristically, optimal values of w are those that are slightly larger than the minimum distance of the codes (which can be estimated with the Gilbert-Varshamov distance). Indeed, this normally guarantees that the sets $A_w(\mathfrak{C}_1)$ and $A_w(\mathfrak{C}_2)$ are moderately small and, at the same time, contain a sufficient number of codewords. A lower bound on the complexity of Leon's algorithm can be estimated as follows.

Proposition 3 ([10])**.** Let $\mathfrak{C}_1 \subseteq \mathbb{F}_q^n$ be a random code with dimension k, $\pi \xleftarrow{\$} \mathsf{S}_n$ and $\mathfrak{C}_2 = \pi(\mathfrak{C}_1)$. The cost of Leon's algorithm, running with parameter $w \in \mathbb{N}$, $w \leq n$, can be estimated[3] as

$$O\big(\ln(N_w) C_{ISD}(q, n, k, w)\big).$$

For the sake of completeness, the proof of Proposition 3 is reported in Appendix B, where we additionally (as a new result) derive a theoretical bound on the required value for w. In practice, the attack is normally optimized when w is slightly larger than the minimum distance (say, by 1 or 2).

Beullens' Algorithm. In a recent work [10], Beullens introduced a novel approach to solve the code equivalence problem. The algorithm can be thought of as a refinement of Leon's algorithm, in which one tries to reduce the computational complexity by avoiding to compute the whole set of codewords with some fixed weight. The algorithm is based on the simple, but effective, intuition that permutations preserve also the multiset entries. Exploiting this observation, one can easily see how Leon's algorithm can be improved, by reducing the size of X and Y. In a nutshell, Beullens' algorithm works by first finding a subset of codewords with weight w from each code, and then searches for collisions among codewords having the same entries multiset. Each found collision is then used to piece-wise reconstruct the action of the permutation: if $\boldsymbol{x} \in \mathfrak{C}_1$ and $\boldsymbol{y} \in \mathfrak{C}_2$ have the same entries multiset and $x_i \neq y_j$, then we guess $\pi(i) \neq j$. When the number of collisions is sufficiently high, one has enough information to fully retrieve the permutation π. As done in [10], we can consider that the algorithm is successful whenever the number of collisions is approximately $2n\ln(n)$.

Note that, differently from Leon, Beullens' algorithm is probabilistic, since it fails in case i) bad collisions are found (i.e., codewords $\boldsymbol{x}$ and $\boldsymbol{y}$ that have the same entries multiset but $\boldsymbol{y} \neq \pi(\boldsymbol{x})$), and ii) the number of collisions is too low. The analysis of these cases and a precise cost estimate (which is missing in Beullens' original paper) are based on several technical aspects, which we detail in Appendix C. A compact and simple analysis of Beullens' algorithm is encapsulated in the following Proposition.

Proposition 4. The time complexity of Beullens' algorithm, running with parameters L and w such that i) $L = \sqrt{2N_w n\ln(n)}$ and ii) $(1-1/N_w)(q-1)L^2\binom{w+q-3}{w-1}^{-1} < 1$ is

$$O\left(\sqrt{\frac{n\ln(n)}{N_w}} C_{ISD}(q, n, k, w)\right).$$

Condition i) implies that we do not find all codewords with weight w, otherwise the algorithm would reduce to Leon. Condition ii) sets a lower bound on the number of codewords we need to find, in order to have enough information to run the permutation recovery algorithm. Finally, condition iii) expresses the fact that bad collisions do not happen.

[3] Here we use the same estimate derived in [10, Section 2.2], which corresponds to a lower bound for the actual complexity since the cost of steps 2 and 3 is neglected. In other words, the proposition takes into account only the cost of the codewords enumeration phase.

5. Solvers for Hard Linear Equivalence Instances

In this section, we recall the procedure of Beullens, which is a starting point for our new algorithm. We provide lower bounds on the complexities of Leon's and Beullens' algorithms to solve LEP, in order to make a comparison with our own estimate. These algorithms have several features which are similar to the ones we have already analyzed in the previous section; yet, using monomial transformations instead of permutations lead, in some cases, to radical differences. Analogously to what we have done for the permutation case, we will study LEP under the hypothesis that a solution always exists and, recalling Assumption 1, that the monomial isomorphism group is trivial. As a consequence, we have that the only solutions for the LEP instance represented by $(\mathfrak{C}_1, \mathfrak{C}_2)$, with $\mathfrak{C}_2 = \tau(\mathfrak{C}_1)$, are the monomial τ and its scalar multiples. Note that all of these transformations use the same permutation, and differ only for the scaling coefficients.

5.1. Leon's algorithm.

Leon's algorithm can be used to solve the linear equivalence problem, with an operating procedure that is essentially identical to the one we have already discussed in Section 4. The only difference is in the fact that, after the codewords enumeration, one searches for a monomial matrix instead of a permutation. When the value of w is properly chosen, this can be reconstructed in polynomial time, so that the bottleneck in the computational complexity is (again) in the codewords enumeration. Hence, also in this case, we can rely on Proposition 3 to have an estimate for the cost of the algorithm.

5.2. Beullens' algorithm.

In [10], Beullens proposed a second algorithm, to solve the linear equivalence problem. The algorithm principle is analogous to the PEP case, but some modifications are necessary, since monomial transformations do not preserve the multisets of codewords entries. To overcome this issue, Beullens first observes that if $\tau \in \mathsf{M}_n$ is such that $\tau(\mathfrak{C}_1) = \mathfrak{C}_2$, then for any subcode $\mathfrak{B}_1 \subseteq \mathfrak{C}_1$ there must exist a subcode $\mathfrak{B}_2 \subseteq \mathfrak{C}_2$ such that $\tau(\mathfrak{B}_1) = \mathfrak{B}_2$. Considering subcodes of small dimension and small support (we expect that very few such subcodes exist) instead of low weight codewords, we have that the same procedure as the one to solve permutation equivalence (plus some tweaks) can retrieve the secret monomial transformation. In particular, as in [10], we analyze the algorithm when two-dimensional subcodes are employed. Note that, to obtain an algorithm solving linear equivalence, we need the following three tweaks:

1. the codewords matching procedure shown in Algorithm 7 is replaced with an algorithm that produces colliding subcodes. To do this, we need to i) tweak ISD so that it returns subcodes with support size w, and ii) introduce the function $\mathsf{Lex}^{(2)}$ to take into account two-dimensional spaces. For an example of how such functions may be computed, we refer the reader to Appendix D while, for the sake of completeness, we report the full subcodes collisions procedure in Algorithm 1. Note that computing $\mathsf{Lex}^{(2)}$ has a cost of $O\left(n(q^2-1)(q^2-q)\right)$;
2. the list P produced by Algorithm 1 contains pairs $\{\boldsymbol{X}, \boldsymbol{Y}\} \in \mathbb{F}_q^{2\times n} \times \mathbb{F}_q^{2\times n}$ for which $\mathsf{Lex}^{(2)}(\boldsymbol{X}) = \mathsf{Lex}^{(2)}(\boldsymbol{Y})$;
3. in the reconstruction phase, one first finds the permutation, and then recovers the scaling factors. To have an efficient permutation recovery method, we can proceed in a way that is analogous to that of the permutation equivalence

case; again, for the sake of completeness, we have reported the procedure in Algorithm 2. Once the permutation has been recovered, the scaling factors can be found in many efficient ways. For instance, the permutation can be applied to a generator for $\mathfrak{C}_1$, obtaining $\boldsymbol{G}'$. Then, we choose a parity-check matrix for $\mathfrak{C}_2$, and aim to determine a non-singular diagonal matrix $\boldsymbol{D}$ such that $\boldsymbol{G}'\boldsymbol{D}\boldsymbol{H}_2^\top = \boldsymbol{0}$. This linear system has $k(n-k)$ equations for n unknowns, so that in general it is over constrained and can be easily solved. The n non-null entries of $\boldsymbol{D}$ are the unknown scaling coefficients $\boldsymbol{v}$, which are used to retrieve the desired monomial as $\pi \rtimes \boldsymbol{v}$.

Algorithm 1: Algorithm to find and match subcodes

```
Data: Number of subcodes $L \in \mathbb{N}$, support size $w \in \mathbb{N}$, ISD routine
Input: linear codes $\mathfrak{C}_1, \mathfrak{C}_2 \subseteq \mathbb{F}_q^n$ with dimension $k$
Output: list $P$ containing pairs $(\boldsymbol{X}, \boldsymbol{Y}) \in \mathbb{F}_q^{2\times n} \times \mathbb{F}_q^{2\times n}$, such that
        $\mathsf{Lex}^{(2)}(\boldsymbol{X}) = \mathsf{Lex}^{(2)}(\boldsymbol{Y})$

/* Produce a list $X$ of $L$ subcodes from $\mathfrak{C}_1$ with support size $w$ */
$X = \varnothing$;
while $|X| < L$ do
    Call ISD to find $\mathfrak{B} \subseteq \mathfrak{C}_1$ with support size $w$;
    $\boldsymbol{X} \leftarrow$ basis of $\mathfrak{B}$;
    $X \leftarrow X \cup \{\mathsf{SF}(\boldsymbol{X})\}$;

/* Produce a list $Y$ of $L$ subcodes from $\mathfrak{C}_2$ with support size $w$ */
$Y = \varnothing$;
while $|Y| < L$ do
    Call ISD to find $\mathfrak{B} \subseteq \mathfrak{C}_2$ with support size $w$;
    $\boldsymbol{Y} \leftarrow$ basis of $\mathfrak{B}$;
    $Y \leftarrow Y \cup \{\mathsf{SF}(\boldsymbol{Y})\}$;

/* Find collisions between the lists $X$ and $Y$ */
for $\{\boldsymbol{X}, \boldsymbol{Y}\} \in X \times Y$ do
    if $\mathsf{Lex}^{(2)}(\boldsymbol{X}) = \mathsf{Lex}^{(2)}(\boldsymbol{Y})$ then
        $P \leftarrow P \cup \{\boldsymbol{X}, \boldsymbol{Y}\}$;

return $P$;
```

We now proceed with the complexity analysis of the algorithm. We first argue that the complexity to find a specific 2-dimensional subcode with support size w is (essentially) the same as finding a specific codeword with weight w. One can indeed apply the same procedure of an ISD algorithm, with only minor tweaks so that the algorithm searches (and returns) a subcode. Namely, the algorithm in [10] can be seen as an adaptation of Lee & Brickell ISD, since it just consists in first applying the typical gaussian elimination and then checking whether couples of rows generate a subcode with support size w. The resulting time complexity is (essentially) the same as Lee & Brickell algorithm to find low weight codewords. For the rest of this work (and coherently with the codewords search version), we will denote by $C_{ISD}(q,n,k,w)$ the corresponding time complexity of finding a solution, in the regime in which a unique solution exists.

Algorithm 2: Fast permutation recovery, for the linear equivalence version of Beullens' algorithm.

```
Input: list $P$, containing $M$ pairs $\{\boldsymbol{X},\boldsymbol{Y}\} \in \mathbb{F}_q^{2\times n} \times \mathbb{F}_q^{2\times n}$ with support size
       $w$ and such that $\mathsf{Values}(\boldsymbol{X}) = \mathsf{Values}(\boldsymbol{Y})$
Output: permutation $\pi$, or report failure

$\boldsymbol{U} \leftarrow n \times n$ matrix made of all ones;
for $\{\boldsymbol{X},\boldsymbol{Y}\} \in P$ do
    for $i \in \{1,\cdots,n\}$ do
        $\boldsymbol{x}_i \leftarrow i$-th column of $\boldsymbol{X}$;
        for $j \in \{1,\cdots,n\}$ do
            $\boldsymbol{y}_j \leftarrow j$-th column of $\boldsymbol{Y}$;

            /* Filter $(i,j)$                                        */
            if $(\boldsymbol{x}_i == \boldsymbol{0}) \neq (\boldsymbol{y}_j == \boldsymbol{0})$ then
                $u_{i,j} = 0$;

/* Use $\boldsymbol{U}$ to reconstruct the permutation; if not possible, report failure */
if $\boldsymbol{U}$ is a permutation matrix then
    $\pi \leftarrow$ permutation described by $\boldsymbol{U}$;
    return $\pi$;
else
    report failure;
```

To estimate the number of two-dimensional subcodes with support size w a random code contains, on average, we use $N_w^{(2)}$, as in Remark 3. Taking all of this into account, we have that the cost of each ISD call can be optimistically assessed as $\frac{C_{ISD}(q,n,k,w)}{N_w^{(2)}}$. Then, we can assess the cost of Algorithm 1 as follows.

Proposition 5. Let $\mathfrak{C}_1 \subseteq \mathbb{F}_q^n$ be a random linear code with dimension k, and let $\mathfrak{C}_2 = \tau(\mathfrak{C}_1)$ with $\tau \overset{\$}{\leftarrow} \mathsf{M}_n$. Let P be the list obtained by running Algorithm 1 with parameters L and w, with $w \leq n-k+2$. The algorithm runs in time

$$O\left(L\left(\log_2(L) + (q^2-q)(q-1)\right) + M' + M'' + \frac{L}{N_w^{(2)}} C_{ISD}(q,n,k,w)\right),$$

and produces a list P with $M = M' + M''$ elements, where $M' = L^2/N_w^{(2)}$ is the average number of good collisions and $M'' \leq \frac{t_w^{(2)}(L^2-M')}{N_w^{(2)}}$ is that of bad collisions.

Proof. See Appendix F. □

Remark 4. We expect $\frac{t_w^{(2)}(L^2-M')}{N_w^{(2)}}$ to be a loose upper bound on the value of M'', especially when q is not high. This is due to the fact that $t_w^{(2)}$ is a rather loose upper bound on the number of equivalent subcodes with support size w that one code possesses.

Note that, with arguments similar to those of Proposition 17, we can estimate the probability with which Algorithm 2 succeeds in retrieving the correct permutation. Yet, to avoid computations that may be too complicated, we omit these details and skip to the more interesting case in which a much more simpler, slightly optimistic of the Algorithm is derived.

5.3. Heuristic analysis. For comparison purposes, it is interesting to provide a crude lower bound on the cost of Beullens' algorithm. Indeed, since we do not have a precise estimate of the number M'' of bad collisions, we need to make sure that heuristic assumptions, made to compare our LEP resolution algorithm described in Section 6 with Beullens' method, are to the advantage of the latter. So, we conservatively neglect M' and assume that bad collisions never happen. Furthermore, we bound from below the number of good collisions we need to reconstruct the permutation. To this end, we consider that we must filter $n(n-1)$ pairs of indexes, and each pair of subcodes gives information about $2w(n-w)$ pairs of indices, so that we need at least $\left\lceil \frac{n(n-1)}{2w(n-w)} \right\rceil$ pairs of subcodes. Since we have that the number of good collisions is heuristically given by $M' = L^2/N_w^{(2)}$, then we can set

$$L = \sqrt{N_w^{(2)} \left\lceil \frac{n(n-1)}{2w(n-w)} \right\rceil}.$$

With this in mind, we can greatly simplify the analysis of the algorithm as follows.

Proposition 6. The time complexity of Beullens' algorithm, running with parameters L and w such that i) $L < N_w^{(2)}$, ii) $L = \sqrt{N_w^{(2)} \left\lceil \frac{n(n-1)}{2w(n-w)} \right\rceil}$ and iii) $\frac{t_w^{(2)} L^2 (1-1/N_w^{(2)})}{N_w^{(2)}} < 1$, is bounded from below by

$$\Omega\left(\frac{L}{N_w^{(2)}} C_{ISD}(q,n,k,w) \right).$$

6. Improving Beullens' algorithm for LEP

In this section, we analyze a method to improve Beullens' approach to solve LEP. Namely, we propose a new algorithm to choose the initial two-dimensional subcodes from which the list P is built, which is based on first finding small weight codewords and then combining them to obtain colliding subcodes.

6.1. Finding subcodes more efficiently. Our idea consists in constructing two dimensional subcodes by first finding codewords with small Hamming weight, say w', and then considering only the subcodes which are generated by pairs of such codewords and have support size w. Before analyzing our algorithm, we briefly sketch the main intuition behind it. Remember that Beullens' algorithm aims to find pairs of subcodes with support size w and to produce a collision in the first lexicographic basis. Note that no additional condition is required, apart from the one on the support size. Heuristically, we expect any such subcodes to behave like a length-w random code plus $n-w$ coordinates that are always null. Consider a pair of matrices such as those in Figure 1a: to have a matching in the computation of Lex, the two matrices must lead to the same orange sub-matrix (which is expected to contain a number of columns rather close to w). If, instead, we consider subcodes generated by a pair of codewords with weight w' (and, still, with support size w),

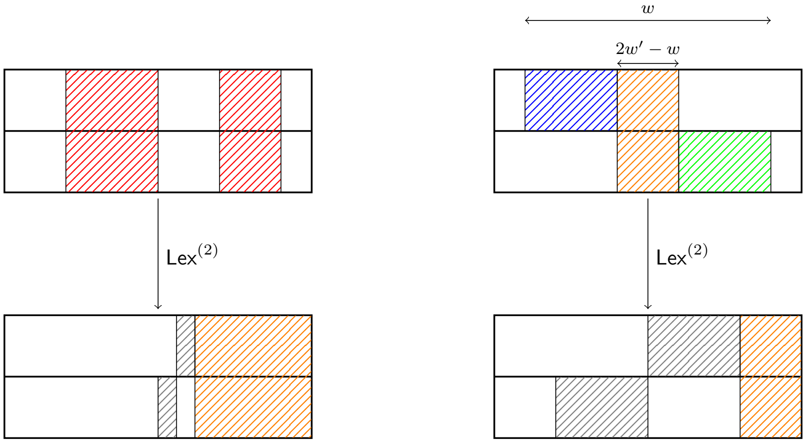

(A)

(B)

FIGURE 1. Computation of Lex for two-dimensional subcodes with support size w. In Figure (A), the subcode is a random one, while in Figure (B) it is generated by a pair of codewords with Hamming weight w'. The rectangles in gray color indicate the parts containing only ones, while the empty rectangles denote portions containing only zeros.

then the situation is depicted in Figure 1b. To have the same Lex value, the portions that must collide now contain $2w' - w$ columns. If we choose w' so that $2w'$ is only slightly larger than w, then we increase the probability to find collisions (because the number of relevant columns gets lower). Also, given how ISD operates, finding a sufficient number of low weight codewords should be easier than directly finding some subcodes with small support. In the end, this reasoning can be summarized as follows: we consider subcodes with a structure that i) allows to easily find them, and ii) increases the probability to find collisions.

To formalize the above intuition, we consider the following procedure to search for subcodes with small support:

1. use ISD to find L' codewords with weight w';
2. form $2 \times n$ matrices using all $\binom{L'}{2}$ pairs of codewords;
3. keep only the matrices which generate a code with support size w.

Clearly, the values of L' and w' have a strong impact on the complexity of this approach, which we derive in the sequel of this section.

We start our analysis with the following technical Lemma, which describes the distribution probability of the support size of a subcode that is generated by two codewords with known weights.

Lemma 2. *Let $\boldsymbol{a} \in \mathbb{F}_q^n$ with Hamming weight w_a. Let $\boldsymbol{b} \in \mathbb{F}_q^n$ be a random vector with Hamming weight w_b. Then, the code generated by $\boldsymbol{a}$, $\boldsymbol{b}$ (i.e., admitting the generator matrix whose rows are $\boldsymbol{a}$ and $\boldsymbol{b}$) has dimension 2 with probability*

$$\begin{cases} 1 & \text{if } w_b \neq w_a, \\ 1 - \frac{1}{\binom{n}{w_a}(q-1)^{w_a - 1}} & \text{if } w_b = w_a, \end{cases}$$

and support size $w \in [0; n]$ with probability

$$\zeta_{w_a,w_b}(w) = \begin{cases} 0 & \text{if } w < \max\{w_a; w_b\}, \\ 0 & \text{if } w > \min\{n; w_a + w_b\}, \\ \frac{\binom{w_a}{w_a+w_b-w}\binom{n-w_a}{w-w_a}}{\binom{n}{w_b}} & \text{otherwise.} \end{cases}$$

Proof. First, we consider the probability that the two chosen vectors do not generate a space with dimension 2. Note that this can happen only if $\boldsymbol{b} = v\boldsymbol{a}$ for some $v \in \mathbb{F}_q^*$. In such a case, we clearly have $w_a = w_b$. There are $q - 1$ distinct values for v (yielding to distinct vectors $v\boldsymbol{a}$), while the number of vectors with weight w_a is given by $\binom{n}{w_a}(q-1)^{w_a}$. Hence, the probability that $\boldsymbol{b}$ is one of them (that is, the probability that $\boldsymbol{a}$ and $\boldsymbol{b}$ generate a space with dimension 1) is given by

$$\frac{q-1}{\binom{n}{w_a}(q-1)^{w_a}} = \frac{1}{\binom{n}{w_a}(q-1)^{w_a-1}}.$$

We now derive the probability distribution for the support size of $\mathfrak{C}$, which we denote by w. Note that $w = w_a + w_b - |\mathsf{Supp}(\boldsymbol{a}) \cap \mathsf{Supp}(\boldsymbol{b})|$, from which we obtain $|\mathsf{Supp}(\boldsymbol{a}) \cap \mathsf{Supp}(\boldsymbol{b})| = w_a + w_b - w$. It is immediately seen that it must be $\max\{0; w_a + w_b - n\} \leq |\mathsf{Supp}(\boldsymbol{a}) \cap \mathsf{Supp}(\boldsymbol{b})| \leq \min\{w_a; w_b\}$, from which we find that the support size of $\mathfrak{C}$ is bounded as

$$\max\{w_a; w_b\} \leq w \leq \min\{w_a + w_b; n\}.$$

For all the admitted values, we have that we can have a support size w if and only if the set entries of $\boldsymbol{b}$ overlap with those of $\boldsymbol{a}$ in exactly $w_a + w_b - w$ positions. Since $\boldsymbol{b}$ is random, this happens with probability

$$\frac{\binom{w_a}{w_a+w_b-w}\binom{n-w_a}{w-w_a}}{\binom{n}{w_b}}.$$

□

Starting from a list of L' codewords with weight w', the number of subcodes with support size w that we can form with pairs of such codewords can be estimated as

$$\binom{L'}{2}\zeta_{w',w'}(w) \approx \frac{L'^2\zeta_{w',w'}(w)}{2}.$$

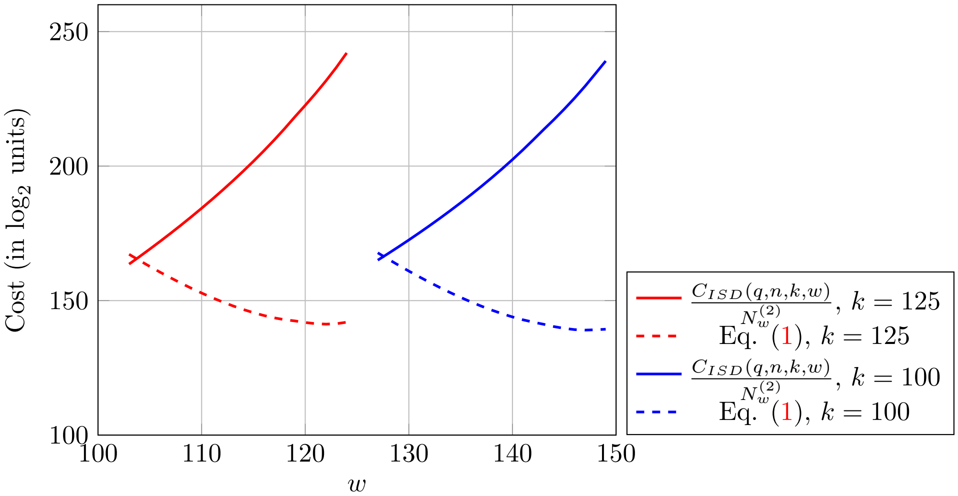

Figure 2. Cost of finding a subcode with support size w. All the considered codes have $q = 29$ and $n = 250$.

Setting $\frac{L'^2\zeta_{w',w'}(w)}{2} \approx 1$, from which $L' \approx \sqrt{\frac{1}{\zeta_{w',w'}(w)}}$, we have that on average the considered approach produces one subcode. Hence, considering the number of ISD calls we need, in order to produce L' distinct codewords with weight w', we have that our proposed approach can find a subcode with support size w with a cost given by

$$\text{(1)} \qquad \frac{\ln\left(1 - \frac{1}{N_{w'}\sqrt{\zeta_{w',w'}(w)}}\right)}{N_{w'}\ln\left(1 - \frac{1}{N_{w'}}\right)} C_{ISD}(q,n,k,w').$$

Note that, by using ISD directly, we need to face a cost given by $\frac{C_{ISD}(q,n,k,w)}{N_w^{(2)}}$. To show that this approach is faster than using ISD to directly search for subcodes (as proposed in [10]), we report a comparison between the costs of these two approaches in Figure 2, where we have considered several values of w and, for our proposed approach, we have computed the value of w' which minimizes (1). In the next section we describe how this reasoning affects the complexity of Beullens' algorithm.

6.2. Improved LEP algorithm. We now analyze the application of our proposed approach to Beullens' algorithm. Technically, we propose to replace Algorithm 1 with Algorithm 3. Notice that the only difference with Algorithm 1 is in how the lists X and Y are constructed. According to the analysis we have performed in the previous section, we expect this approach to be faster when the values of w' and L' are properly chosen.

Notice that, as a little technical caveat, we have that the probability to have bad collisions gets modified, because we are considering subcodes having a particular structure. To this end, we consider the following Proposition.

Algorithm 3: Our algorithm to find and match subcodes

```
Data: Number of codewords $L' \in \mathbb{N}$, weight $w' \in \mathbb{N}$, support size $w \in \mathbb{N}$,
      ISD routine
Input: linear codes $\mathfrak{C}_1, \mathfrak{C}_2 \subseteq \mathbb{F}_q^n$ with dimension $k$
Output: list $P$ containing pairs $(\boldsymbol{X}, \boldsymbol{Y}) \in \mathbb{F}_q^{2\times n} \times \mathbb{F}_q^{2\times n}$, such that
        $\mathsf{Values}(\mathsf{Lex}(\boldsymbol{X})) = \mathsf{Values}(\mathsf{Lex}(\boldsymbol{Y}))$

/* Produce a list $X'$ of $L'$ codewords from $\mathfrak{C}_1$ with weight $w'$ */
$X' = \varnothing$;
while $|X| < L$ do
    Call ISD to find $\boldsymbol{x} \in \mathfrak{C}_1$ with weight $w'$;
    $X' \leftarrow X' \cup \{\mathsf{Lex}(\boldsymbol{x})\}$;
/* Use pairs of codewords to produce subcodes with support size $w$ */
$X \leftarrow \varnothing$;
for $\boldsymbol{a} \in X'$ do
    for $\boldsymbol{b} \in X' \setminus \{\boldsymbol{a}\}$ do
        $\boldsymbol{X} \leftarrow$ matrix with rows $(\boldsymbol{a}, \boldsymbol{b})$;
        if Support of $\boldsymbol{X}$ has size $w$ then
            $X \leftarrow X \cup (\mathsf{Lex}(\boldsymbol{X}))$;

/* Produce a list $Y'$ of $L'$ codewords from $\mathfrak{C}_2$ with weight $w'$ */
$Y' = \varnothing$;
while $|X| < L$ do
    Call ISD to find $\boldsymbol{y} \in \mathfrak{C}_2$ with weight $w'$;
    $Y' \leftarrow Y' \cup \{\mathsf{Lex}(\boldsymbol{y})\}$;
/* Use pairs of codewords to produce subcodes with support size $w$ */
$Y \leftarrow \varnothing$;
for $\boldsymbol{a} \in Y'$ do
    for $\boldsymbol{b} \in Y' \setminus \{\boldsymbol{a}\}$ do
        $\boldsymbol{Y} \leftarrow$ matrix with rows $(\boldsymbol{a}, \boldsymbol{b})$;
        if Support of $\boldsymbol{Y}$ has size $w$ then
            $Y \leftarrow Y \cup (\mathsf{Lex}(\boldsymbol{Y}))$;

/* Find collisions between the lists $X$ and $Y$ */
for $\{\boldsymbol{X}, \boldsymbol{Y}\} \in X \times Y$ do
    if $\mathsf{Values}(\mathsf{Lex}\boldsymbol{X}) = \mathsf{ValuesLex}(\boldsymbol{Y})$ then
        $P \leftarrow P \cup \{\boldsymbol{X}, \boldsymbol{Y}\}$;

return $P$;
```

Proposition 7. Consider Algorithm 3, applied on two codes $\mathfrak{C}_1$ and $\mathfrak{C}_2$, where $\mathfrak{C}_1$ is random and $\mathfrak{C}_2 = \tau(\mathfrak{C}_1)$. Then, on average, P contains $M' = \frac{\zeta_{w',w'}(w)}{2}\left(\frac{L'^2}{N_{w'}}\right)^2$ good collisions and $M'' = p_{w'}(w)\frac{L'^4\zeta_{w',w'}(w)}{4}\left(\zeta_{w',w'}(w) - \frac{2}{N^2_{w'}}\right)$ bad collisions, where

$$p_{w'}(w) = \frac{\binom{n}{w-w'}\binom{n-(w-w')}{w-w'}\binom{n-2(w-w')}{2w'-w}(2w'-w)!(q-1)^{w-2w'+1}}{2\binom{n}{w'}\binom{n-w'}{w-w'}\binom{w'}{2w'-w}}.$$

Proof. We first derive the average number of good collisions. We consider that the number of codewords in X' which are mapped into codewords in Y' (through τ) can be estimated as $\widetilde{M} = \frac{L'^2}{N_{w'}}$. Indeed, for any codeword in X, we have only codeword (among all the $N_{w'}$ ones in $\mathfrak{C}_2$) which is its image through τ. Using any pair of such codewords to construct subcodes, we obtain good collisions. Considering that any of such subcodes will have the desired support size with probability $\zeta_{w',w'}(w)$, we can estimate the number of good collisions as

$$M' = \binom{\widetilde{M}}{2}\zeta_{w',w'}(w) \approx \frac{\widetilde{M}^2}{2}\zeta_{w',w'}(w) = \frac{\zeta_{w',w'}(w)}{2}\left(\frac{L'^2}{N_{w'}}\right)^2.$$

We now comment about the number of bad collisions. For any code, we dispose, on average, of $\binom{L'}{2}\zeta_{w',w'}(w) \approx \frac{L'^2}{2}\zeta_{w',w'}(w)$ subcodes with support size w. Hence, we have a total of $\left(\binom{L'}{2}\zeta_{w',w'}(w)\right)^2 \approx \left(\frac{L'^2}{2}\zeta_{w',w'}(w)\right)^2$ subcode pairs (one from X, one from Y): since M' of these pairs are good collisions, the number of pairs which may arise in bad collisions is $\left(\frac{L'^2}{2}\zeta_{w',w'}(w)\right)^2 - M' = \frac{L'^4\zeta_{w',w'}(w)}{4}\left(\zeta_{w',w'}(w) - \frac{2}{N_{w'}^2}\right)$. If each of these pairs is a bad collision with probability $p_{w'}(w)$, then we can estimate the number of bad collisions as

$$p_{w'}(w)\frac{L'^4\zeta_{w',w'}(w)}{4}\left(\zeta_{w',w'}(w) - \frac{2}{N_{w'}^2}\right).$$

To conclude the proof, we need to estimate $p_{w'}(w)$. Any subcode in X is generated by a $2 \times n$ matrix in which i) the rows have weight w', and ii) overlap in $x = 2w' - w$ positions (since the support has size w). The number of matrices with these properties is obtained as

$$U_{w'}(w) = \binom{n}{w'}\binom{n-w'}{w'-x}\binom{w'}{x}(q-1)^{2w'} = \binom{n}{w'}\binom{n-w'}{w-w'}\binom{w'}{2w'-w}(q-1)^{2w'}.$$

Indeed, the term $\binom{n}{w'}\binom{n-w'}{w'-x}\binom{w'}{x}$ counts all the possible supports for the rows of the generator matrix, while the term $(q-1)^{2w'}$ is due to the fact that, for each row, there are $(q-1)^{w'}$ choices for the set coefficients. We divide $U_{w'}(w)$ by $2(q-1)(q-1)$ to avoid multiple counting of the same matrix (since $(q-1)(q-1)$ is the number of matrices we can obtain by scaling each row, and the factor 2 is due to row swapping). Finally, we multiply $\frac{U_{w'}(w)}{2(q-1)^2}$ by $\frac{\begin{bmatrix}k\\2\end{bmatrix}_q}{\begin{bmatrix}n\\2\end{bmatrix}_q}$ to consider the probability that a matrix generates indeed a subcode of $\mathfrak{C}_1$. Now, we need to consider the number of subcodes of $\mathfrak{C}_1$ we can obtain by applying a monomial transformation to one of the generator matrices in X. This quantity can be set as

$$\begin{aligned}\tilde{t}_w^{(2)} &= \binom{n}{w'-x}\binom{n-w'+x}{w'-x}(q-1)^{2(w'-x)}\binom{n-2(w'-x)}{x}x!(q-1)^{x-1}\frac{\begin{bmatrix}k\\2\end{bmatrix}_q}{\begin{bmatrix}n\\2\end{bmatrix}_q}\\ &= \binom{n}{w-w'}\binom{n-(w-w')}{w-w'}\binom{n-2(w-w')}{2w'-w}(2w'-w)!(q-1)^{w-1}\frac{\begin{bmatrix}k\\2\end{bmatrix}_q}{\begin{bmatrix}n\\2\end{bmatrix}_q}\end{aligned}$$

Indeed, for $\boldsymbol{V} \in X$, we consider all possible matrices $\boldsymbol{V}'$ which can be obtained as $\boldsymbol{V}' = \sigma(\boldsymbol{V})$, where $\sigma \in \mathsf{M}_n$. In each $\boldsymbol{V}'$ we have $w' - x$ columns in which the entry in the first row is set and the one in the second row is null; also, we must have $w' - x$ other columns in which the entry in the first row is null, and the one in the second row is set. The number of such columns is counted as $\binom{n}{w'-x}\binom{n-w'+x}{w'-x}(q-1)^{2*(w'-x)}$.

We then consider the number of monomial transformations of the columns containing two set entries, which is equal to x: this number cannot be larger than $\binom{n-2(w'-x)}{x}x!(q-1)^x$. Indeed, it may happen that distinct transformations produce the same matrix, but we neglect such a possibility to simplify the analysis. Finally, we multiply again by the probability that the subcode generated by such a matrix is contained in $\mathfrak{C}_1$, and divide by $(q-1)$, to avoid multiple counting of matrices that generate the same subcode. Given the above reasoning, we can set

$$\begin{aligned} p_{w'}(w) = \frac{\tilde{t}_w^{(2)}}{\frac{U_{w'}(w)}{(q-1)}\frac{\begin{bmatrix}k\\2\end{bmatrix}_q}{\begin{bmatrix}n\\2\end{bmatrix}_q}} &= \frac{\binom{n}{w-w'}\binom{n-(w-w')}{w-w'}\binom{n-2(w-w')}{2w'-w}(2w'-w)!(q-1)^{w-1}}{2\binom{n}{w'}\binom{n-w'}{w-w'}\binom{w'}{2w'-w}(q-1)^{2w'-2}} \\ &= \frac{\binom{n}{w-w'}\binom{n-(w-w')}{w-w'}\binom{n-2(w-w')}{2w'-w}(2w'-w)!(q-1)^{w-2w'+1}}{2\binom{n}{w'}\binom{n-w'}{w-w'}\binom{w'}{2w'-w}}. \end{aligned}$$

□

We are now ready to evaluate the complexity of our new LEP algorithm.

Proposition 8. The time complexity of our LEP algorithm, using lists of L' codewords of weight w' such that i) $L' = \sqrt[4]{\frac{4N_{w'}^2}{\zeta_{w',w'}(w)}n\ln(n)}$ and ii) $M'' < 1$, is in

$$O\left(\frac{\ln\left(1 - L'/N_{w'}\right)}{N_{w'}\ln\left(1 - 1/N_{w'}\right)}C_{ISD}(q,n,k,w')\right).$$

Conditions i) and ii) sets an estimate on the number of good collisions we need; notice that condition ii) is obtained by setting $M' = 2n\ln(n)$. Finally, condition iii) guarantees that, with high probability, bad collisions do not happen. For large inputs, we have that $L'/N_{w'} = o(1)$, and therefore, the first order approximation of the cost of our LEP algorithm simplifies as

$$O\left(\frac{C_{ISD}(q,n,k,w')}{\sqrt{N_{w'}}}\sqrt[4]{\frac{n\log(n)}{\zeta_{w',w'}(w)}}\right).$$

6.3. Performance of our new LEP algorithm. In this section we comment about the effectiveness of our new approach. First, we present the results of numerical simulations, to validate the statement of Proposition 7.

(n,k,q)	(L',w',w)	Num subcodes		M'		M''	
		th.	emp.	th.	emp.	th.	emp.
$(40,20,7)$	$(10,12,19)$	7.55	7.30	1.30	0.43	3.20	2.57
	$(100,13,20)$	626.86	614.00	58.76	59.40	18558.74	15016.60
$(30,10,13)$	$(11,15,21)$	8.88	8.37	10.77	7.98	0.036	0.02
	$(40,16,24)$	207.30	208.32	24.48	24.78	25.42	26.22
$(30,10,19)$	$(25,16,24)$	79.73	82.06	76.20	78.56	0.22	0.44
	$(50,17,24)$	341.36	343.15	5.82	4.70	1.11	1.30

Table 2. Comparison between numerical results and theoretical estimates on the composition of the list P obtained with Algorithm 3. For each triplet (n,k,q), the empirical results have been averaged over 100 random codes.

To this end, for each parameter set, we have considered 100 different pairs of codes $\mathfrak{C}_1$ and $\mathfrak{C}_2$. Then, for each pair, we have simulated Algorithm 3; in Table 2 we compare the empirical values of M' and M'' (averaged over all the trials) with the theoretical ones, estimated though Proposition 7. As we can see from the table, the theoretical estimates match the empirical ones; this provides a validation of the heuristic we have employed to assess the performances of Algorithm 3.

In Figure 3 we compare the complexity arising from Proposition 8 with those of Leon's and Beullens' algorithms. For our algorithm we rely on the average complexity estimate resulting from Proposition 8, while for the other algorithms we have considered the lower bounds resulting from Propositions 3 and 6. Remember that for Leon's algorithm we are underestimating the weight value which is necessary to run the attack, so that, in practice, the actual complexity of the algorithm may be much larger. For Beullens' method, the lower bound comes from the fact that Proposition 6 is derived assuming bad collisions never happen (i.e. M'' is set to 0).

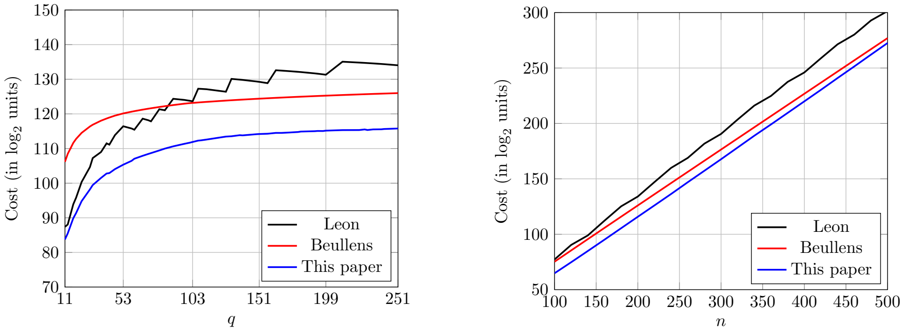

(A) $n = 200$, $k = 100$, several q

(B) $q = 251$, $k = \frac{1}{2}n$, several n

FIGURE 3. Comparison between several methods to solve LEP.

This comparison shows that our proposed algorithm performs better than the state-of-the-art solvers. Note that, in Figure 3, the cost of Beullens' method gets closer to ours for fixed q and $n \to \infty$. In this regime, the simplifications made in the cost analysis of Beullens' algorithm might underestimate the real cost by a significant margin. Indeed, the term M'' that we removed is dominated by $t_w^{(2)} \sim n!/(n-w)!$, which is exponential in n. To give more insight on how our algorithm operates, in Table 3 we have reported the optimal setting for the attack, for some of the instances we have considered in Figure 3.

q	L'	w'	w	Cost
11	13	56	97	$2^{83.76}$
53	183	62	111	$2^{105.34}$
103	$8.83 \cdot 10^7$	80	124	$2^{111.90}$
151	$5.57 \cdot 10^9$	83	128	$2^{114.19}$
199	$4.48 \cdot 10^{12}$	86	126	$2^{115.17}$
251	$3.39 \cdot 10^{14}$	88	127	$2^{115.78}$

TABLE 3. Optimal setting and resulting complexity for our algorithm to solve LEP. All the considered codes have $n = 200$, $k = 100$.

7. Quantum Solvers

In this section, we show how to use existing quantum claw finding methods to solve permutation and linear code equivalence problems. Indeed, both Leon and Beullen's algorithms can be easily recast as claw finding procedures that assume access to an oracle for quantum ISD. We present a quantum ISD algorithm based on Kachigar and Tillich's work [18] in Appendix A. To the best of our knowledge, this is the first analysis of the performance of quantum algorithms for the resolution of permutation and linear code equivalence. We also note that the literature was lacking a formal presentation of a q-ary variant of Kachigar and Tillich's quantum ISD method, which is the computational bottleneck of our quantum method for solving the code equivalence. Since quantum ISD methods upon which we rely were optimized for quantum time complexity (i.e. circuit depth), we chose this metric in our analysis. In particular, we assume that getting the i-th element in quantum memory of size N takes $O(\log(N))$ time, following the assumptions made in [18]. The real cost of quantum memory in cryptanalysis has been the topic of recent works, in particular from Jaques and Schanck [16]. It appears that focusing on depth only tends to underestimate the cost of quantum algorithms. On the other hand, for cryptographic purposes, security parameters derived from a cost analysis based on circuit depth tend to be more conservative.

7.1. Permutation Code Equivalence.

First, let us focus on the case of permutation equivalence. We collect sets $X \subseteq \mathfrak{C}_1$ and $Y \subseteq \mathfrak{C}_2$ of weight-w codewords. Let $m = \log_2(|X|) = \log_2(|Y|) = \log_2(L)$, and assume that the quantum version of Algorithm 6 for SDP returns a weight-w codeword that is distributed uniformly at random. To collect 2^m different weight-w codewords in $\mathfrak{C}_1$ (resp. $\mathfrak{C}_2$), we call the quantum ISD algorithm $O(m2^m)$ times (this is an instance of the coupon collector's problem). Then we store these 2^m elements in memory and we have two functions

$$f : \boldsymbol{x} \in \{0,1\}^m \mapsto \mathsf{Lex} \text{ of weight-}w \text{ codeword number } i \text{ given by Algorithm 6 on } \mathfrak{C}_1$$

$$g : \boldsymbol{x} \in \{0,1\}^m \mapsto \mathsf{Lex} \text{ of weight-}w \text{ codeword number } i \text{ given by Algorithm 6 on } \mathfrak{C}_2$$

As we have seen in the previous sections, we can derive the permutation between $\mathfrak{C}_1$ and $\mathfrak{C}_2$ from codewords of matching Lex in X and Y. Finding these matching Lex values corresponds to finding claws of f and g, i.e. pairs $\boldsymbol{x}, \boldsymbol{y} \in \{0,1\}^m$ such that $f(\boldsymbol{x}) = g(\boldsymbol{y})$. Tani's quantum algorithm [28] allows us to find such claws. Beullens showed that $p = \Theta(\log(n))$ claws are needed to solve the code equivalence problem. The procedure to reconstruct the permutation from these claws is described in Algorithm 8.

Theorem 1. *The time to find p unique pairs of vectors in $X \times Y$ with matching* Lex *value with Tani's claw finding algorithm is in*

$$O(m2^m C_{ISD}^{N_w}(q,n,k,w)) + \tilde{O}\left(C_{\mathsf{Lex}}(n,q,w)2^{m\frac{2p}{2p+1}}\right).$$

where $C_{\mathsf{Lex}}(n,q,w) = O(w(q-1))$ is the cost of computing the Lex *value of a vector, and $C_{ISD}^{N_w}(q,n,k,w)) = C_{ISD}^{1}(q,n,k,w))/\sqrt{N_w}$ is the cost of finding a codeword of weight w among $(q-1)N_w$ solutions with the quantum algorithm described in Appendix A.*

Proof. We construct the lists X and Y by calling the quantum ISD $O(m2^m)$ times in $\mathfrak{C}_1$ and $O(m2^m)$ times in $\mathfrak{C}_2$. Then from [28], we know that we can find p unique claws in $O(2^{m\frac{2p}{2p+1}})$ oracle calls. The cost of an oracle call is the calculation of the Lex value of the i-th precomputed codeword. □

As specified before, the cost of storing 2^m codewords in quantum memory is underestimated by focusing only on the time complexity. A naïve memory-less version could consist in the hard-coding of all the 2^m codewords of each set X in the circuit of the oracle for f, and the 2^m codewords of Y in the circuit of the oracle for g. This would result in an extra 2^m factor in the circuit depth of the claw-finding subroutine, i.e. a cost of $\tilde{O}\left(C_{\mathsf{Lex}}(n,q,w)2^{2m\frac{2p}{2p+1}}\right)$.

7.2. Linear Code Equivalence. Beullens [10, Sec. 4] also proposed a claw-finding procedure to solve the linear code equivalence. In this case, it is not sufficient to compare vectors of low weight from $\mathfrak{C}_1$ and $\mathfrak{C}_2$ to infer information on the hidden map from $\mathfrak{C}_1$ to $\mathfrak{C}_2$. However, this can be done if we consider 2-dimensional subspaces of $\mathfrak{C}_1$ and $\mathfrak{C}_2$ of support bounded by w. We denote

$$X_1(w) = \{V \subset \mathfrak{C}_1 \mid \dim(V) = 2 \text{ and } |\operatorname{Supp}(V)| \leq w\}$$
$$X_2(w) = \{V \subset \mathfrak{C}_2 \mid \dim(V) = 2 \text{ and } |\operatorname{Supp}(V)| \leq w\}$$

Testing whether $\mu(V) = W$ for $V \in X_1(w)$ and $W \in X_2(w)$ and μ the secret monomial permutation, is done by comparing $\operatorname{lex}(V)$ and $\operatorname{lex}(W)$. From pairs of matching subcodes, we retrieve the secret monomial permutation as in the approach designed for classical computers. Below, we propose quantum algorithms that achieve a speed-up over the approach of [10, Sec. 4], as well as over the improvements we proposed in Section 6.

7.2.1. *Beullens' approach.* The bulk of the work in this procedure is the search for elements of $X_1(w)$ and $X_2(w)$. Beullens proposes an adaptation of the general high level routine of Algorithm 6. Here we present a quantum adaptation of this method. First, assume we fix $V \in X_1(w)$, and let $\pi \in S_n$ be chosen at random. Then the probability that 2 indices of $\operatorname{Supp}(V)$ get mapped to $[1,k]$ while the $w-2$ remaining ones get mapped to $[k+1,n]$ is

$$P := \frac{\binom{n-k}{w-2}\binom{k}{2}}{\binom{n}{w}}.$$

For each good permutation π, we apply π to the generating matrix of the code and compute its row echelon form according to the first k columns (assuming linear independence of its restriction to these columns). Then one of the $\binom{k}{2}$ vector spaces spanned by two rows of the resulting matrix is $\pi(V)$.

Proposition 9. Using a Grover search, the cost $C_{\text{Q-LB}}^{N_w^{(2)}}(q,n,k,w)$ of Algorithm 4 on an instance with $N_w^{(2)}$ solutions is in

$$O\left(\frac{\log(n)}{\sqrt{\varepsilon}}k^2 n\log(q)\right)$$

where $\varepsilon := PN_w^{(2)}$.

Proof. Let $g : S_n \to \{0,1\}$ be the function that returns 1 if and only if a $V \in X_1(w)$ is found through the procedure of Steps 3 to 7. There are an average of $|X_1(w)| = N_w^{(2)}$ different V's to be found, and thus the probability that a given $\pi \in S_n$ yields some $V \in X_1(w)$ is $\varepsilon := P|X_1(w)|$. Thus, the cost of finding a V

Algorithm 4: 2-dimensional Quantum Lee-Brickell (Q-LB)

```
Input: $\boldsymbol{H}$ parity check matrix of $\mathfrak{C}$, $w$.
Output: $V \subseteq \mathfrak{C}$ of dimension 2 and support weight $w$.
1: for all $\pi \in \mathsf{S}_n$ do
2:   Compute row reduction $M$ of $\pi(\boldsymbol{H})$.
3:   for all $i, j \in [1, k]$, $i \neq j$ do
4:     $V \leftarrow \mathrm{Span}(M_i, M_j)$ ($V$ is spanned by rows of indices $i, j$).
5:     if $V$ has support of weight $w$ then
6:       return $V$
7:     end if
8:   end for
9: end for
```

is in $O\left(\frac{1}{\sqrt{\varepsilon}}\mathrm{Cost}(\mathcal{O}_g)\right)$ where $\mathcal{O}_g$ is the quantum circuit implementing g. To assess $\mathrm{Cost}(\mathcal{O}_g)$, we see that Step 2 costs $O(k^2)$ row operations (which cost $O(n\log(q))$ each), while a Grover search can perform the search of Steps 3 to 7 (over a search space of size $O(k^2)$) in $O(k)$ operations. □

To find the list of 2^m spaces $V \in X_1(w)$ and 2^m spaces $W \in X_2(w)$, we simply apply the above algorithm $O(m2^m)$ times sequentially. Then we use a claw finding method to recover $p \in \Theta(\log(n))$ pairs V, W with $\mathsf{Lex(V)} = \mathsf{Lex(W)}$, which costs $O\left(C_{\mathsf{Lex}}(n, q, w)2^{m\frac{2p}{2p+1}}\right)$ time with Tani's claw finding algorithm, as noted before. Then the functions f, g used in Tani's algorithm are:

$$f : \boldsymbol{i} \in \{0,1\}^m \mapsto \mathsf{Lex\ of\ } i\mathsf{th\ space\ of\ } X$$
$$g : \boldsymbol{j} \in \{0,1\}^m \mapsto \mathsf{Lex\ of\ } j\mathsf{th\ space\ of\ } Y$$

Proposition 10. **(Quantized version of Beullens' linear code equivalence)** Overall, the time complexity of finding p pairs of matching Lex value is in

$$O\left(m2^m C_{\text{Q-LB}}^{N_w^{(2)}}(q, n, k, w)\right) + \tilde{O}\left(C_{\mathsf{Lex}}(n, q, w)2^{m\frac{2p}{2p+1}}\right),$$

where $C_{\mathsf{Lex}}(n, q, w) = O(n(q^2 - q)(q^2 - 1))$ is the cost of computing the Lex value of a subcode given by two codes.

7.2.2. *Using pairs of codewords.* We also introduce a quantum variant of our new method for solving the linear code equivalence problem that consists in deriving $L' = 2^{m'}$ codewords of weight w', and then creating a list of $L = 2^m$ subcodes of dimension 2 with support size w. From a high level standpoint, the steps are summarized in Algorithm 5

Step 2 is dealt with as a claw finding problem within $\mathfrak{C}_1$ (resp. $\mathfrak{C}_2$). Let X' and Y' the sets of codewords from $\mathfrak{C}_1$ (resp. $\mathfrak{C}_2$) created in Step 1. We define the functions

$$f = g : \boldsymbol{i} \in \{0,1\}^m \mapsto i-\text{th codeword of } X'.$$

Then finding p' pairs of codewords of X' generating a support size w subcode is a p, q-subset finding problem solved by Tani's algorithm [28] where $p = q = p'$ and $(f(\boldsymbol{x}_1), \ldots, f(\boldsymbol{x}_{p'}), g(\boldsymbol{y}_1), \ldots, g(\boldsymbol{y}_{p'}) \in R$ if the subcode formed by all $\boldsymbol{x}_i, \boldsymbol{y}_i$ has support size w.

Algorithm 5: Linear code equivalence

```
Input: Codes $\mathfrak{C}_1$, $\mathfrak{C}_1$, weights $w, w'$, $p$.
Output: $p$ pairs of 2-dimensional subcodes $V_1 \subseteq \mathfrak{C}_1$, $V_2 \subseteq \mathfrak{C}_2$ with
    $\mathsf{Lex}^{(2)}(V_2) = \mathsf{Lex}^{(2)}(V_2)$
1: Create $2^{m'}$ weight-$w'$ codewords in $\mathfrak{C}_1$ and $\mathfrak{C}_2$.
2: In each code, find $p' = 2^m$ pairs of codewords that form a subcode of support
   size $w$.
3: Find $p$ claws for the functions $f$ and $g$ defined above.
```

Proposition 11. **(Quantized linear code equivalence from pairs of codewords)** The time complexity of Step 2 is $O\left(\log(q)n^2 2^{m'\frac{2p'}{2p'+1}}\right)$. Hence the total cost is

$$O(m'2^{m'}C_{ISD}^{N_{w'}}(q,n,k,w')) + O\left(\log(q)n^2 2^{m'\frac{2p'}{2p'+1}}\right) + O\left(C_{\mathsf{Lex}}(n,q,w)2^{m\frac{2p}{2p+1}}\right).$$

where $C_{\mathsf{Lex}}(n,q,w) = O(n(q^2-q)(q^2-1))$ is the cost of computing the Lex value of a subcode given by two codes.

The parameters w and w' need to be optimized, while p is still $O(\log(n))$. The other parameters follow from

- $2^{m'} = \left(\frac{2N_{w'}^2 p}{\zeta_{w',w'}(w)}\right)^{1/4}$.
- $2^m = \zeta_{w',w'}(w)\binom{2^{m'}}{2}$.
- $p' = 2^m$.

The quantum algorithms we have presented (for both permutation and linear code equivalence) focus on harvesting pairs of codewords (resp. subcodes) from $\mathfrak{C}_1$ and $\mathfrak{C}_2$ with matching $\mathsf{Lex}^{(2)}$ value. This is the bottleneck of the computation which is then concluded by using Algorithm 8 (resp. Algorithm 2) to retrieve the secret map between $\mathfrak{C}_1$ and $\mathfrak{C}_2$.

7.3. Further directions. The quantum algorithms for solving the code equivalence problem presented in this section are the first to have been described, to the best of our knowledge. While they do provide a speed-up with respect to the classical methods discussed in this paper for solving the hardest instances of the code equivalence problem (in particular: the linear code equivalence problem in codes over a field of cardinality $q \geq 5$), these approaches might not be optimal. In the following, we present a few directions for future improvements on our quantum methods.

7.3.1. *ISD in superposition.* The ISD algorithm returns a random solution to the decoding problem. In the cases of interest for the resolution of the code equivalence problem, there is a potentially large amount of solutions. The algorithms presented in this paper require a list of codewords of a given weight w. In the quantum algorithms of Section 7.1 and Section 7.2, such a list is collected by repeatedly querying a quantum ISD algorithm, and thus constructing a list of codewords (stored in classical memory). If the list is of small size (i.e. if the weight w is small), then most of the effort is spent on the ISD calls, and the cost of matching codewords is negligible. However, it is possible to consider a different trade-off where a longer list of codewords with large weight w is created, and then processed through a claw-finding procedure. In this case, it might be useful to incorporate the ISD routine to the claw-finding algorithm. To do this, we need to be able to implement in superposition

$$g : \boldsymbol{x} \in \{0,1\}^m \mapsto \mathsf{Lex\ of\ weight}\text{-}w \text{ codeword number } i \text{ given by Algorithm 6 on } \mathfrak{C}_2$$

Then, to derive a claw, we can use an approach based on the collision-finding algorithm of Brassard, Høyer and Tapp [13]. In the following, we refer to this approach as BHT. It is optimal in terms of *query complexity*, but requires access to the $O(2^m)$ elements stored in memory. The approach to find one claw can be summarized by the following steps:

1. Store $\{f(\boldsymbol{x}_1), \ldots, f(\boldsymbol{x}_{2^{m/3}})\}$ in memory for random $\boldsymbol{x}_1, \ldots, \boldsymbol{x}_{2^{m/3}} \in \{0,1\}^m$.
2. Use Grover's search algorithm to find $\boldsymbol{y} \in \{0,1\}^m$ such that

$$g(\boldsymbol{y}) \in \{f(\boldsymbol{x}_1), \ldots, f(\boldsymbol{x}_{2^{m/3}})\}.$$

Let $C^{N_w}_{\text{ISD-S}}(q, n, k, w)$ be the cost of ISD run in superposition on an instance with $(q-1)N_w$ solutions, and m such that $N_w = 2^m$ (i.e. X and Y are the sets containing all of the weight-w codewords up to multiplication by $\lambda \in \mathbb{F}_q^*$). The probability ε to draw $\boldsymbol{y} \in \{0,1\}^m$ such that $g(\boldsymbol{y}) \in \{f(\boldsymbol{x}_1), \ldots, f(\boldsymbol{x}_{2^{m/3}})\}$ is $\varepsilon \approx N_w/N_w^{1/3} = N_w^{2/3}$ because each element of X has only one match in Y (up to bad collisions which we neglect here to simplify the analysis). Then the run time of the search for one claw is in

$$\begin{aligned} & O\left(\frac{1}{\sqrt{\varepsilon}}\left(C^{N_w}_{\text{ISD-S}}(q,n,k,w) + C_{\mathsf{Lex}}(n,q,w)\right)\right) \\ & = O\left(N_w^{1/3}\left(C^{N_w}_{\text{ISD-S}}(q,n,k,w) + C_{\mathsf{Lex}}(n,q,w)\right)\right). \end{aligned}$$

The memory requirements of this procedure are $O(N_w^{1/3})$ vectors to store the list $\{f(\boldsymbol{x}_1), \ldots, f(\boldsymbol{x}_{2^{m/3}})\}$ in addition to the memory cost of ISD in superposition. We can see that the query complexity is lower than that of the approach described in Section 7.1 ($\tilde{O}(N_w^{1/3})$ instead of $\tilde{O}(N_w^{1/2})$ while the memory bottleneck is likely the ISD procedure. Currently, ISD returns a random solution to the decoding problem, but in order to use it as a subroutine in superposition within a claw-finding algorithm, we need it to consistently return solution number i on input i.

7.3.2. *Low memory routines.* The approaches presented in this section require a non trivial amount of memory, and make optimistic assumptions regarding the cost of memory access. Time-memory trade-offs (and memory-less solutions) are important to provide solutions to the code equivalence problem under conservative assumptions regarding the cost of memory. There has been a significant amount of effort devoted to find efficient collision finding methods with low memory requirement (see for example [14, 17]), and it is likely that the strategies developed in the context of collision for hash functions or claw finding for isogeny computation can be adapted to the resolution of the code equivalence problem.

Additionally, the quantum ISD methods presented in this paper also have a high memory cost due to the fact that they use quantum walks on the Johnson graph of a large set. These methods are clearly optimized for time (i.e. circuit depth), and better time-memory trade-offs can be achieved.

7.3.3. *A quantum SSA algorithm.* To conclude, we have a look at the quantum complexity of the Support Splitting Algorithm. Recall that this essentially, boils down to the computation of the Weight Enumerator Function (WEF) on the hull of the considered codes. The hull computation requires simple linear algebra, and comes with a cost of $O(n^3)$ operations in the finite field. The computation of the WEF of a code is the bottleneck of SSA. By definition, $\mathsf{Wef}(\mathfrak{C})$ is a bivariate polynomial given by

$$\mathsf{Wef}(\mathfrak{C})(x,y) = \sum_{w=0}^{n} (q-1)N_w x^w y^{n-w} \quad \text{where } (q-1)N_w = |\{c \in \mathfrak{C} \mid \mathrm{wt}(c) = w\}.$$

Hence, the computation of $\mathsf{Wef}(\mathfrak{C})$ reduces to the counting of elements of weight w for all $w \in [0,n]$ of a code. This can be seen as n instances of the quantum counting problem, which is defined as follows.

Definition 1. ***(Counting problem)*** *Given a set X and a function $f : X \to \{0,1\}$, find $f^{-1}(\{1\})$.*

The computation of $\mathsf{Wef}(\mathfrak{C})$ is thus directly rephrased as n counting problems defined by the functions $f_w : \boldsymbol{c} \in \mathfrak{C} \mapsto 1$ if and only if $\mathrm{wt}(\boldsymbol{c}) = w$. It is pretty clear that the cost of evaluating f_w is in $O(n)$. We denote by $\mathcal{O}_f$ the quantum circuit that reversibly evaluates f. The cost of finding an approximation of N_w can be obtained from a result of Brassard, Høyer and Tapp [12].

Proposition 12. (Cor. 4 of [12]) Let $f : X \to \{0,1\}$ be a function, $N = |X|$, and $t = f^{-1}(\{1\})$. Then there is an algorithm requiring an expected number of $\Theta\left(\sqrt{tN}\right)$ evaluations of f an estimate $\bar{t}$ such that $\bar{t} = t$ with probability at least $3/4$ using space linear in $\log(N)$.

Beals, Buhrman, Cleve, Mosca, and de Wolf [7] proved that any quantum algorithm capable of deciding with high probability whether not a function $F : \{0,\dots,N-1\} \to \{0,1\}$ is such that $\mid F^{-1}(\{1\}) \mid \leq t$, given some $0 < t < N/2$, must query F at least $\Omega(\sqrt{Nt})$ times, showing the optimality of Proposition 12. The issue to directly apply this result to the computation of $\mathsf{Wef}(\mathfrak{C})$ is the probability of success of quantum counting. Indeed, $\mathsf{Wef}(\mathfrak{C})$ is only successfully if the n instances of the counting problems return the correct N_w. In [12, Sec. 4], it is shown that an exact count can be reached with high probability through the use of $\Theta(t)$ additional quantum memory. Further work could be done to analyse which trade-off between success probability of quantum counting and quantum memory required would procure us an acceptable probability of success for the computation of $\mathsf{Wef}(\mathfrak{C})$. Below, we analyse the circuit depth that could be achieved if these precision issues were resolved.

Proposition 13. Let $\mathfrak{C}$ be an $[n,k]$ linear code over $\mathbb{F}_q$. There is a quantum algorithm for computing $\mathsf{Wef}(\mathfrak{C})$ in time

$$O\left(n^2 q^{k/2} \sqrt{(q-1)\max_{w \leq n} N_w}\right),$$

where $(q-1)N_w$ is the number of codewords of weight w.

Proof. We simply apply the counting algorithm for each possible weight (i.e. at most n times). For each call, the query complexity is $O(\sqrt{q^k(q-1)N_w})$ while the cost of the oracle is $O(n)$, hence the final result. □

Then, solving the code equivalence problem via SSA can be rephrased as finding claws for

$$f : i \in [1, n] \longmapsto \mathcal{S}(\mathfrak{C}_1, i)$$
$$g : j \in [1, n] \longmapsto \mathcal{S}(\mathfrak{C}_2, j)$$

This time, the space size is in $O(n)$, and therefore, this step can be performed in classical polynomial time in n rather than with expensive quantum methods.

8. Conclusions

The code equivalence problem is seeing an increasing presence in cryptographic literature. Since protocols based on code equivalence have the potential to be very efficient and lead to good solutions for code-based signature schemes (as well as other functionalities), it is important to properly assess the hardness of the problem in practical instances. In this paper, we provided a detailed analysis of the various approaches for solvers, for both permutation and linear code equivalence.

We have briefly explained why solvers that exploit particular properties, such as that of Bardet et al. [5] and Sendrier's support splitting algorithm [26], do not perform well in most instances of LEP, including the ones used in cryptography. In fact, both solvers are only truly efficient for the case of codes with trivial hulls, and it is thus easy to find hard instances. With regards to the latter, for example, it is worth mentioning that, in the linear case, SSA needs to be applied to the *closure* of the considered codes; however, for $q \geq 5$, the closure of a code is always weakly-self dual, and thus has a hull of maximum dimension k, leading to exactly the hardest instances for SSA to solve.

As a consequence of the above considerations, we gave an extensive treatment only to techniques that exploit the Hamming weight as an invariant, and utilize ISD as a subroutine for searching codewords. We have summarized and given a precise cost estimate of the two main algorithms of this type, Leon's [19] and Beullens' [10], that can be originally applied to the case of permutation equivalence. We have then shown how both can be adapted to the linear equivalence case, and produced a concrete technical analysis, which was lacking in the original works. Furthermore, we have presented an improved routine, that can considerably reduce the cost of Beullens' algorithm. We have given accurate complexity formulae, in all cases.

Finally, we have given consideration to the possibility of applying quantum techniques to solve code equivalence as well. To do so, we have described a dedicated technique, that uses a quantum version of ISD as a subroutine to find claws. The analysis of such a quantum ISD, as well as a description of a q-ary version of it, are both appearing for the first time, and therefore go to fill a significant gap in literature.

References

[1] M. R. Albrecht et al. "Classic McEliece: conservative code-based cryptography". In: (). URL: https://classic.mceliece.org/.

[2] N. Aragon et al. "BIKE: Bit Flipping Key Encapsulation". In: *NIST Post-Quantum Standardization, 3rd Round* (2021). URL: https://bikesuite.org/.

[3] L. Babai. "Graph Isomorphism in Quasipolynomial Time". In: *CoRR* abs/1512.03547 (2015). arXiv: 1512.03547. URL: http://arxiv.org/abs/1512.03547.

[4] M. Baldi et al. “A Finite Regime Analysis of Information Set Decoding Algorithms”. In: *Algorithms* 12.10 (2019). ISSN: 1999-4893. URL: https://www.mdpi.com/1999-4893/12/10/209.

[5] M. Bardet, A. Otmani, and M. Saeed-Taha. “Permutation Code Equivalence is Not Harder Than Graph Isomorphism When Hulls Are Trivial”. In: *IEEE ISIT 2019*. July 2019, pp. 2464–2468.

[6] A. Barenghi et al. “LESS-FM: Fine-tuning Signatures from the Code Equivalence Problem”. In: *International Conference on Post-Quantum Cryptography*. Springer. 2021, pp. 23–43.

[7] R. Beals et al. “Quantum Lower Bounds by Polynomials”. In: *39th Annual Symposium on Foundations of Computer Science, FOCS ’98, November 8-11, 1998, Palo Alto, California, USA*. IEEE Computer Society, 1998, pp. 352–361. DOI: 10.1109/SFCS.1998.743485. URL: https://doi.org/10.1109/SFCS.1998.743485.

[8] A. Becker et al. “Decoding Random Binary Linear Codes in $2^{n/20}$: How $1 + 1 = 0$ Improves Information Set Decoding”. In: *Advances in Cryptology - EUROCRYPT 2012 - 31st Annual International Conference on the Theory and Applications of Cryptographic Techniques, Cambridge, UK, April 15-19, 2012. Proceedings*. Ed. by D. Pointcheval and T. Johansson. Vol. 7237. Lecture Notes in Computer Science. Springer, 2012, pp. 520–536. DOI: 10.1007/978-3-642-29011-4_31. URL: https://doi.org/10.1007/978-3-642-29011-4%5C_31.

[9] D. J. Bernstein et al. “Quantum Algorithms for the Subset-Sum Problem”. In: *Post-Quantum Cryptography - 5th International Workshop, PQCrypto 2013, Limoges, France, June 4-7, 2013. Proceedings*. Ed. by P. Gaborit. Vol. 7932. Lecture Notes in Computer Science. Springer, 2013, pp. 16–33. DOI: 10.1007/978-3-642-38616-9_2. URL: https://doi.org/10.1007/978-3-642-38616-9%5C_2.

[10] W. Beullens. “Not Enough LESS: An Improved Algorithm for Solving Code Equivalence Problems over $\mathbb{F}_q$”. In: *International Conference on Selected Areas in Cryptography*. Springer. 2020, pp. 387–403.

[11] J.-F. Biasse et al. “LESS is More: Code-Based Signatures Without Syndromes”. In: *AFRICACRYPT*. Ed. by A. Nitaj and A. Youssef. Springer, 2020, pp. 45–65.

[12] G. Brassard, P. Høyer, and A. Tapp. “Quantum Counting”. In: *Automata, Languages and Programming, 25th International Colloquium, ICALP’98, Aalborg, Denmark, July 13-17, 1998, Proceedings*. Ed. by K. G. Larsen, S. Skyum, and G. Winskel. Vol. 1443. Lecture Notes in Computer Science. Springer, 1998, pp. 820–831. DOI: 10.1007/BFb0055105. URL: https://doi.org/10.1007/BFb0055105.

[13] G. Brassard, P. Høyer, and A. Tapp. “Quantum Cryptanalysis of Hash and Claw-Free Functions”. In: *LATIN ’98: Theoretical Informatics, Third Latin American Symposium, Campinas, Brazil, April, 20-24, 1998, Proceedings*. Ed. by C. L. Lucchesi and A. V. Moura. Vol. 1380. Lecture Notes in Computer Science. Springer, 1998, pp. 163–169. DOI: 10.1007/BFb0054319. URL: https://doi.org/10.1007/BFb0054319.

[14] A. Chailloux, M. Naya-Plasencia, and A. Schrottenloher. “An Efficient Quantum Collision Search Algorithm and Implications on Symmetric Cryptography”. In: *Advances in Cryptology - ASIACRYPT 2017 - 23rd International Conference on the Theory and Applications of Cryptology and Information Security, Hong Kong, China, December 3-7, 2017, Proceedings, Part II*. Ed. by T. Takagi and T. Peyrin. Vol. 10625. Lecture Notes in Computer Science. Springer, 2017, pp. 211–240. DOI: 10.1007/978-3-319-70697-9_8. URL: https://doi.org/10.1007/978-3-319-70697-9%5C_8.

[15] A. Fiat and A. Shamir. “How to prove yourself: Practical solutions to identification and signature problems”. In: *CRYPTO*. Springer. 1986, pp. 186–194.

[16] S. Jaques and J. M. Schanck. “Quantum Cryptanalysis in the RAM Model: Claw-Finding Attacks on SIKE”. In: *Advances in Cryptology - CRYPTO 2019 - 39th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2019, Proceedings, Part I*. Ed. by A. Boldyreva and D. Micciancio. Vol. 11692. Lecture Notes in Computer Science. Springer, 2019, pp. 32–61. DOI: 10.1007/978-3-030-26948-7_2. URL: https://doi.org/10.1007/978-3-030-26948-7%5C_2.

[17] S. Jaques and A. Schrottenloher. “Low-Gate Quantum Golden Collision Finding”. In: *Selected Areas in Cryptography - SAC 2020 - 27th International Conference, Halifax, NS, Canada (Virtual Event), October 21-23, 2020, Revised Selected Papers*. Ed. by O. Dunkelman, M. J. J. Jr., and C. O’Flynn. Vol. 12804. Lecture Notes in Computer Science. Springer, 2020, pp. 329–359. DOI: 10.1007/978-3-030-81652-0_13. URL: https://doi.org/10.1007/978-3-030-81652-0%5C_13.

[18] G. Kachigar and J.-P. Tillich. “Quantum Information Set Decoding Algorithms”. In: *Post-Quantum Cryptography - 8th International Workshop, PQCrypto 2017, Utrecht, The Netherlands, June 26-28, 2017, Proceedings*. Ed. by T. Lange and T. Takagi. Vol. 10346. Lecture Notes in Computer Science. Springer, 2017, pp. 69–89. DOI: 10.1007/978-3-319-59879-6_5. URL: https://doi.org/10.1007/978-3-319-59879-6%5C_5.

[19] J. Leon. “Computing automorphism groups of error-correcting codes”. In: *IEEE Transactions on Information Theory* 28.3 (May 1982), pp. 496–511.

[20] A. May, A. Meurer, and E. Thomae. “Decoding Random Linear Codes in $2^{0.054n}$”. In: *Advances in Cryptology - ASIACRYPT 2011 - 17th International Conference on the Theory and Application of Cryptology and Information Security, Seoul, South Korea, December 4-8, 2011. Proceedings*. Ed. by D. H. Lee and X. Wang. Vol. 7073. Lecture Notes in Computer Science. Springer, 2011, pp. 107–124. DOI: 10.1007/978-3-642-25385-0_6. URL: https://doi.org/10.1007/978-3-642-25385-0%5C_6.

[21] C. A. Melchor et al. “HQC: Hamming Quasi-Cyclic”. In: *NIST Post-Quantum Standardization, 3rd Round* (2021). URL: http://pqc-hqc.org/.

[22] C. Peters. “Information-set decoding for linear codes over $\mathbb{F}_q$”. In: *International Workshop on Post-Quantum Cryptography*. Springer. 2010, pp. 81–94.

[23] E. Petrank and R. M. Roth. “Is code equivalence easy to decide?” In: *IEEE Transactions on Information Theory* 43.5 (Sept. 1997), pp. 1602–1604.

[24] E. Prange. “The use of information sets in decoding cyclic codes”. In: *IRE Trans. Inf. Theory* 8.5 (Sept. 1962), pp. 5–9.

[25] M. A. Saeed. In: *PhD thesis* (2017).

[26] N. Sendrier. “The Support Splitting Algorithm”. In: *Information Theory, IEEE Transactions on* (Aug. 2000), pp. 1193–1203.

[27] N. Sendrier and P. Symbolique. “On the Dimension of the Hull”. In: *SIAM Journal on Discrete Mathematics* 10 (Nov. 1995). DOI: 10.1137/S0895480195294027.

[28] S. Tani. “An Improved Claw Finding Algorithm Using Quantum Walk”. In: *Mathematical Foundations of Computer Science 2007, 32nd International Symposium, MFCS 2007, Ceský Krumlov, Czech Republic, August 26-31, 2007, Proceedings*. Ed. by L. Kucera and A. Kucera. Vol. 4708. Lecture Notes in Computer Science. Springer, 2007, pp. 536–547. DOI: 10.1007/978-3-540-74456-6_48. URL: https://doi.org/10.1007/978-3-540-74456-6%5C_48.

Appendix A. Information-Set Decoding algorithms

We begin this section by reviewing the well-known syndrome decoding problem and its solvers, the information-set decoding algorithms.

Definition A.1 (Syndrome Decoding Problem (SDP))**.** Let $\mathfrak{C}$ be a code of length n and dimension k defined by a parity-check matrix $\boldsymbol{H} \in \mathbb{F}_q^{(n-k)\times n}$, and let $\boldsymbol{s} \in \mathbb{F}_q^{n-k}$, $w \leq n$, find $\boldsymbol{e} \in \mathbb{F}_q^n$ such that $\boldsymbol{He}^T = \boldsymbol{s}^T$ and $\boldsymbol{e}$ is of weight w.

Given that the matrix $\boldsymbol{H}$ has n columns but only $n-k$ rows, SDP is equivalent to solving an underdetermined linear system. As long as $w < n-k$, we can hope that the solution $\boldsymbol{e}$ has k zero coefficients. Let $\pi \in \mathsf{S}_n$ be the permutation of the columns of $\boldsymbol{H}$ that brings all these zeros to the last coefficients, then we have

$$\boldsymbol{He}^T = \begin{pmatrix}\boldsymbol{H}_1 & \boldsymbol{H}_2\end{pmatrix}\begin{pmatrix}\boldsymbol{e}_1^T \\ \boldsymbol{0}\end{pmatrix} = \boldsymbol{H}_1\boldsymbol{e}_1^T = \boldsymbol{s}^T$$

where $\boldsymbol{H}_1 \in \mathbb{F}_q^{(n-k)\times(n-k)}$, and $\boldsymbol{e}_1 \in \mathbb{F}_q^{n-k}$. We can then solve for $\boldsymbol{e}_1$ and recover $\boldsymbol{e}$. The original strategy to solve SDP due to Prange [24] (which is usually recognized as the first instance of an ISD algorithm) consists in sampling random $\pi \in \mathsf{S}_n$, and for each π, apply the permutation to $\boldsymbol{H}$ and attempt to solve the system of $n-k$ unknowns $\boldsymbol{H}_1\boldsymbol{e}_1^T = \boldsymbol{s}^T$. When an appropriate π is found (which is the difficult part), this yields a solution to SDP at little extra cost.

The above strategy puts the entire burden of the computation on the search for a good permutation π and almost none on the resolution of the subsequent linear system $\boldsymbol{H}_1\boldsymbol{e}_1^T = \boldsymbol{s}^T$. Instead, one can use a tradeoff where more permutations are considered (thus making the search for permutations more likely to succeed) at the cost of a more difficult system to solve. More specifically, we introduce two parameters l, p, and denote by $\mathscr{S}$ a set of $k+l$ indices, where instead of assuming that all entries of the solution are zero on these indices, we rather assume that only p of them are non-zero. Thus, the first construction corresponds to $l = p = 0$. The probability that exactly p non-zero coordinates of *a fixed* $\boldsymbol{e}$ belong to $\mathscr{S}$ while the remaining $w-p$ are outside of it is denoted by

$$P_{l,p} := \frac{\binom{k+l}{p}\binom{n-k-l}{w-p}}{\binom{n}{w}}.$$

Assuming a *good* permutation is found, one now needs to solve a new linear algebra problem. By "good", we mean that it moves the $k+l$ coordinates of indices in $\mathscr{S}$ to the left of the matrix. We apply this permutation π to the columns of $\boldsymbol{H}$ and assume that the restriction of $\pi(\boldsymbol{H})$ to its last $n-k-l$ columns is of full row rank. We perform a Gaussian elimination, which corresponds to the multiplication of $\pi(\boldsymbol{H})$ by an invertible matrix $\boldsymbol{U}$ on the left such that the resulting matrix has the shape:

$$\boldsymbol{U} \cdot \pi(\boldsymbol{H}) = \begin{pmatrix} \boldsymbol{H}' & \boldsymbol{0}_l \\ \boldsymbol{H}'' & \boldsymbol{I}_{n-k-l} \end{pmatrix}$$

We can write $\pi(\boldsymbol{e}) = (\boldsymbol{e}'||\boldsymbol{e}'')$ where $\boldsymbol{e}' \in \mathbb{F}_q^{l+k}$ has weight p and $\boldsymbol{e}'' \in \mathbb{F}_q^{n-k-l}$. Then, given a good permutation, we can reduce our problem to the resolution of the overdetermined system defined by $\boldsymbol{H}' \in \mathbb{F}_q^{l\times(k+l)}$. Indeed. we have

$$\boldsymbol{U} \cdot \boldsymbol{s}^T = \begin{pmatrix} \boldsymbol{s}'^T \\ \boldsymbol{s}''^T \end{pmatrix} = \boldsymbol{U} \cdot \boldsymbol{H}\boldsymbol{e}^T = \begin{pmatrix} \boldsymbol{H}'\boldsymbol{e}'^T \\ \boldsymbol{H}''\boldsymbol{e}'^T + \boldsymbol{e}''^T \end{pmatrix}$$

Once we find a solution $\boldsymbol{e}'$ to the system $\boldsymbol{H}'\boldsymbol{e}'^T = \boldsymbol{s}'^T$ (which is somewhat expensive), we can immediately derive $\boldsymbol{e}''^T$ as $\boldsymbol{s}''^T - \boldsymbol{H}''\boldsymbol{e}'^T$.

Algorithm 6: Generic ISD approach

```
Input: $\boldsymbol{H}$, $s$, $l$, $p$, $w$.
Output: $\boldsymbol{e}$ of weight $w$ such that $\boldsymbol{H}\boldsymbol{e}^T = \boldsymbol{s}^T$.
for all $\pi \in \mathsf{S}_n$ do
    Compute row reduction of $\pi(\boldsymbol{H})$: $\boldsymbol{U} \cdot \pi(\boldsymbol{H}) = \begin{pmatrix} \boldsymbol{H}' & \boldsymbol{0}_l \\ \boldsymbol{H}'' & \boldsymbol{I}_{n-k-l} \end{pmatrix}$
    Compute $\boldsymbol{s}', \boldsymbol{s}''$ defined by $\boldsymbol{U} \cdot \boldsymbol{s}^T = \begin{pmatrix} s'^T \\ s''^T \end{pmatrix}$
    Find $\boldsymbol{e}'$ of weight $p$ such that $\boldsymbol{H}'\boldsymbol{e}'^T = \boldsymbol{s}'^T$. (method to be determined)
    $\boldsymbol{e}''^T \leftarrow \boldsymbol{s}''^T - \boldsymbol{H}''e'^T$, $\boldsymbol{e} \leftarrow (e'||\boldsymbol{e}'')$.
    if $\boldsymbol{e}$ has weight $w$ then
        break
    end if
end for
return $\boldsymbol{e}$.
```

We are now concerned with solving the overdetermined linear system of Step 4 of Algorithm 6. This means that we are given $\boldsymbol{H}' \in \mathbb{F}_q^{l\times(k+l)}$, and $\boldsymbol{s}' \in \mathbb{F}_q^l$, and we are trying to find $\boldsymbol{e}' \in \mathbb{F}_q^{l+k}$ of weight p such that $\boldsymbol{H}'\boldsymbol{e}'^T = \boldsymbol{s}'^T$. This problem can be rephrased as a k-sum problem.

Definition A.2 (Generalized k-sum problem)**.** Consider an abelian group G, and arbitrary set E, a map $f : E \to G$, k subsets $V_0, V_1, \ldots, V_{k-1} \subseteq E$, a map $g : E^k \to \{0, 1\}$, and an element $s \in G$. Find $(v_0, \ldots, v_{k-1}) \in V_0 \times \ldots \times V_{k-1}$ such that

1. $f(v_0) + f(v_1) + \ldots + f(v_{k-1}) = s$.
2. $g(v_0, \ldots, v_{k-1}) = 0$.

The search for $\boldsymbol{e}'$ can be reduced to the 2-sum problem with the following parameters:

$$\begin{aligned}
G &= \mathbb{F}_q^l, E = \mathbb{F}_q^{l+k}, f(v) = \boldsymbol{H}'^T\boldsymbol{v}^T, s = \boldsymbol{s}'\\
V_0 &= \{(\boldsymbol{e}_0, \boldsymbol{0}_{(k+l)/2}) \in \mathbb{F}_q^{l+k}, \boldsymbol{e}_0 \in \mathbb{F}_q^{(l+k)/2}, \mathrm{wt}(\boldsymbol{e}_0) = p/2\}\\
V_1 &= \{(\boldsymbol{0}_{(k+l)/2}, \boldsymbol{e}_1) \in \mathbb{F}_q^{l+k}, \boldsymbol{e}_1 \in \mathbb{F}_q^{(l+k)/2}, \mathrm{wt}(\boldsymbol{e}_1) = p/2\}
\end{aligned}$$

and g defined as $g(v_0, v_1) = 0$ if and only if $\boldsymbol{e} = (\boldsymbol{e}'||\boldsymbol{e}'')$ is of weight w where $\boldsymbol{e}' = v_0 + v_1$ and $\boldsymbol{e}''^T = \boldsymbol{s}''^T - \boldsymbol{H}''\boldsymbol{e}'^T$. The cost of the resolution of the 2-sum problem is a tradeoff between the size of the V_i's and the size of G.

To further reduce the size of the sets involved and create more possibilities for tradeoffs, we can reduce the search for $\boldsymbol{e}' \in \mathbb{F}_q^{l+k}$ of weight p such that $\boldsymbol{H}'\boldsymbol{e}'^T = \boldsymbol{s}'^T$ to a 4-sum problem with the following parameters.

$$\begin{aligned}
G &= \mathbb{F}_q^l, E = \mathbb{F}_q^{l+k}, f(v) = \boldsymbol{H}'^T v^T, s = s'\\
V_{00} &= \{(\boldsymbol{e}_{00}, \boldsymbol{0}_{3(k+l)/4}) \in \mathbb{F}_q^{l+k}, \boldsymbol{e}_{00} \in \mathbb{F}_q^{(l+k)/4}, \mathrm{wt}(\boldsymbol{e}_{00}) = p/4\}\\
V_{01} &= \{(\boldsymbol{0}_{(k+l)/4}, \boldsymbol{e}_{01}, \boldsymbol{0}_{(k+l)/2}) \in \mathbb{F}_q^{l+k}, \boldsymbol{e}_{01} \in \mathbb{F}_q^{(l+k)/4}, \mathrm{wt}(\boldsymbol{e}_{01}) = p/4\}\\
V_{10} &= \{(\boldsymbol{0}_{(k+l)/2}, \boldsymbol{e}_{10}, \boldsymbol{0}_{(k+l)/4}) \in \mathbb{F}_q^{l+k}, \boldsymbol{e}_{10} \in \mathbb{F}_q^{(l+k)/4}, \mathrm{wt}(\boldsymbol{e}_{10}) = p/4\}\\
V_{11} &= \{(\boldsymbol{0}_{3(k+l)/4}, \boldsymbol{e}_{11}) \in \mathbb{F}_q^{l+k}, \boldsymbol{e}_{11} \in \mathbb{F}_q^{(l+k)/4}, \mathrm{wt}(\boldsymbol{e}_{11}) = p/4\}
\end{aligned}$$

and g defined as $g(v_{00}, v_{01}, v_{10}, v_{11}) = 0$ if and only if $\boldsymbol{e} = (\boldsymbol{e}'||\boldsymbol{e}'')$ is of weight w where $\boldsymbol{e}' = v_{00} + v_{01} + v_{10} + v_{11}$ and $\boldsymbol{e}''^T = \boldsymbol{s}''^T - \boldsymbol{H}''\boldsymbol{e}'^T$.

A.1. Quantum Algorithms for SDP. Kachigar and Tillich in [18] presented a method for solving SDP using a dedicated algorithm for a quantum computer, which consists in rephrasing it as a 4-sum problem and then using a quantum walk (and Grover search). To solve the 4-sum problem with a combination of Grover's search algorithm and a random walk in a product of 4 copies of a Johnson graph, we make the assumption that $|V_i| = V$ for all i, and that $G = G_0 \times G_1$. We denote by $\pi_i : G \to G_i$ the projection of G onto one of its components. We keep the same notation as in the formulation of the search for $\boldsymbol{e}'$ such that $\boldsymbol{H}'\boldsymbol{e}'^T = \boldsymbol{s}'^T$ as a 4-sum problem. The algorithm is a Grover search for an element $r \in G_1$ such that $g(r) = 1$ where $g : G_1 \to \{0, 1\}$ evaluates to 1 if and only if there exists $(\boldsymbol{v}_{00}, \boldsymbol{v}_{01}, \boldsymbol{v}_{10}, \boldsymbol{v}_{11}) \in V_{00} \times V_{01} \times V_{10} \times V_{11}$ such that

$$\begin{aligned}
\pi_1(f(\boldsymbol{v}_{00})) + \pi_1(f(\boldsymbol{v}_{01})) &= r\\
\pi_1(f(\boldsymbol{v}_{10})) + \pi_1(f(\boldsymbol{v}_{11})) &= \pi_1(s) - r\\
\pi_0(f(\boldsymbol{v}_{00})) + \pi_0(f(\boldsymbol{v}_{01})) + \pi_0(f(\boldsymbol{v}_{10})) + \pi_0(f(\boldsymbol{v}_{11})) &= \pi_0(s)\\
g(\boldsymbol{v}_{00}, \boldsymbol{v}_{01}, \boldsymbol{v}_{10}, \boldsymbol{v}_{11}) &= 0
\end{aligned}$$

The overall cost of this procedure is $O\left(\sqrt{|G_1|}\mathrm{Cost}(\mathcal{O}_g)\right)$. The cost of the Grover oracle $\mathcal{O}_g$ is determined by the strategy we employ to find a quadruplet with the desired properties. In [18], Kachigar and Tillich use a quantum walk in the product of the 4 Johson graphs defined by the V_i's and subsets $U_i \subseteq V_i$ of cardinality $U = \Theta\left(V^{4/5}\right)$. By using a similar data structure and update strategy as in [9], they show the following statement.

Proposition 1 (Prop. 3 of [18]). *Assuming that* $|G_1| = \Omega(V^{4/5})$, *and* $|G| = \Omega(V^{8/5})$, *it is possible to set up a data structure of size* $O(U)$ *such that the above quantum walk takes time* $\tilde{O}\left(V^{4/5}\right)$.

Hence the time (i.e. cost expressed in terms of circuit depth) taken to solve our 4-sum problem is in $\tilde{O}\left(|G_1|^{1/2}V^{4/5}\right)$ (see [18, Prop. 2]).

Let us see how this applies to the time complexity of Algorithm 6 for solving the ISD problem. We use Grover's quantum search to determine an appropriate $\pi \in \mathsf{S}_n$ (main "for" loop from Step 1 to Step 9). The oracle we denote by $g_{\text{perm}} : \mathsf{S}_n \to \{0,1\}$ satisfies $g_{\text{perm}}(\pi) = 1$ if and only if Steps 2 to 7 lead to the creation of an appropriate $e \in \mathbb{F}_q^n$. Hence, the cost of $\mathcal{O}_{g_{\text{perm}}}$ is dominated by that of solving the 4-sum problem (i.e. Step 4 of Algorithm 6). We denote by ε the proportion of marked elements (i.e. good permutation π leading to a solution of the 4-sum problem). Hence the total cost of the procedure is

$$O\left(\frac{1}{\sqrt{\varepsilon}}\mathrm{Cost}\left(\mathcal{O}_{g_{\text{perm}}}\right)\right) = \tilde{O}\left(\frac{1}{\sqrt{\varepsilon}}|G_1|^{1/2}V^{4/5}\right).$$

For a given e of weight w, a permutation π yields a solution to the 4-sum problem if the $k+l$ positions of $\mathscr{S}$ are split into 4 sets of size $\frac{k+l}{4}$ containing exactly $p/4$ non-zero coefficients each. We denote by $(q-1)N_w$ the number of possible e's of weight w, which yields

$$\varepsilon = (q-1)N_w \frac{\binom{\frac{k+l}{4}}{\frac{p}{4}}^4 \binom{n-k-l}{w-p}}{\binom{n}{w}}.$$

Meanwhile, $G = \mathbb{F}_q^l$ and $G_1 = \mathbb{F}_q^{l/2}$, while the cardinality V of the V_i's satisfies

$$V = (q-1)^{\frac{p}{4}}\binom{\frac{k+l}{4}}{\frac{p}{4}}.$$

Finally, the condition $|G| = \Omega\left(V^{8/5}\right)$ induces the constraint on l and p:

$$l \geq \frac{8}{5}\log_q\left((q-1)^{\frac{p}{4}}\binom{\frac{k+l}{4}}{\frac{p}{4}}\right).$$

The total time is obtained by finding the optimum of this cost when l and p vary.

A.2. Representation Technique and $1+1=0$. The solution relying on the quantum algorithm to solve the 4-sum problem can be optimized using techniques of [20, 8]. This consists in restricting the search space of elements in G that yield a solution $(\boldsymbol{v}_{00}, \boldsymbol{v}_{01}, \boldsymbol{v}_{10}, \boldsymbol{v}_{11})$ to the 4-sum problem. To do this, we need to increase the ways we can represent a solution $\boldsymbol{e}'$ to the system $\boldsymbol{H}'\boldsymbol{e}'^T = \boldsymbol{s}'^T$. The first way we can do this, introduced in [20], is called the *representation technique*. It consists in relaxing the conditions on the positions of the positions of the $p/4$ non-zero coordinates of the solutions $\boldsymbol{v}_{00}, \boldsymbol{v}_{01}, \boldsymbol{v}_{10}, \boldsymbol{v}_{11}$. Previously, we assumed the permutation π mapped the p indices in $\mathscr{S}$ to 4 groups of size $p/4$ each within

$[1,(k+l)/4]$, $[(k+l)/4+1,(k+l)/2]$, $[(k+l)/2+1,3(k+l)/4]$, $[3(k+l)/4+1,k+l]$. Thus, we requested that the $\boldsymbol{v}_i$' be of the form

$$\begin{aligned}\boldsymbol{v}_{00} &= (\boldsymbol{e}_{00}, \mathbf{0}_{3(k+l)/4})\\ \boldsymbol{v}_{01} &= (\mathbf{0}_{(k+l)/4}, \boldsymbol{e}_{01}, \mathbf{0}_{(k+l)/2})\\ \boldsymbol{v}_{10} &= (\mathbf{0}_{(k+l)/2}, \boldsymbol{e}_{10}, \mathbf{0}_{(k+l)/4})\\ \boldsymbol{v}_{11} &= (\mathbf{0}_{3(k+l)/4}, \boldsymbol{e}_{11})\end{aligned}$$

for $\mathrm{wt}(\boldsymbol{e}_{00}) = \mathrm{wt}(\boldsymbol{e}_{01}) = \mathrm{wt}(\boldsymbol{e}_{10}) = \mathrm{wt}(\boldsymbol{e}_{11}) = p/4$. Instead of this, we may only assume that π maps the p indices in $\mathscr{S}$ to 2 groups of size $p/2$ each within $[1,(k+l)/2]$ and $[(k+l)/2+1,k+l]$. In this case, we can write $\boldsymbol{e}' = \boldsymbol{v}_{00} + \boldsymbol{v}_{01} + \boldsymbol{v}_{10} + \boldsymbol{v}_{11}$ with $\boldsymbol{v}_{00}, \boldsymbol{v}_{01}, \boldsymbol{v}_{10}, \boldsymbol{v}_{11}$ of the shape

$$\begin{aligned}\boldsymbol{v}_{00} &= (\boldsymbol{e}_{00}, \mathbf{0}_{(k+l)/2})\\ \boldsymbol{v}_{01} &= (\mathbf{0}_{(k+l)/2}, \boldsymbol{e}_{01})\\ \boldsymbol{v}_{10} &= (\boldsymbol{e}_{10}, \mathbf{0}_{(k+l)/2})\\ \boldsymbol{v}_{11} &= (\mathbf{0}_{(k+l)/2}, \boldsymbol{e}_{11})\end{aligned}$$

for $\mathrm{wt}(\boldsymbol{e}_{00}) = \mathrm{wt}(\boldsymbol{e}_{01}) = \mathrm{wt}(\boldsymbol{e}_{10}) = \mathrm{wt}(\boldsymbol{e}_{11}) = p/4$. This way, $\boldsymbol{v}_{00}+\boldsymbol{v}_{01}$ and $\boldsymbol{v}_{10}+\boldsymbol{v}_{11}$ both have weight $p/2$ with half of their non-zero coordinates in $[1,(k+l)/2]$, and the other half in $[(k+l)/2+1,k+l]$. Each choice of $p/4$ coordinates of $\boldsymbol{e}'$ within its $p/2$ non-zero coordinates in $[1,(k+l)/2]$ and $p/4$ coordinates within its $p/2$ non-zero coordinates in $[(k+l)/2+1,k+l]$ fixes a quadruplet $\boldsymbol{v}_{00}, \boldsymbol{v}_{01}, \boldsymbol{v}_{10}, \boldsymbol{v}_{11}$ with the shape described above such that $\boldsymbol{e}' = \sum_i \boldsymbol{v}_i$. Therefore we can re-write $\boldsymbol{e}'$ in $\binom{p/2}{p/4}^2$ different ways. With this relaxation, a subset of the original search space over the parameter $r \in G$ yields the solution $\boldsymbol{e}'$ to the system $\boldsymbol{H}'\boldsymbol{e}'^T = \boldsymbol{s}'^T$. In [8], a further refinement of this technique was introduced to take advantage of potential cancellations of coefficients in the sum $(\boldsymbol{v}_{00}+\boldsymbol{v}_{01}) + (\boldsymbol{v}_{10}+\boldsymbol{v}_{11})$ when the weight of $\boldsymbol{v}_{00}+\boldsymbol{v}_{01}$ and of $\boldsymbol{v}_{10}+\boldsymbol{v}_{11}$ are $\frac{p}{2}+\Delta p$ for some Δp. Indeed, if the weight of $\boldsymbol{v}_{00}+\boldsymbol{v}_{01}$ is $\frac{p}{2}+\Delta p$, then if the Δp extra non-zero coefficients of $\boldsymbol{v}_{00}+\boldsymbol{v}_{01}$ are on the same indices as the Δp extra non-zero coefficients of $\boldsymbol{v}_{10}+\boldsymbol{v}_{11}$, then these will cancel and thus will not contribute to the weight of $\boldsymbol{e}'$. This method was described as "$1+1=0$" since it took advantage of the fact that over $\mathbb{F}_2$, 1's in matching indices canceled out. Over $\mathbb{F}_q$, this could be rephrased as "$x+(q-x)=0$" meaning that if $\boldsymbol{v}_{00}+\boldsymbol{v}_{01}$ has coefficient $x \neq 0$ at the index i, and if $\boldsymbol{v}_{10}+\boldsymbol{v}_{11}$ has coefficient $q-x$ at index i, then the coefficient i of $\boldsymbol{e}' = (\boldsymbol{v}_{00}+\boldsymbol{v}_{01}) + (\boldsymbol{v}_{10}+\boldsymbol{v}_{11})$ is zero and thus does not contribute to $\mathrm{wt}(\boldsymbol{e}')$. The search for a solution of the 4-sum problem is therefore over the new V_i's given by

$$\begin{aligned}V_{00} = V_{10} &= \left\{(\boldsymbol{e}_0, \mathbf{0}_{(k+l)/2}) \in \mathbb{F}_q^{l+k}, \boldsymbol{e}_0 \in \mathbb{F}_q^{(l+k)/2}, \mathrm{wt}(\boldsymbol{e}_0) = \frac{p}{4} + \frac{\Delta p}{2}\right\}\\ V_{01} = V_{11} &= \left\{(\mathbf{0}_{(k+l)/2}, \boldsymbol{e}_1) \in \mathbb{F}_q^{l+k}, \boldsymbol{e}_1 \in \mathbb{F}_q^{(l+k)/2}, \mathrm{wt}(\boldsymbol{e}_1) = \frac{p}{4} + \frac{\Delta p}{2}\right\}\end{aligned}$$

The solution $\boldsymbol{e}' \in \mathbb{F}_q^{l+k}$ of weight p can be represented in $\binom{p/2}{p/4}^2 \binom{\frac{k+l}{2} - \frac{p}{2}}{\frac{\Delta p}{2}}^2 (q-1)^{\Delta p}$ different ways as $\boldsymbol{e}' = v_{00} + v_{01} + v_{10} + v_{11}$ where $v_i \in V_i$. This is due to the fact that for each choice of the $p/2$ non-zero coordinates of $v_{00}+v_{01}$, we can choose an additional Δp indices among the $k+l-p$ indices where e' has a zero coefficient (split evenly between $[1,(k+l)/2]$ and $[(k+l)/2+1,k+l]$), together with Δp non-zero coordinates of $v_{00}+v_{01}$ at these indices.

Now that more r's in G can yield a 4-tuple solution to the 4-sum problem, we restrict the search space accordingly by writing $G = G_0 \times G_1 \times G_2$. In this new setting, G_1 where we search r is replaced by $G_1 \times G_2$, and we only search for an r_1 in G_1 (having fixed the r_2 coordinate of $r = (r_1, r_2)$ arbitrarily). We denote by π_0, π_2, π_{12} the projections of G onto G_0, G_1, and $G_1 \times G_2$ respectively. The size of G_2 is adjusted so that for a given (arbitrary) choice of $r_2 \in G_2$, there is only one $r = (r_1, r_2) \in G_1 \times G_2$ such that there is $(v_{00}, v_{01}, v_{10}, v_{11}) \in V_{00} \times V_{01} \times V_{10} \times V_{11}$ such that

$$\begin{aligned}\pi_{12}(f(\boldsymbol{v}_{00})) + \pi_{12}(f(\boldsymbol{v}_{01})) &= r\\ \pi_{12}(f(\boldsymbol{v}_{10})) + \pi_{12}(f(\boldsymbol{v}_{11})) &= \pi_{12}(s) - r\\ \pi_0(f(\boldsymbol{v}_{00})) + \pi_0(f(\boldsymbol{v}_{01})) + \pi_0(f(\boldsymbol{v}_{10})) + \pi_0(f(\boldsymbol{v}_{11})) &= \pi_0(s)\\ g(\boldsymbol{v}_{00}, \boldsymbol{v}_{01}, \boldsymbol{v}_{10}, \boldsymbol{v}_{11}) &= 0\end{aligned}$$

For a choice of r_2, we define the search oracle $g_{r_2} : G_1 \to \{0,1\}$ where $g_{r_2}(r_1) = 1$ if and only if $r = (r_1, r_2)$ satisfies the above conditions for a 4-tuple $(\boldsymbol{v}_i)$. As before, the overall cost of this procedure is $O\left(\sqrt{|G_1|}\mathrm{Cost}(\mathcal{O}_{g_{r_2}})\right)$. The cost of the Grover oracle $\mathcal{O}_{g_{r_2}}$ is determined by the cost of the quantum walk in the product of the 4 Johnson graphs defined by the V_i's and subsets $U_i \subseteq V_i$ of cardinality $U = \Theta\left(V^{4/5}\right)$.

Proposition 2 (Prop. 4 of [18])**.** *Assuming that $|G_1||G_2| = \Omega(V^{4/5})$, $|G| = \Omega(V^{8/5})$, and that there are $\Omega(|G_2|)$ solutions to the 4-sum problem, then the above quantum walk takes time $\tilde{O}\left(V^{4/5}\right)$.*

Hence the time (i.e. cost expressed in terms of circuit depth) taken to solve our 4-sum problem is in $\tilde{O}\left(|G_1|^{1/2}V^{4/5}\right)$.

Let us see how this applies to the time complexity of Algorithm 6 for solving SDP. We use Grover's quantum search to determine an appropriate $\pi \in \mathsf{S}_n$ (main "for" loop from Step 1 to Step 9). The oracle we denote by $g_{\mathrm{perm}} : \mathsf{S}_n \to \{0,1\}$ satisfies $g_{\mathrm{perm}}(\pi) = 1$ if and only if Steps 2 to 7 lead to the creation of an appropriate $\boldsymbol{e} \in \mathbb{F}_q^n$. Hence, the cost of $\mathcal{O}_{g_{\mathrm{perm}}}$ is dominated by that of solving the 4-sum problem (i.e. Step 4 of Algorithm 6). We denote by ε the proportion of marked elements (i.e. good permutation π leading to a solution of the 4-sum problem). Hence the total cost of the procedure is

$$O\left(\frac{1}{\sqrt{\varepsilon}}\mathrm{Cost}\left(\mathcal{O}_{g_{\mathrm{perm}}}\right)\right) = \tilde{O}\left(\frac{1}{\sqrt{\varepsilon}}|G_1|^{1/2}V^{4/5}\right).$$

For a given $\boldsymbol{e}$ of weight w, a permutation π yields a solution to the 4-sum problem if the $k+l$ positions of $\mathscr{S}$ are split into 2 sets of size $\frac{k+l}{2}$ containing exactly $p/2$ non-zero coefficients each. Still denoting the number of weight-w solutions by $(q-1)N_w$, we get

$$\varepsilon = (q-1)N_w \frac{\binom{\frac{k+l}{2}}{\frac{p}{2}}^2 \binom{n-k-l}{w-p}}{\binom{n}{w}}.$$

Meanwhile, we still have $G = \mathbb{F}_q^l$ and to ensure that there are $\Omega(|G_2|)$ solutions to the 4-sum problem, we set $G = \mathbb{F}_q^{l_2}$ for

$$l_2 := \log_q \left(\underbrace{\binom{p/2}{p/4}^2 \binom{\frac{k+l}{2} - \frac{p}{2}}{\frac{\Delta p}{2}}^2 (q-1)^{\Delta p}}_{\text{number of representations of } e'} \right)$$

This yields $G_1 = \mathbb{F}_q^{\frac{l}{2} - l_2}$, while the cardinality V of the V_i's satisfies

$$V = (q-1)^{\frac{p}{4} + \frac{\Delta p}{2}} \binom{\frac{k+l}{2}}{\frac{p}{4} + \frac{\Delta p}{2}}.$$

Finally, the condition $|G| = \Omega\left(V^{8/5}\right)$ induces the constraint on l and p:

$$l \geq \frac{8}{5} \log_q \left((q-1)^{\frac{p}{4} + \frac{\Delta p}{2}} \binom{\frac{k+l}{2}}{\frac{p}{4} + \frac{\Delta p}{2}} \right).$$

The total time is obtained by finding the optimum of this cost when l, p, and Δp vary.

Appendix B. Leon's algorithm

In this Appendix we provide further details about Leon's algorithm. We start by proving Proposition 3.

Proposition 3 ([10])**.** Let $\mathfrak{C}_1 \subseteq \mathbb{F}_q^n$ be a random code with dimension k, $\pi \xleftarrow{\$} \mathsf{S}_n$ and $\mathfrak{C}_2 = \pi(\mathfrak{C}_1)$. The cost of Leon's algorithm, running with parameter $w \in \mathbb{N}$, $w \leq n$, can be estimated[4] as

$$O\big(\ln(N_w) C_{ISD}(q, n, k, w)\big).$$

Proof. (*Heuristic*) We first consider the cost of the codewords enumeration in Step 1. For both codes $\mathfrak{C}_1$, $\mathfrak{C}_2$, we need to find all of the N_w codewords with weight w. To this end, we model an ISD algorithm as an oracle that, in each call, returns a random weight-w codeword. We first focus on $\mathfrak{C}_1$: the first ISD call will take time complexity $C_{ISD}(q, n, k, w)/N_w$. In the second call we desire to find a distinct codeword, so that the time complexity of this second call is $C_{ISD}(q, n, k, w)/(N_w - 1)$. If we iterate this reasoning, we get that the codewords enumeration for $\mathfrak{C}_1$ takes time

$$O\bigg(C_{ISD}(q, n, k, w) \cdot \sum_{i=1}^{N_w} \frac{1}{i}\bigg).$$

When N_w is large, we consider that $\sum_{i=1}^{N_w} \frac{1}{i} \approx \ln(N_w)$. The codewords enumeration is repeated for $\mathfrak{C}_2$, with analogous cost: this yields a constant factor 2 in the complexity.

Under the assumption that w is properly chosen, we have that $\mathsf{Mor}_{\mathsf{S}_n}\left(A_w(\mathfrak{C}_1), A_w(\mathfrak{C}_2)\right)$ contains a very small number of elements (ideally, only one). So, we can neglect the cost of steps 2 and 3, and consider only the cost of codewords enumeration. $\square$

[4] Here we use the same estimate derived in [10, Section 2.2], which corresponds to a lower bound for the actual complexity since the cost of steps 2 and 3 is neglected. In other words, the proposition takes into account only the cost of the codewords enumeration phase.

We now give some insight on how the choice of w is expected to affect the algorithm. In the following Proposition we derive a heuristic lower bound on the size of $\mathsf{Mor}_{\mathsf{S}_n}\big(A_w(\mathfrak{C}_1), A_w(\mathfrak{C}_2)\big)$, which is obtained under the (realistic) assumption that weight-w codewords of random codes have random supports.

Proposition 14. Let $\mathfrak{C}_1 \subseteq \mathbb{F}_q^n$ be a random code with dimension k, $\pi \xleftarrow{\$} \mathsf{S}_n$ and $\mathfrak{C}_2 = \pi(\mathfrak{C}_1)$. Then, the set $\mathsf{Mor}_{\mathsf{S}_n}\big(A_w(\mathfrak{C}_1), A_w(\mathfrak{C}_2)\big)$ contains at least $u!$ elements, where $u = \max\left\{1\,,\,\left\lfloor n\left(1-\frac{w}{n}\right)^{N_w}\right\rfloor\right\}$.

Proof. Since $\mathfrak{C}_1$ is random, we use N_w to estimate the number of codewords with weight w. For $i = 1, 2$, let $B_i = \{j \in [1;n] \mid \forall \boldsymbol{x} \in A_w(\mathfrak{C}_i) :\; c_j = 0\}$ and $\bar{B}_i = \{1, \cdots, n\} \setminus B_i$. Note that, for any index $j \in B_i$, we have that all the codewords in $A_w(\mathfrak{C}_i)$ have a null entry in position j, while for any $j \in \bar{B}_i$ there is at least a codeword in $A_w(\mathfrak{C}_i)$ whose j-th entry is non null. Let us now consider a permutation $\sigma \neq \pi$ such that $\sigma(j) = \pi(j)$, for all $j \in \bar{B}_1$: this implies that $\sigma(j) \in B_2$ for all $j \in \bar{B}_1$. Then, clearly $\sigma \in \mathsf{Mor}_{\mathsf{S}_n}\big(A_w(\mathfrak{C}_1), A_w(\mathfrak{C}_2)\big)$. Notice that the number of valid permutations σ is equal to the number of bijections from B_1 to B_2, which is $|B_1|! = |B_2|!$. Note that this is only a lower bound, since there may exist permutations that map $A_w(\mathfrak{C}_1)$ into $A_w(\mathfrak{C}_2)$ even if $\sigma(j) \neq \pi(j)$ for some $j \in \bar{B}_1$.

To complete the proof, we need to estimate the size of B_1. To this end, we rely on the following estimate $|B_1| = n\left(1-\frac{w}{n}\right)^{N_w}$. Indeed, since $\mathfrak{C}_1$ is random, we see any of its codewords as a random vector. Consequently, the probability that an index j is in the support of a codeword with weight w is w/n. Since $\mathfrak{C}_1$ has N_w codewords with weight w, the probability that an index never appear in the supports of all N_w weight-w codewords is $\left(1-\frac{w}{n}\right)^{N_w}$. Multiplying the above probability by n, we obtain an estimate for the average size of B_1. □

The result in the above Proposition de facto sets a theoretical lower bound on the value of w which must be used when running Leon's algorithm.

Appendix C. Beullens' algorithm for PEP

In this appendix we provide details about Beullens' algorithm to solve PEP. Given a pair $(\boldsymbol{x}, \boldsymbol{y}) \in \mathfrak{C}_1 \times \mathfrak{C}_2$ such that $\mathsf{Values}(\boldsymbol{x}) = \mathsf{Values}(\boldsymbol{y})$, we will say that $(\boldsymbol{x}, \boldsymbol{y})$ is a *good collision* if $\pi(\boldsymbol{x}) = \boldsymbol{y}$, and a *bad collision* if $\pi(\boldsymbol{x}) \neq \boldsymbol{y}$. Once colliding pairs of codewords have been obtained, one can employ a probabilistic procedure to retrieve the permutation with some probability. In particular, in [10], the author has considered an approach which works only in case bad collisions do not happen.

C.1. Finding matching codewords. We start by analyzing the routine which produces pairs of colliding codewords. We briefly recall the approach of [10] and provide a heuristic analysis on the number of bad and good collisions which one expects to have, on average. To begin, we observe that for any pair of vectors such that $\boldsymbol{y} = \pi(\boldsymbol{x})$, it must also be that $v\boldsymbol{y} = \pi(v\boldsymbol{x})$ for all $v \in \mathbb{F}_q^*$. This means that, given a pair $(\boldsymbol{x}, \boldsymbol{y})$ of colliding codewords, we are able to produce additional $q-1$ pairs of colliding codewords which, however, do not bring any new information about π. Considering all of these pairs in the permutation reconstruction algorithm is useless. Hence, we can get rid of such additional collisions with the following approach (proposed by Beullens in [10]). Let Lex denote the function that on input a vector $\boldsymbol{a}$ returns $b\boldsymbol{a}$, with $b \in \mathbb{F}_q^*$ such that $\mathsf{Values}(b\boldsymbol{a})$ comes first, in lexicographical order, among the multiset entries of all scalar multiples of $\boldsymbol{a}$. To understand how this function operates, we have reported an example in Figure 4. Embedding the function Lex into the codewords finding algorithm, one can get rid of all unnecessary codewords.

$$\begin{aligned}\mathsf{Values}(1\boldsymbol{a}) &= \{1,2,2,3,3,3,4\}\\ \mathsf{Values}(2\boldsymbol{a}) &= \{1,1,1,2,3,4,4\}\\ \mathsf{Values}(3\boldsymbol{a}) &= \{1,1,2,3,4,4,4\}\\ \mathsf{Values}(4\boldsymbol{a}) &= \{1,2,2,2,3,3,4\}\end{aligned}$$

(A)

$$\begin{array}{lll}\mathbf{1}^\circ & \text{-} & \mathsf{Values}(2\boldsymbol{a}) = \{1,1,1,2,3,4,4\}\\ \mathbf{2}^\circ & \text{-} & \mathsf{Values}(3\boldsymbol{a}) = \{1,1,2,3,4,4,4\}\\ \mathbf{3}^\circ & \text{-} & \mathsf{Values}(4\boldsymbol{a}) = \{1,2,2,2,3,3,4\}\\ \mathbf{4}^\circ & \text{-} & \mathsf{Values}(1\boldsymbol{a}) = \{1,2,2,3,3,3,4\}\end{array}$$

(B)

Figure 4. Example of lexicograph ordering, for the finite field with $q = 5$ elements and a vector $\boldsymbol{a} = (0,3,2,0,0,3,3,2,4,1)$, for which $\mathsf{Lex}(\boldsymbol{a}) = 2\boldsymbol{a}$. Figure (A) shows the multisets of entries for all scalar multiples of $\boldsymbol{a}$, while figure (B) reports the lexicographic order of such multisets.

The full subroutine for finding colliding codewords is shown in Algorithm 7. We observe that including the computation of Lex into the codewords search guarantees that we do not put scalar multiples into the lists X and Y and, consequently, into P. The number of codewords we draw from each code is indicated as L, while w is the Hamming weight of the found codewords.

In the next proposition we compute the average size of P, as well as the number of good and bad collisions.

Proposition 15. Let $\mathfrak{C}_1 \subseteq \mathbb{F}_q^n$ be a random linear code with dimension k, and let $\mathfrak{C}_2 = \pi(\mathfrak{C}_1)$ with $\pi \overset{\$}{\leftarrow} \mathsf{S}_n$. Then, on average P contains $M = M' + M''$ elements, when $M' = L^2/N_w$ is the average number of good collisions and $M'' = (1 - 1/N_w)(q-1)L^2\binom{w+q-3}{w-1}^{-1}$ is that of bad collisions.

Proof. We consider that $\mathfrak{C}_1$ contains N_w codewords with weight w and, according to Assumption 1, assume that the automorphism group of the code is trivial. We first determine the number of good collisions. For each $\boldsymbol{x} \in \mathfrak{C}_1$, we have only one codeword $\boldsymbol{y} \in \mathfrak{C}_2$ such that $\pi(\boldsymbol{x}) = \boldsymbol{y}$. Since ISD returns random codewords of weight w, we have that on average the number of good collisions is given by

$$M' = \sum_{i=1}^{L} i \cdot \frac{\binom{L}{i}\binom{N_w - L}{L-i}}{\binom{N_w}{L}} = \frac{L^2}{N_w}.$$

We now count the number of bad collisions; let us first make some preliminary considerations. First, for any vector $\boldsymbol{a}$, we have that $\mathsf{Lex}(\boldsymbol{a})$ contains at least a 1. Hence, for each $\boldsymbol{x} \in X$, we may assume that $\mathsf{Values}(\boldsymbol{x})$ is a random multiset with one entry equal to 1, and the other $w-1$ ones picked at random over $\mathbb{F}_q^*$. The same goes for each $\boldsymbol{y} \in Y$. To have $\{\boldsymbol{x}, \boldsymbol{y}\} \in P$, it must be $\mathsf{Values}(\boldsymbol{x}) = \mathsf{Values}(\boldsymbol{y})$: since there are $\binom{q+w-3}{w-1}$ ways to choose $w-1$ elements from $\mathbb{F}_q^*$ with repetitions,

Algorithm 7: Algorithm to find and match codewords

```
Data: Number of codewords $L \in \mathbb{N}$, weight $w \in \mathbb{N}$, ISD routine
Input: linear codes $\mathfrak{C}_1, \mathfrak{C}_2 \subseteq \mathbb{F}_q^n$ with dimension $k$
Output: list $P$ containing pairs $(\boldsymbol{x}, \boldsymbol{y}) \in \mathfrak{C}_1 \times \mathfrak{C}_2$, such that
        $\mathsf{Values}(\boldsymbol{x}) = \mathsf{Values}(\boldsymbol{y})$

/* Produce a list $X$ of $L$ codewords from $\mathfrak{C}_1$ with weight $w$ */
$X = \varnothing$;
while $|X| < L$ do
    Call ISD to find $\boldsymbol{x} \in \mathfrak{C}_1$ with weight $w$;
    $X \leftarrow X \cup \{\mathsf{Lex}(\boldsymbol{x})\}$;

/* Produce a list $Y$ of $L$ codewords from $\mathfrak{C}_2$ with weight $w$ */
$Y = \varnothing$;
while $|Y| < L$ do
    Call ISD to find $\boldsymbol{y} \in \mathfrak{C}_2$ with weight $w$;
    $Y \leftarrow Y \cup \{\mathsf{Lex}(\boldsymbol{y})\}$;

/* Find collisions between the lists $X$ and $Y$ */
for $\{\boldsymbol{x}, \boldsymbol{y}\} \in X \times Y$ do
    if $\mathsf{Values}(\boldsymbol{x}) = \mathsf{Values}(\boldsymbol{y})$ then
        $P \leftarrow P \cup \{\boldsymbol{x}, \boldsymbol{y}\}$;

return $P$;
```

assuming that such elements are drawn at random, we have that a collision between $\mathsf{Values}(\boldsymbol{x})$ and $\mathsf{Values}(\boldsymbol{y})$ is expected to happen with probability $\binom{q+w-3}{w-1}^{-1}$. Hence, the number of bad collisions can be estimated as

$$M'' = (L^2 - M')(q-1)\binom{w+q-3}{w-1}^{-1}.$$

Indeed, we have $L^2 - M'$ possible pairs that do not give rise to a good collision, and a fraction $\binom{w+q-3}{w-1}^{-1}$ of these is expected to yield to a bad collision. Then, the list size of P is given by $M' + M''$. □

A confirmation of the heuristics we have used for the Proposition is shown in Section C.3, where we compare the performances of the algorithm with those of numerical simulations. The complexity of executing Algorithm 7 is computed in the next proposition.

Proposition 16. *Let $\mathfrak{C}_1 \subseteq \mathbb{F}_q^n$ be a random code with dimension k, and $\mathfrak{C}_2 = \pi(\mathfrak{C}_1)$, with π being a randomly picked permutation. Then, the complexity of running algorithm 7 with parameter L and w is*

$$O\left(L\left(\log_2(L) + (q-1)w\log_2^2(q)\right) + M' + M'' + \frac{C_{ISD}(q,n,k,w)}{N_w} \cdot \frac{\ln\left(1-\frac{L}{N_w}\right)}{\ln\left(1-\frac{1}{N_w}\right)}\right).$$

Proof. (Heuristic) We start by estimating the number of ISD calls to find L distinct codewords. We consider that each ISD call costs $C_{ISD}(n,k,q,w)/N_w$. Now: the average number of distinct codewords we find, after u calls, is $N_w\left(1-\left(1-\frac{1}{N_w}\right)^u\right)$. Since we want this quantity to be equal to L, it must be $u=\frac{\ln\left(1-\frac{L}{N_w}\right)}{\ln\left(1-\frac{1}{N_w}\right)}$. Then, the average cost of calling ISD is $O\left(\frac{C_{ISD}(q,n,k,w)}{N_w}\cdot\frac{\ln\left(1-\frac{L}{N_w}\right)}{\ln\left(1-\frac{1}{N_w}\right)}\right)$. For each found codeword we compute the value of Lex, which comes with a cost of $(q-1)w\log_2^2(q)$ (since we must compute the $(q-1)$ scalar multiple of each found codeword, having weight w). Then, we have to produce the merged list, which can be done efficiently if one firsts hashes the entries of lists X and Y and then uses a binary search algorithm. This comes with a cost that is estimated as $L\log_2(L)$. Finally, we consider that the list P contains $M'+M''$ elements, and consider such a value as the estimate for the complexity to build the list. □

Remark 5. If $L \ll N_w$ (as we expect), then $\ln(1-\frac{L}{N_w}) \approx -\frac{L}{N_w}$ and $\ln(\frac{1}{N_w}) \approx -\frac{1}{N_w}$. Then, the cost of ISD becomes $O\left(\frac{LC_{ISD}(q,n,k,w)}{N_w}\right)$: this means that we call ISD for L times, and that every call costs $\frac{C_{ISD}(q,n,k,w)}{N_w}$. Embedding this simplification into the expression of 16, we obtain a time complexity of

$$O\left(L\log_2(L)+L(q-1)w\log_2^2(q)+M'+M''+\frac{LC_{ISD}(q,n,k,w)}{N_w}\right).$$

Furthermore, considering that the cost of ISD is expected to be prevalent, with respect to the other terms, we can simply assume that Algorithm 7 costs

$$O\left(\frac{LC_{ISD}(q,n,k,w)}{N_w}\right)$$

C.2. Probabilistic permutation recovery. We now move on to assessing the performance of the permutation reconstruction phase described in [10]. The algorithm exploits the following crucial observation: if we know that $\pi(\boldsymbol{x})=\boldsymbol{y}$ for some permutation π, and we have $x_i \neq y_j$, then we know that π does not map i to j. Considering all pairs of indexes (i,j) for which $x_i \neq y_j$, we gather a significant amount of information about π or, to put it differently, we filter out a wide number of candidate permutations. Exploiting all pairs in P, and putting all the information together, it may become possible to recover the secret π with a procedure as simple as the one in Algorithm 8.

We note that, however, such an efficient permutation recovery is characterized by a certain failure probability. Indeed, when P is not populated with a sufficient number of elements, it may happen that the algorithm is not able to return a valid matrix. To estimate the probability that Algorithm 8 is successful, and to additionally derive the minimum number of elements that P should contain, we rely on the following proposition.

Proposition 17. Let $\mathfrak{C}_1 \subset \mathbb{F}_q^n$ be a random linear code with dimension k and $\mathfrak{C}_2=\pi(\mathfrak{C}_1)$, with $\pi \xleftarrow{\$} \mathsf{S}_n$. Let P be the list produced by Algorithm 7 with input parameters L and w. Let L and w such that $M''\left(1-\frac{1}{N_w}\frac{L^2}{\binom{w+q-3}{w-1}}\right)\approx 0$. Then, Algorithm 8 runs in time $O\left(n^2\frac{L^2}{N_w}\right)$, and retrieves the correct permutation with probability $(1-\rho^{\frac{L^2}{N_w}})^{n(n-1)}$, where $\rho=\left(1-\frac{w}{n}\right)^2+\frac{1}{q-1}\left(\frac{w}{n}\right)^2$.

Algorithm 8: Probabilistic permutation recovery, for the permutation equivalence version of Beullens' algorithm

```
Input: list $P$, containing pairs $\{\boldsymbol{x}, \boldsymbol{y}\} \in \mathbb{F}_q^n \times \mathbb{F}_q^n$
Output: permutation $\pi \in \mathsf{S}_n$, or report failure

$\boldsymbol{U} \leftarrow n \times n$ matrix made of all ones;
for $\{\boldsymbol{x}, \boldsymbol{y}\} \in P$ do
    for $i \in \{1, \cdots, n\}$ do
        for $j \in \{1, \cdots, n\}$ do
            if $x_i \neq y_j$ then
                $u_{i,j} = 0$

/* Use $\boldsymbol{U}$ to reconstruct the permutation; if not possible, report failure */
if $\boldsymbol{U}$ is a permutation matrix then
    return $\pi$;
else
    report failure
```

Proof. We assume that L and w are such that P does not contain bad collisions. Hence, we have $|P| = M' = \frac{L^2}{N_w}$. To check each pair in P, the algorithm uses $O(n^2)$ operations (since it goes through all pairs of indexes $(i,j) \in \{1, \ldots, n\}^2$). We now proceed to estimate the success probability. We consider a pair of codewords $\{\boldsymbol{x}, \boldsymbol{y}\} \in P$ and a pair of indexes (j, ℓ), and consider the probability that we have $x_j = y_\ell$. This probability is given by

$$\rho = \left(1 - \frac{w}{n}\right)^2 + \frac{1}{q-1}\left(\frac{w}{n}\right)^2.$$

The algorithm will succeed if, for all possible pairs (j, ℓ), we have at least a couple of codewords $\{\boldsymbol{x}, \boldsymbol{y}\}$ for which $x_j \neq y_\ell$. Given that P contains $\frac{L^2}{N_w}$ pairs, and that we have $x_j \neq y_\ell$ must happen for $n(n-1)$ pairs of indexes, we have that the success probability can be estimated as $\left(1 - \rho^{\frac{L^2}{N_w}}\right)^{n(n-1)}$. □

Remark 6. The probability in Proposition 17 is an approximation of the actual success probability. Indeed, the proof of the proposition assumes that all pairs of indexes (j, ℓ) for which we have $x_j = y_\ell$ behave as random and uncorrelated variables, which is clearly ideal. Indeed, the actual distribution depends on the supports of the codewords which are present in P and, especially when P is small, the pairs of indexes are heavily correlated. Taking this phenomenon into account in a more accurate way would require a much more involved analysis. In any case, the probability expressed by Proposition 17 offers a crude, but appropriate and simple, approximation of the actual probability.

C.3. Numerical confirmation. In this section we present the results of some numerical simulations we have run, in order to validate the analysis of Beullens' algorithm we have performed in the previous sections. For our simulations, we have fully implemented the algorithm using Sage; the code we have used for the experiments is made fully available[5].

[5] https://github.com/paolo-santini/LESS_project

We start by considering the codewords finding algorithm; we have run the following experiment:

1. generate random pairs of codes, one being a permutation of the other;
2. for each couple of codes, cal ISD to find L distinct codewords with weight w;
3. for each couple of codes, run Algorithm 7 and have empirically measured the values of M' and M'';
4. average the obtained values of M' and M'', and compare with the theoretical estimates in Proposition 15.

In Table 4 we have reported a comparison between the theoretical values and the empirical ones.

(n,k,q)	w	L	M'		M''		$\lvert P\rvert$	
			th.	emp.	th.	emp.	th.	emp.
$(50,25,5)$	13	12	7.2	8.2	1.2	3.3	8.4	11.5
	14	100	47.3	47.6	71.1	268.9	118.4	316.5
$(40,10,11)$	23	40	31.5	31.0	$7.78\cdot 10^{-4}$	0	31.5	31
	24	145	58.4	58.2	$7.4\cdot 10^{-3}$	0.1	58.3	58.3
$(40,10,23)$	26	250	52.7	54.2	$1.9\cdot 10^{-7}$	0	52.7	54.2
	28	2000	28.9	29.1	$3.9\cdot 10^{-6}$	0	28.9	29.1
$(30,10,31)$	8	90	51.9	51.9	$2.9\cdot 10^{-2}$	0	51.9	51.9
	9	200	3.5	4.5	$3.1\cdot 10^{-2}$	0	3.5	4.5

Table 4. Comparison between numerical results and theoretical estimates on the composition of list P. For each triplet (n,k,q), the empirical results have been averaged over 10 random codes.

Finally, we have considered also the success probability of the permutation reconstruction algorithm; in Figure 5 we compare the obtained values with the theoretical ones, obtained through Proposition 17. The results show that Proposition 17 offer indeed a realistic approximation of the actual probability.

Appendix D. Computing First Lexicographic Basis

In this section we describe how the computation of Lex can be extended to the case of two-dimensional spaces; to avoid confusion with the one-dimensional case, we will refer to the operation as $\mathsf{Lex}^{(2)}$. We define $\mathsf{Lex}^{(2)}$ as the function that, on input $\boldsymbol{V}\in\mathbb{F}_q^{2\times n}$, returns the matrix $\boldsymbol{BV}$, with $\boldsymbol{B}\in\mathsf{GL}_2$ and such that the lexicographic minimum of $\{\tau(\boldsymbol{BV})\mid\tau\in\mathsf{M}_n\}$ does not come after the lexicographic minimum of each $\{\tau(\boldsymbol{B}^*\boldsymbol{V})\mid\tau\in\mathsf{M}_n\}$, with $\boldsymbol{B}^*\in\mathsf{GL}_2\setminus\boldsymbol{B}$. Note that this definition is a generalization to the two-dimensional case of the Lex function we have already considered in Section C.1.

To compute the first lexicographic basis of all possible monomial transformations of a matrix $\boldsymbol{V}$, we can operate as follows. We multiply each column of $\boldsymbol{V}$ by the inverse of the element in the first row, in order to remain with either zeros or ones in the first row. Now, we permute the columns of the obtained matrix with the goal of placing the zeros in the leftmost part of the first row. To do this, we consider the element in the second row: when two columns have the same element in the first row, we look at the element in the second row, and put on the left the one with the lowest entry. Finally, if we have some non null entry in the second row which

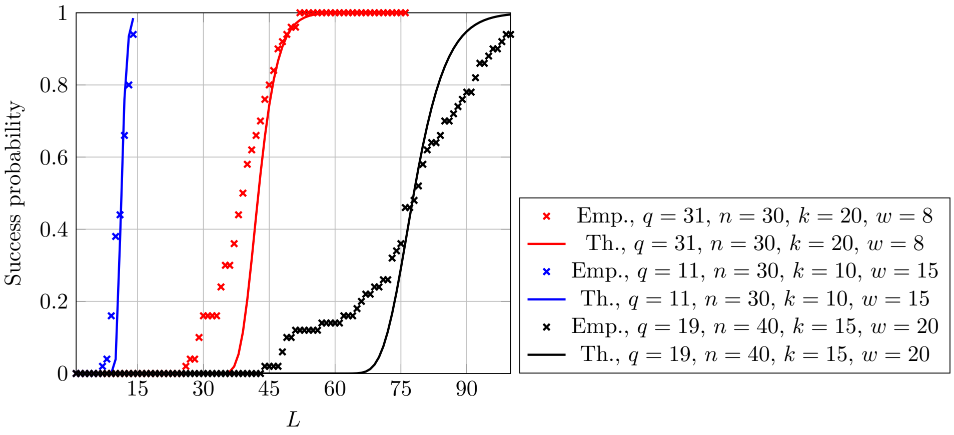

FIGURE 5. Success probability of the probabilistic permutation recovery as a function of L, for codes with different parameters. For every configuration, we have tested the attack on 50 codes. The empirical success probability has been computed by averaging over the trials.

corresponds to a null entry in the first row, we can scale the corresponding column to put a one in the second row. For the sake of clarity, in Figure 6 we show an example of this procedure.

Then, given an input matrix $\boldsymbol{V} \in \mathbb{F}_q^{2\times n}$, we compute the matrices $\boldsymbol{BV}$, with $\boldsymbol{B} \in \mathsf{GL}_2$. We then perform the operations shown in Figure 6 and, for each $\boldsymbol{BV}$, keep only the resulting lexicographically minimum matrix. Finally, we compare all of such matrices and pick the one which comes first, in the lexicographically order. Notice that, for each input matrix, we test a total of $(q^2-1)(q^2-q)$ basis, and for each basis we use $O(n)$ operations to find the lexicographic minimum matrix. Each comparison requires $O(2n)$ operations, so that the computation of $\mathsf{Lex}^{(2)}$ costs $O\left(n(q^2-1)(q^2-q)\right)$ operations.[6]

APPENDIX E. CONSIDERATIONS ABOUT $\mathsf{Lex}^{(2)}$ AND TWO-DIMENSIONAL EQUIVALENT CODES

In this appendix we estimate the probability to have bad collisions when considering the computation of $\mathsf{Lex}^{(2)}$ and Values on two-dimensional subcodes. We start with the following technical Lemma.

Lemma 3. Let $\boldsymbol{V}_1, \boldsymbol{V}_2 \in \mathbb{F}_q^{2\times n}$, with $\boldsymbol{A}_1 = \mathsf{Lex}^{(2)}(\boldsymbol{V}_1)$ and $\boldsymbol{A}_2 = \mathsf{Lex}^{(2)}(\boldsymbol{V}_2)$ being such that $\mathsf{Values}(\boldsymbol{A}_1) = \mathsf{Values}(\boldsymbol{A}_2)$. Then, the codes generated by $\boldsymbol{V}_1$ and $\boldsymbol{V}_2$ are equivalent.

Proof. We observe that, if $\mathsf{Values}(\boldsymbol{A}_1) = \mathsf{Values}(\boldsymbol{A}_2)$, then this means that there exist two monomials $\tau_1, \tau_2 \in \mathsf{M}_n$ such that $\tau_1(\boldsymbol{A}_1) = \tau_2(\boldsymbol{A}_2)$. Let $\boldsymbol{Q}_1$ and $\boldsymbol{Q}_2$ be the associated matrices, and let $\boldsymbol{A}_1 = \boldsymbol{B}_1\boldsymbol{V}_1$, $\boldsymbol{A}_2 = \boldsymbol{B}_2\boldsymbol{V}_2$. Then, we have

$$\boldsymbol{B}_1\boldsymbol{V}_1\boldsymbol{Q}_1 = \boldsymbol{B}_2\boldsymbol{V}_2\boldsymbol{Q}_2 \implies \boldsymbol{V}_1 = \boldsymbol{B}_1^{-1}\boldsymbol{B}_2\boldsymbol{V}_2\boldsymbol{Q}_2\boldsymbol{Q}_1^{-1},$$

[6] Clearly, one can improve the computation of $\mathsf{Lex}^{(2)}$ using small weight codewords: this avoids to consider all $(q^2-1)(q^2-q)$ changes of basis. However, taking this into account would burden the description. Since this aspect does not affect significantly the complexity of the algorithms we analyze, we chose to consider only the trivial (and, perhaps, naive) computation for $\mathsf{Lex}^{(2)}$.

$$\begin{pmatrix} 0 & 1 & 0 & 0 & 2 & 3 & 2 & 0 & 4 \\ 1 & 0 & 0 & 2 & 0 & 3 & 4 & 0 & 2 \end{pmatrix}$$

(A)

$$\begin{pmatrix} 0 & 1 & 0 & 0 & 1 & 1 & 1 & 0 & 1 \\ 1 & 0 & 0 & 2 & 0 & 1 & 2 & 0 & 3 \end{pmatrix}$$

(B)

$$\begin{pmatrix} 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 & 1 \\ 0 & 0 & 1 & 2 & 0 & 0 & 1 & 2 & 3 \end{pmatrix}$$

(C)

$$\begin{pmatrix} 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 & 1 \\ 0 & 0 & 1 & 1 & 0 & 0 & 1 & 2 & 3 \end{pmatrix}$$

(D)

Figure 6. Example of lexicographic ordering of a basis, for the finite field with $q = 5$ elements. In figure (A) we show the initial basis, while in the other figures we detail the steps we perform to find the corresponding lexicographic minimum. The matrix in figure (B) is obtained by scaling all columns so that the entry in the first row is a 1. To obtain the matrix in figure (C), we sort the columns. Finally, we see that we have some degrees of freedom, since the third and fourth columns have a zero in the first row and a non null entry in the second row. Hence, we scale these columns and finally obtain the minimum lexicograph basis as in figure (D).

which corresponds to the definition of linear equivalence. □

Proposition 3. *Let $\mathfrak{C} \subseteq \mathbb{F}_q^n$ be a random linear code with dimension k. Let $\mathfrak{B} \subseteq \mathfrak{C}$ be a randomly-chosen, two-dimensional subcode with support size w. Then, the average number of two-dimensional subcodes $\mathfrak{B}' \subseteq \mathfrak{C}$ which are linearly equivalent to $\mathfrak{B}$ is upper bounded by*

$$t_w^{(2)} = \binom{n}{w} w!(q-1)^{w-1} \frac{\begin{bmatrix} k \\ 2 \end{bmatrix}_q}{\begin{bmatrix} n \\ 2 \end{bmatrix}_q}.$$

Proof. The bound is trivially obtained by considering all distinct monomial transformations of a code whose support is w, which is given by $\binom{n}{w} w!(q-1)^{w-1}$ (excluding monomial transformations which are identical, up to a scalar multiplication). Then, we consider that for every such code, there is a probability equal to $\frac{\begin{bmatrix} k \\ 2 \end{bmatrix}_q}{\begin{bmatrix} n \\ 2 \end{bmatrix}_q}$ that it is in $\mathfrak{C}$ (since $\mathfrak{C}$ is random). □

Appendix F. Proof of Proposition 5

Proposition 5. Let $\mathfrak{C}_1 \subseteq \mathbb{F}_q^n$ be a random linear code with dimension k, and let $\mathfrak{C}_2 = \tau(\mathfrak{C}_1)$ with $\tau \xleftarrow{\$} \mathsf{M}_n$. Let P be the list obtained by running Algorithm 1 with parameters L and w, with $w \leq n-k+2$. The algorithm runs in time

$$O\left(L\left(\log_2(L)+(q^2-q)(q-1)\right)+M'+M''+\frac{L}{N_w^{(2)}}C_{ISD}(q,n,k,w)\right),$$

and produces a list P with $M = M' + M''$ elements, where $M' = L^2/N_w^{(2)}$ is the average number of good collisions and $M'' \leq \frac{t_w^{(2)}(L^2-M')}{N_w^{(2)}}$ is that of bad collisions.

Proof. For the running time of the algorithm, it is enough to repeat the same computation performed for Proposition 16. Hence, we only show how the number of good and bad collisions can be computed. The value of M' is estimated with the same reasoning we have adopted for the proof of Proposition 15. To estimate M'', we consider the following facts:

1) let $\boldsymbol{V}_1$ be the basis of a subcode $\mathfrak{V}_1 \in \mathfrak{C}_1$, and assume that $\mathfrak{V}_1$ is equivalent to u subcodes in $\mathfrak{C}_1$. According to Lemma 3, this means that there are u subcodes in $\mathfrak{C}_1$ which lead to a collision in the computation of Lex;

2) since $\mathfrak{C}_2$ is a monomial transformation of $\mathfrak{C}_1$, this means that also $\mathfrak{C}_2$ contains u subcodes that are equivalent to $\mathfrak{B}_1$;

2) we can use $t_w^{(2)}$ to upper bound the value of u (to see how $t_w^{(2)}$ is obtained, check Appendix E). Then, for any pair of drawn subcodes $\mathfrak{V}_1 \subseteq \mathfrak{C}_1$, $\mathfrak{V}_2 \subseteq \mathfrak{C}_2$, the probability that they are equivalent is upper bounded as $\frac{t_w^{(2)}}{N_w^{(2)}}$;

4) given that we draw L subcodes from each code, we have a total of L^2 pairs. We know that M' of them are good collisions, while for each remaining one there is a probability $t_w^{(2)}/N_w^{(2)}$ that it is a bad collision. Then, on average, the number of bad collisions is given by $\frac{t_w^{(2)}(L^2-M')}{N_w^{(2)}}$.

□