A Descent to Hades: Attacks on PKP and PEP over Extension Fields

Alessandro Budroni[1], Jesús-Javier Chi-Domínguez[1], Jorge Chavez-Saab[1], and Andre Esser[1⋆]

Cryptography Research Center, Technology Innovation Institute, Abu Dhabi, UAE,
{alessandro.budroni,jesus.dominguez,jorge.saab,andre.esser}@tii.ae

Abstract. The Permutation Code Equivalence Problem (PEP) and Permuted Kernel Problem (PKP) are two notorious computational problems over linear codes used for building post-quantum digital signature schemes. Although traditionally analyzed over prime fields, recent proposals, such as the signature schemes PERK and SETH, have considered extension fields to improve efficiency and compactness. In this work, we analyze the hardness of solving PKP and PEP when instantiated over field extensions. For binary field extensions, by exploiting a reduction to a structured variant of the Regular Syndrome Decoding Problem (RSD), we uncover new polynomial-time parameter regimes for both PKP and PEP, including families of self-orthogonal PEP instances and all self-dual instances over extensions of degree $\nu > 4$. We also adapt the permutation-based Regular-ISD algorithm of Esser and Santini for RSD (CRYPTO '24) to PKP-derived instances, and uncover regimes of parameters for which it improves upon the state-of-the-art. Moreover, we present a reduction from a broad family of PEP instances over extension fields with odd characteristic to the Graph Isomorphism Problem, yielding a polynomial-time algorithm to solve those instances. Overall, our results invalidate the use of PEP over extension fields for most of the scenarios, and provide novel insights into the security of PKP over extension fields.

Keywords: Permuted Kernel Problem, Permutation Code Equivalence, Regular Syndrome Decoding, Code-based Cryptography, Post-quantum Cryptography

1 Introduction

Replacing a prime field with an extension field is often an attractive design choice in public-key cryptography as it can yield compact representations and efficient arithmetic while preserving the apparent size of the underlying search space. However, this replacement might come at a security cost. The introduced

⋆ Supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – Project-ID MA 2536/12.

algebraic structure can enable more efficient attacks and, in some cases, lead to a complete collapse of hardness.

For instance, Weil descent attacks on the elliptic curve discrete logarithm problem provide an early and influential illustration of this principle. In his seminal work [31], Frey showed how to transfer discrete logarithm instances on elliptic curves over binary extension fields to instances on Jacobians of higher-genus curves defined over smaller subfields. This approach was subsequently refined and extended in a series of works [33,32,38,37]. Further investigations identified broad classes of binary curves for which the resulting higher-genus curve discrete-logarithm problems admit subexponential-time algorithms, thereby providing substantially less security than a generic analysis would suggest [42,41,34,25,26]. These results show that the additional extension-field structure must be carefully accounted for when assessing the security of elliptic curves over non-prime fields. Similar lessons have emerged in code-based cryptography, notably from structural attacks against variants of the McEliece cryptosystem defined using codes over extension fields [27,46].

These examples motivate revisiting the security of cryptographic problems when their base field is replaced by an extension field, even when the underlying computational problem is unchanged. In this work, we study the complexity of two well-known problems in code-based cryptography when defined over extension fields.

The first problem we study is the Permutation Equivalence Problem (PEP) which asks whether two linear codes differ only by a permutation of their coordinates. This is a special case of the more general Linear Code Equivalence Problem (LEP) that underlies the security of the signature scheme LESS [6], which was a second-round contender of the NIST call for additional post-quantum signatures [44], and the threshold signature scheme LEAST [7], a preview submission to the ongoing NIST call for multi-party threshold schemes [45], and several other proposals [12,35,36]. In fact, an earlier version of LESS [16] relied directly on the hardness of PEP, as well as a more recently proposed digital signature SPECK [9]. The updatable public-key encryption scheme of Albrecht, Benčina, and Lai [3] combines LWE and PEP instantiated with self-orthogonal codes. Its authors identify a module-LWE variant as a natural route towards practical instantiations, but leave its security analysis open precisely because it gives rise to PEP instances over field extensions [3, Sec. 1.4], whose complexity has not yet been studied.

The second problem we address is the Permuted Kernel Problem (PKP). Introduced by Shamir [52], it asks for a permutation that maps one or more public vectors into the kernel of a public matrix. In recent years, it has been used to construct signature schemes [29,17,15,30,9], including PERK [1] which was also a second-round contender of the NIST call for additional post-quantum signatures. PKP over extension fields was introduced first in PERK v2.1 [2] to enable an efficient implementation of the VOLE-in-the-Head paradigm using PKP. More recently, for similar reasons and under the equivalent formulation of the Subcode Equivalence Problem, it was used to instantiate the signature

scheme SETH [8]. Despite this motivation, PKP over extension fields has not yet been studied in depth.

In this work, we present new cryptanalytic results for PKP and PEP over extension fields, showing that their use significantly affects the complexity of both problems. By establishing a connection between these two problems and syndrome decoding problems, we obtain new polynomial-time regimes for both problems, as well as a novel approach for solving PKP. Furthermore, for the case of PEP between self-orthogonal codes defined over extension fields, we provide a polynomial time algorithm for a wide family of parameters.

Related Work on PKP. Shamir proved in 1992 that PKP is NP-complete in the one-dimensional case [52], and Bergen, Gueye, and Klamti later extended this result to the multidimensional setting [14]. On the cryptanalytic side, the current state of the art was set by Koussa, Macario-Rat, and Patarin (KMP) [39], who introduced a meet-in-the-middle algorithm for solving PKP. More recently, Santini, Baldi, and Chiaraluce [50], and Budroni, Defranceschi, and Pintore [21], proposed further improvements over KMP that reduce the practical cost of solving PKP in the finite regime.

Related Work on PEP. Leon was the first to propose an algorithm for solving PEP, in 1982 [40]. Later, Petrank and Roth showed that the problem is unlikely to be NP-complete [47]. Sendrier then proposed the first efficient algorithm for solving PEP when the hull of the code is small [51], and this setting was subsequently cryptanalyzed further [49,11]. As a result of these works, PEP is nowadays considered hard only when the hull is large, e.g., with self-orthogonal codes with hull dimension greater than the square root of the code length [13]. For these cases, the best known algorithms are improvements of Leon's work [10,22].

1.1 Our contributions.

We study PKP and PEP over field extensions through descent to the base field. Our main results are as follows.

1. We introduce *Permutation Syndrome Decoding* (PSD), a variant of the syndrome decoding problem in which the solution is the vectorization of a permutation matrix. We provide a transformation that maps a PKP or PEP instance over $\mathbb{F}_{2^\nu}$ to a structured PSD instance over $\mathbb{F}_2$. This provides a common setting for studying these problems.
2. For binary field extensions, we uncover new families of parameters for which PKP and PEP over self-orthogonal codes are solvable in polynomial time. In particular, we show that PEP with self-dual codes (the usual choice in most constructions) over $\mathbb{F}_{2^\nu}$ is solvable in polynomial time whenever $\nu > 4$.
3. For field extensions of odd characteristic p, we apply the Frobenius conjugation to a PEP instance and reduce the original PEP instance to a graph isomorphism problem (GIP) instance that can be solved in polynomial time. This result holds for random (weakly) self-dual and self-orthogonal codes of length n such that $p \not\equiv 1 \bmod 4$ or $n \not\equiv 2 \bmod 4$.

4. We introduce a new algorithm for solving PKP. We first extend the permutation-based Regular-ISD algorithm of Esser and Santini [28], originally developed for RSD, to PSD, explicitly accounting for the additional permutation structure of the sought solution. We then apply the resulting algorithm to PSD instances derived from PKP. Our algorithm improves upon the KMP algorithm [39] in parameter regimes where either the extension degree or the number of solutions is large. The latter regime is particularly noteworthy, as the multiplicity of solutions is not exploited by existing PKP solvers. By contrast, the time complexity of our algorithm decreases linearly in this multiplicity.

Overall, PKP remains hard for carefully chosen parameter regimes. We remark that PERK [2] and SETH [8] are not affected by our attack. In hindsight, their PKP parameters were chosen in regimes where the expected number of solutions is approximately one, and our algorithm therefore does not outperform the best previously known attacks. On the other hand, a module-based instantiation of the construction of Albrecht et al. [3] would necessarily rely on the hardness of PEP over extension fields, which we show insecure for binary field extensions, and for a wide family of parameter sets in odd characteristic field extensions. Finally, our results give a cautionary tale against the use of field extensions in cryptographic applications of PEP.

Artifacts. The scripts to run and reproduce the experiments presented in this manuscript are available at [23].

Organization. After giving the necessary preliminaries in Section 2, we introduce PSD in Section 3. In Section 4 we show how to transform PKP and PSD instances into PSD and uncover new polynomial-time parameter regimes for these problems. In Section 5 we introduce our new attack tailored for PSD and study its impact on PKP. Finally, in Section 6, we present the reduction from PEP to GIP that allows for quasi-polynomial-time algorithms for solving PEP under certain combinations of parameters.

2 Preliminaries

We denote with $\mathbb{N}$ and $\mathbb{Z}$ the sets of natural and integer numbers, respectively. Let $q \in \mathbb{N}$ be a prime or a power of a prime. We write with $\mathbb{F}_q$ the finite field with q elements, and with $\mathbb{F}_q^n$ the vector space over $\mathbb{F}_q$ of dimension n. Given two positive integers $a, b \in \mathbb{N}$, we write $a|b$ to say that a divides b.

We denote with capital letters matrices $\boldsymbol{M}$ and with lower case column vectors $\boldsymbol{v}$. Given a square matrix $\boldsymbol{M} \in \mathbb{F}_q^{n \times n}$, we write $\mathsf{vec}(\boldsymbol{M}) \in \mathbb{F}_q^{n^2}$ the column vector formed by concatenating the columns of $\boldsymbol{M}$. For two matrices $\boldsymbol{A} \in \mathbb{F}_q^{n_1 \times m_1}$ and $\boldsymbol{B} \in \mathbb{F}_q^{n_2 \times m_2}$, the operator $\otimes$ denotes the Kronecker product between $\boldsymbol{A}$ and $\boldsymbol{B}$, i.e., $[\boldsymbol{A} \otimes \boldsymbol{B}] \in \mathbb{F}_q^{n_1 n_2 \times m_1 m_2}$. For $\boldsymbol{U} \in \mathbb{F}_q^{m_1 \times n_2}$, the following property holds

$$[\boldsymbol{A} \otimes \boldsymbol{B}]\mathsf{vec}(\boldsymbol{U}) = \mathsf{vec}(\boldsymbol{B}\boldsymbol{U}\boldsymbol{A}^\top). \tag{1}$$

We denote with vec^{-1} the inverse operation of vec, that is, for a given input vector $\boldsymbol{u} \in \mathbb{F}_q^{n^2}$, it returns the matrix $\boldsymbol{U} \in \mathbb{F}_q^{n \times n}$ whose columns are formed by the entries of $\boldsymbol{v}$ in sequence.

The Hamming weight of a vector $\boldsymbol{v}$ is the number of non-null entries of the vector, and it is denoted with $|\boldsymbol{v}|$. We denote with $\mathsf{Perm}_n(\mathbb{F}_q)$ the set of all $n \times n$ permutation matrices and with $\mathsf{Mono}_n(\mathbb{F}_q)$ the set of all monomial matrices $\boldsymbol{M} = \boldsymbol{P}\boldsymbol{D}$, where $\boldsymbol{P} \in \mathsf{Perm}_n(\mathbb{F}_q)$ and $\boldsymbol{D} \in \mathbb{F}_q^{n \times n}$ is diagonal and full-rank.

An $[n, k]_q$ linear code $\mathcal{C}$ of length n and dimension k is a linear subspace of dimension $k < n$ of $\mathbb{F}_q^n$. We say that $\mathcal{C}$ has *rate* $R \coloneqq k/n$. $\mathcal{C}$ can be represented by either a *generator matrix* $\boldsymbol{G} \in \mathbb{F}_q^{k \times n}$ whose are a basis of $\mathcal{C}$, or by a *parity-check matrix*, that is, a full-rank matrix $\boldsymbol{H} \in \mathbb{F}_q^{(n-k) \times n}$ such that $\boldsymbol{H}\boldsymbol{c} = \boldsymbol{0}$, for every $\boldsymbol{c} \in \mathcal{C}$. The code generated by the rows of $\boldsymbol{H}$ is called the *dual code* of $\mathcal{C}$ and is denoted by $\mathcal{C}^\perp$. The *hull* of a code is defined as $\mathcal{H}(\mathcal{C}) = \mathcal{C} \cap \mathcal{C}^\perp$. We say that $\mathcal{C}$ is *weakly self-dual* if $\dim(\mathcal{H}(\mathcal{C})) > 1$ and that it is self-orthogonal if $\dim(\mathcal{H}(\mathcal{C})) = k$. Moreover, we say that $\mathcal{C}$ is *self-dual* if $\dim(\mathcal{H}(\mathcal{C})) = k = n/2$.

A set $J \subset \{1, \ldots, n\}$ of size k is an information set for $\mathcal{C}$ if any two distinct $\boldsymbol{c}, \boldsymbol{c}' \in \mathcal{C}$ differ in at least one coordinate indexed by J.

2.1 Hard Problems on Linear Codes

We now recall a fundamental problem in coding theory, which underlies the security of several cryptographic constructions in post-quantum cryptography. Since in this work we focus on binary codes, we give the definition already tailored to this case.

Definition 1 (Syndrome Decoding). *Let $k, n, w \in \mathbb{N}$ be such that $k, w \leq n$. The* Syndrome Decoding (SD) Problem *is to find a vector $\boldsymbol{e} \in \mathbb{F}_2^n$ of Hamming weight w such that $\boldsymbol{H}\boldsymbol{e} = \boldsymbol{s}$, given a parity-check matrix $\boldsymbol{H} \in \mathbb{F}_2^{(n-k) \times n}$ and a* syndrome *vector $\boldsymbol{s} \in \mathbb{F}_2^{n-k}$.*

We write $\mathsf{SD}(n, k, w)$ to denote the SD problem with parameters n, k, and w. We now give the definition of a structured variant of the SD problem, which has recently found increasing application in code-based cryptography. We say that a vector $\boldsymbol{e} \in \mathbb{F}_2^n$ of Hamming weight w, where $w|n$, is *regular* if it can be written as $\boldsymbol{e} = (\boldsymbol{e}_1, \boldsymbol{e}_2, \ldots, \boldsymbol{e}_w)$, with each block $\boldsymbol{e}_i \in \mathbb{F}_2^{\frac{n}{w}}$ satisfying $|\boldsymbol{e}_i| = 1$. In other words, $\boldsymbol{e}$ is regular if it contains exactly one nonzero entry in each of the w disjoint blocks of size $\frac{n}{w}$.

Definition 2 (Regular Syndrome Decoding). *Let $k, n, w \in \mathbb{N}$ be such that $k, w \leq n$ and $w|n$. The* Regular Syndrome Decoding (RSD) Problem *is to find a regular vector $\boldsymbol{e} \in \mathbb{F}_2^n$ of Hamming weight w such that $\boldsymbol{H}\boldsymbol{e} = \boldsymbol{s}$, given a parity-check matrix $\boldsymbol{H} \in \mathbb{F}_2^{(n-k) \times n}$ and a* syndrome *vector $\boldsymbol{s} \in \mathbb{F}_2^{n-k}$.*

We write $\mathsf{RSD}(n, k, w)$ to denote the RSD problem with parameters n, k, and w. Motivated by cryptographic applications, we assume that one solution exists always by construction. In particular, a *random* RSD instance is sampled follows:

i) sample $\boldsymbol{H}$ uniformly at random among all full-rank matrices in $\mathbb{F}_2^{(n-k)\times n}$,
ii) sample $\boldsymbol{e} \in \mathbb{F}_2^n$ uniformly at random among all regular vectors of weight w,
iii) return $(\boldsymbol{H},\ \boldsymbol{s} = \boldsymbol{H}\boldsymbol{e})$.

Hence, the expected number of solutions to an RSD problem instance is given by

$$S_{n,k,w}^{(\mathsf{RSD})} \coloneqq 1 + \left(\frac{n}{w}\right)^w 2^{-(n-k)}, \tag{2}$$

where $\left(\frac{n}{w}\right)^w$ is the number of all regular vectors, and $2^{-(n-k)}$ is the probability that any of these vectors is a solution to the linear system $\boldsymbol{H}\boldsymbol{x} = \boldsymbol{s}$. It was proven in [28, Thm. 3.1] that if $w \geq k$, then $\mathsf{RSD}(n, k, w)$ is solvable in polynomial time with high probability. We refer the reader to the work of Esser and Santini [28] for an in-depth study of the concrete and asymptotic complexity, hardness regimes and parameter classification of RSD.

Next, we give the definition of another hard problem in code-based cryptography, originally proposed by Shamir in 1989 [52], which has recently gained increased attention for constructing digital signatures.

Definition 3 (Permuted Kernel Problem). *Let $n, k, \ell \in \mathbb{N}$ be parameters such that $\ell < k < n$. The* Permuted Kernel Problem (PKP) *is to find a permutation $\boldsymbol{P} \in \mathsf{Perm}_n(\mathbb{F}_q)$ such that $\boldsymbol{H}\boldsymbol{P}\boldsymbol{V} = \boldsymbol{0}$, given a full-rank matrix $\boldsymbol{V} \in \mathbb{F}_q^{n\times\ell}$ and a parity-check matrix $\boldsymbol{H} \in \mathbb{F}_q^{(n-k)\times n}$.*

We write $\mathsf{PKP}(n, k, \ell, q)$ the PKP problem with parameters n, k, ℓ, and q. Similarly to above, we assume that one solution always exists by construction. In particular, a *random* PKP instance is sampled as follows:

i) sample a full-rank matrix $\boldsymbol{H} \in \mathbb{F}_2^{(n-k)\times n}$ uniformly at random,
ii) sample ℓ linearly independent vectors uniformly at random from the right-kernel of $\boldsymbol{H}$ and form $\tilde{\boldsymbol{V}} \in \mathbb{F}_q^{n\times\ell}$,
iii) sample $\boldsymbol{P} \in \mathsf{Perm}_n(\mathbb{F}_q)$ uniformly at random and return $(\boldsymbol{H},\ \boldsymbol{V} = \boldsymbol{P}^{-1}\tilde{\boldsymbol{V}})$.

The estimated number of solutions is

$$S_{n,k,\ell,q}^{(\mathsf{PKP})} \coloneqq 1 + n!\, q^{-(n-k)\ell},$$

where $n!$ counts the number of all permutations of order n, and the second factor is the probability that a random vector is orthogonal to the space generated by the rows of $\boldsymbol{H}$. The conditions to ensure that such a solution is unique with high probability are i) $n!\, q^{-(n-k)\ell} < 1$ and ii) both $\boldsymbol{V}$ and $\boldsymbol{H}$ have no repeated columns, implying that $q^\ell \geq n$ and $q^{n-k} \geq n$.

Let us give now the definition of another famous computational problem over linear codes.

Definition 4 (Permutation Equivalence Problem). *Let $n, k \in \mathbb{N}$ be parameters such that $k < n$. Given two generator matrices $\boldsymbol{G}, \boldsymbol{G'} \in \mathbb{F}_q^{k\times n}$, the* Permutation Equivalence Problem (PEP) *is to find a permutation $\boldsymbol{P} \in \mathsf{Perm}_n(\mathbb{F}_q)$ and an invertible matrix $\boldsymbol{S} \in \mathbb{F}_q^{k\times k}$ such that $\boldsymbol{G'} = \boldsymbol{S}\boldsymbol{G}\boldsymbol{P}$.*

We write $\mathsf{PEP}(n, k, q)$ the PEP problem with parameters n, k and q. Notably, PEP is easy when the dimension of the hull of the code generated by $\boldsymbol{G}$ is small [11,51]. For this reason, in this work, we consider self-orthogonal codes, for which the dimension of the hull is maximal. We sample a *random* PEP instance as follows

i) sample at random an $[n, k]_q$ self-orthogonal linear code and let $\boldsymbol{G}$ be a generator,
ii) sample an invertible $\boldsymbol{S} \in \mathbb{F}_q^{k\times k}$ and a permutation $\boldsymbol{P} \in \mathsf{Perm}_n(\mathbb{F}_q)$ uniformly at random,
iii) return $(\boldsymbol{G}, \boldsymbol{G}' = \boldsymbol{SGP})$.

In the following sections, we will consider an equivalent reformulation of PEP that is as follows: find a permutation $\boldsymbol{P} \in \mathsf{Perm}_n(\mathbb{F}_q)$ such that $\boldsymbol{H}'\boldsymbol{P}\boldsymbol{G}^\top = \boldsymbol{0}$, for a given parity-check $\boldsymbol{H}' \in \mathbb{F}_q^{(n-k)\times n}$ and a generator matrix $\boldsymbol{G} \in \mathbb{F}_q^{k\times n}$. Here, $\boldsymbol{H}'$ is a parity-check matrix of the code generated by $\boldsymbol{G}'$ in Definition 4. With this formulation, PEP can be seen as a generalization of PKP for $\ell = k$.

3 Permutation Syndrome Decoding

In this section, we first introduce a variant of the RSD problem in which the secret vector carries additional structure. We then present a reduction from PKP to this new problem and discuss the resulting problem hierarchy.

Let $n \in \mathbb{N}$ be a perfect square. We say that a vector $\boldsymbol{e} \in \mathbb{F}_2^n$ is *permutation regular* if $\mathsf{vec}^{-1}(\boldsymbol{e}) \in \mathsf{Perm}_{\sqrt{n}}(\mathbb{F}_2)$.

Proposition 1. *If a vector $\boldsymbol{e} \in \mathbb{F}_2^n$ is permutation regular, then it is also regular.*

Proof. Let $\boldsymbol{E} \coloneqq \mathsf{vec}^{-1}(\boldsymbol{e})$ be the corresponding permutation matrix, and let $\boldsymbol{e}_1, \ldots, \boldsymbol{e}_{\sqrt{n}} \in \mathbb{F}_2^{\sqrt{n}}$ denote its columns. Since $\boldsymbol{E}$ is a permutation matrix, we have $|\boldsymbol{e}_i| = 1$ for all $i = 1, \ldots, \sqrt{n}$. It follows that $\boldsymbol{e} = \mathsf{vec}(\boldsymbol{E}) = (\boldsymbol{e}_1^\top, \ldots, \boldsymbol{e}_{\sqrt{n}}^\top)$ is a regular vector. □

Notice that a permutation regular vector $\boldsymbol{e}$ of length n always has Hamming weight equal to $\sqrt{n}$. In addition, given the permutation matrix $\boldsymbol{E} \coloneqq \mathsf{vec}^{-1}(\boldsymbol{e})$, then $\tilde{\boldsymbol{e}} = \mathsf{vec}(\boldsymbol{E}^\top)$ is also permutation regular and therefore regular as the transpose of a permutation matrix is also a permutation matrix.

Definition 5 (Permutation Syndrome Decoding). *Let $k, n \in \mathbb{N}$ be such that $k \leq n$ and n is a perfect square. The* Permutation Syndrome Decoding (PSD) problem *is to find a permutation regular vector $\boldsymbol{e} \in \mathbb{F}_2^n$ such that $\boldsymbol{H}\boldsymbol{e} = \boldsymbol{s}$, given a parity-check matrix $\boldsymbol{H} \in \mathbb{F}_2^{(n-k)\times n}$ and a syndrome vector $\boldsymbol{s} \in \mathbb{F}_2^{n-k}$.*

We write $\mathsf{PSD}(n, k)$ to denote the PSD problem with parameters n and k. A random $\mathsf{PSD}(n, k)$ instance is sampled as follows:

i) sample $\boldsymbol{H}$ uniformly at random among all full-rank matrices in $\mathbb{F}_2^{(n-k)\times n}$,

ii) sample $\boldsymbol{P} \in \mathsf{Perm}_n(\mathbb{F}_2)$ uniformly at random and return $(\boldsymbol{H}, \boldsymbol{s} = \boldsymbol{H} \cdot \mathsf{vec}(\boldsymbol{P}))$.

The expected number of solutions to a PSD instance is given by

$$S_{n,k}^{(\mathsf{PSD})} \coloneqq 1 + (\sqrt{n})!\, 2^{-(n-k)},$$

where $\sqrt{n}!$ counts the number of permutation regular vectors of length n, or equivalently, the number of elements in $\mathsf{Perm}_{\sqrt{n}}(\mathbb{F}_2)$, and $2^{-(n-k)}$ is the probability that a fixed permutation regular vector $\boldsymbol{e}$ satisfies the linear system $\boldsymbol{H}\boldsymbol{e} = \boldsymbol{s}$. Throughout, we assume that at least one solution always exists.

Theorem 1 ($\mathsf{PSD} \leq \mathsf{RSD}$). *Let n be a perfect square and set $w \coloneqq \sqrt{n}$. Assume*

$$\binom{n}{w}^w 2^{-(n-k)} = \mathsf{negl}(n).$$

Then $\mathsf{PSD}(n,k) \leq \mathsf{RSD}(n,k,w)$.

Proof. Let $\mathcal{A}^{(\mathsf{RSD})}$ be an oracle that solves $\mathsf{RSD}(n,k,w)$. We define an algorithm $\mathcal{A}^{(\mathsf{PSD})}$ that, on input a random $\mathsf{PSD}(n,k)$ instance $(\boldsymbol{H}, \boldsymbol{s})$, runs $\mathcal{A}^{(\mathsf{RSD})}$ on $(\boldsymbol{H}, \boldsymbol{s})$ and outputs the result. Let $\boldsymbol{e}$ denote the permutation regular solution of the PSD instance. Since $\boldsymbol{e}$ is regular of weight w, it is a valid solution to the corresponding $\mathsf{RSD}(n,k,w)$ instance. Notice that, as the $\mathsf{PSD}(n,k)$ instance is random, the parity-check matrix $\boldsymbol{H}$ is also random, and so $(\boldsymbol{H}, \boldsymbol{s})$ seen as a $\mathsf{RSD}(n,k,\sqrt{n})$ instance expects $\binom{n}{w}^w 2^{-(n-k)}$ solutions (see Equation (2)) that are $\mathsf{negl}(n)$ by assumption. Hence, with overwhelming probability, $\boldsymbol{e}$ is the unique regular solution and $\mathcal{A}^{(\mathsf{RSD})}$ returns $\boldsymbol{e}$. Therefore, $\mathcal{A}^{(\mathsf{PSD})}$ solves $\mathsf{PSD}(n,k)$. □

Encoding the permutation structure. Let $\boldsymbol{e} \in \mathbb{F}_2^n$ be a permutation regular vector and let us denote $w \coloneqq \sqrt{n}$. Then, we have that the sum of the elements along each row and each column of $\mathsf{vec}^{-1}(\boldsymbol{e})$ equals 1. Encoding these constraints into linear equations we obtain the following

$$[\mathbf{1}_w^\top \otimes \boldsymbol{I}_w]\boldsymbol{e} = \mathbf{1}_w \qquad \text{and} \qquad [\boldsymbol{I}_w \otimes \mathbf{1}_w^\top]\boldsymbol{e} = \mathbf{1}_w, \tag{3}$$

where $\mathbf{1}_w \in \mathbb{F}_2^w$ denotes the column vector whose entries are all 1. Let us define the matrix

$$\boldsymbol{K} \coloneqq \begin{bmatrix} \boldsymbol{I}_w \otimes \mathbf{1}_w^\top \\ \mathbf{1}_w^\top \otimes \boldsymbol{I}_w \end{bmatrix} \in \mathbb{F}_2^{2w \times n}.$$

The matrix $\boldsymbol{K}$ has rank $2w-1$. Indeed, among the $2w$ linear constraints, exactly one is redundant, since the sum of the row-sum equations equals the sum of the column-sum equations. Therefore, the system $\boldsymbol{K}\boldsymbol{x} = \mathbf{1}_{2w}$ provides $2w-1$ additional equations that are, for a random $\mathsf{PSD}(n,k)$ instance, generally linearly independent with the system $\boldsymbol{H}\boldsymbol{x} = \boldsymbol{s}$.

Proposition 2. *If a regular vector $\boldsymbol{e} \in \mathbb{F}_2^n$ is such that $\boldsymbol{K}\boldsymbol{e} = \mathbf{1}_{2w}$, for $w = \sqrt{n}$, then it is also permutation regular.*

Proof. Let $\boldsymbol{E} := \mathsf{vec}^{-1}(\boldsymbol{e})$. Then each column of $\boldsymbol{E}$ has Hamming weight 1. The second block of Equation (3) is equivalent to

$$(\mathbf{1}_w^\top \otimes \boldsymbol{I}_w)\boldsymbol{e} = \mathbf{1}_w \qquad \Longrightarrow \qquad \boldsymbol{E}^\top \cdot \mathbf{1}_w = \mathbf{1}_w,$$

so each column sum of $\boldsymbol{E}$ equals 1, i.e., each column has odd weight. Since $\boldsymbol{E}$ has exactly w ones in total (one per row by regularity) and the w columns each contain at least one 1 (because their sums equal 1), it follows that every column contains exactly one 1. Therefore, $\boldsymbol{E} \in \mathsf{Perm}_w(\mathbb{F}_2)$ and $\boldsymbol{e}$ is permutation regular. □

As an alternative to the approach followed in the proof of Theorem 1, one may append the permutation equations $\boldsymbol{K}\boldsymbol{x} = \mathbf{1}_{2w}$ to the syndrome equation. In this case, any regular solution to the augmented system is permutation regular by Proposition 2. This gives a direct way to enforce the permutation structure, at the cost of changing the effective dimension of the underlying RSD instance.

Polynomial-time parameters regime. The following proposition follows from [28, Thm. 3.1] and identifies a range of parameters for which PSD is solvable in polynomial time.

Proposition 3. *Let n, k be parameters such that $k < n$ and n is a perfect square. If $2\sqrt{n}-1 \geq k$, then a random $\mathsf{PSD}(n,k)$ instance is solvable in polynomial time with high probability.*

Proof. Consider the linear system obtained by appending the permutation constraints to the equation $\boldsymbol{H}\boldsymbol{x} = \boldsymbol{s}$:

$$\underbrace{\begin{bmatrix}\boldsymbol{H}\\ \boldsymbol{K}\end{bmatrix}}_{\boldsymbol{H}'}\boldsymbol{x} = \underbrace{\begin{bmatrix}\boldsymbol{s}\\ \mathbf{1}_{2\sqrt{n}}\end{bmatrix}}_{\boldsymbol{s}'}. \tag{4}$$

The system $\boldsymbol{H}'\boldsymbol{x} = \boldsymbol{s}'$ has affine solution space of dimension $\epsilon = n - \mathsf{rank}(\boldsymbol{H}')$. If $\epsilon = \mathcal{O}(\log n)$, then the instance can be solved in polynomial time by enumerating all $2^\epsilon = \mathrm{poly}(n)$ solutions until one finds one that is permutation regular. Since $\mathsf{rank}(\boldsymbol{K}) = 2\sqrt{n}-1$, the remaining rank to be contributed by the random matrix $\boldsymbol{H}$ is at most $n - (2\sqrt{n}-1)$. In the setting of $2\sqrt{n}-1 \geq k$, we have $n-k \geq n-(2\sqrt{n}-1)$, and standard rank estimates for random binary matrices imply that $\mathsf{rank}(\boldsymbol{H}') = n - \epsilon$ with $\epsilon = \mathcal{O}(\log n)$ with high probability. Therefore the above enumeration runs in polynomial time with high probability. □

4 From PKP/PEP over Binary Field Extensions to PSD

We describe how to convert a PKP or a PEP instance over binary field extensions into a structured instance of PSD. For simplicity, we give a description for PKP first, and summarize later how this extends to PEP. Moreover, in this section we

focus on binary field extensions. While the transformation easily generalizes to non-binary field extensions, some of the results presented here do not directly apply.

Let us consider a $\mathsf{PKP}(n, k, \ell, q = 2^\nu)$ instance with full-rank public matrices $\boldsymbol{H} \in \mathbb{F}_{2^\nu}^{(n-k)\times n}$ and $\boldsymbol{V} \in \mathbb{F}_{2^\nu}^{n\times \ell}$, and a secret permutation matrix $\boldsymbol{E} \in \mathsf{Perm}_n(\mathbb{F}_2)$ such that

$$\boldsymbol{HEV} = \boldsymbol{0}.$$

Notice that we consider $\boldsymbol{E}$ over $\mathbb{F}_2$ rather than $\mathbb{F}_{2^\nu}$, since its entries are binary. By the property stated in Equation (1), we have

$$[\boldsymbol{V}^\top \otimes \boldsymbol{H}]\mathsf{vec}(\boldsymbol{E}) = \boldsymbol{0}.$$

Let us denote

$$\tilde{\boldsymbol{H}} = [\boldsymbol{V}^\top \otimes \boldsymbol{H}] \in \mathbb{F}_{2^\nu}^{(n-k)\ell\times n^2} \quad \text{and} \quad \boldsymbol{e} = \mathsf{vec}(\boldsymbol{E}) \in \mathbb{F}_2^{n^2}.$$

It is well known that $\mathbb{F}_{2^\nu}$ can be viewed as a vector space of dimension ν over $\mathbb{F}_2$. Hence, there exist ν linearly independent $\mathbb{F}_2$-linear projections[1] $\varphi_i \colon \mathbb{F}_{2^\nu} \to \mathbb{F}_2$, for $i \in \{0, \dots, \nu - 1\}$, such that the matrix

$$\bar{\boldsymbol{H}} \coloneqq \begin{bmatrix} \varphi_0(\tilde{\boldsymbol{H}}) \\ \vdots \\ \varphi_{\nu-1}(\tilde{\boldsymbol{H}}) \end{bmatrix} \in \mathbb{F}_2^{(n-k)\ell\nu\times n^2} \tag{5}$$

has rank at most $\min\{(n-k)\ell\nu,\ n^2\}$. More precisely, since $\mathsf{rank}_{\mathbb{F}_{2^\nu}}(\tilde{\boldsymbol{H}}) = \mathsf{rank}_{\mathbb{F}_{2^\nu}}(\boldsymbol{V})\mathsf{rank}_{\mathbb{F}_{2^\nu}}(\boldsymbol{H}) = (n-k)\ell$, it follows that

$$\mathsf{rank}_{\mathbb{F}_2}(\bar{\boldsymbol{H}}) \le \nu \cdot \mathsf{rank}_{\mathbb{F}_{2^\nu}}(\tilde{\boldsymbol{H}}) = (n-k)\ell\nu.$$

Since $\boldsymbol{e} \in \mathbb{F}_2^{n^2}$, we have

$$\boldsymbol{0} = \varphi_i(\tilde{\boldsymbol{H}}\boldsymbol{e}) = \varphi_i(\tilde{\boldsymbol{H}})\boldsymbol{e}, \quad \text{for all } i \in \{0, \dots, \nu - 1\}.$$

Stacking these equations yields

$$\bar{\boldsymbol{H}}\boldsymbol{e} = \boldsymbol{0}. \tag{6}$$

One may see Equation (6) as an instance of $\mathsf{PSD}(n^2, n^2-(n-k)\ell\nu)$. However, this whole transformation does not yield a generic reduction from PKP to PSD, but rather to PSD instances drawn from a specific, structured distribution induced by the PKP instance. The following proposition shows that solving the derived PSD instance also solves the original PKP instance.

[1] For example, after fixing an $\mathbb{F}_2$-basis $\{\beta_0, \dots, \beta_{\nu-1}\}$ of $\mathbb{F}_{2^\nu}$, one may define the canonical projections $\varphi_i \colon \mathbb{F}_{2^\nu} \to \mathbb{F}_2$ by writing $x = \sum_{j=0}^{\nu-1} x_j\beta_j$ with $x_j \in \mathbb{F}_2$ and setting $\varphi_i(x) \coloneqq x_i$.

Proposition 4. *Let* $(\bar{\boldsymbol{H}}, \boldsymbol{0})$ *be the* $\mathsf{PSD}(n^2, n^2 - (n-k)\ell\nu)$ *instance derived from a* $\mathsf{PKP}(n, k, \ell, 2^\nu)$ *instance* $(\boldsymbol{H}, \boldsymbol{V})$*. If a permutation regular vector* $\boldsymbol{e}$ *is such that* $\bar{\boldsymbol{H}}\boldsymbol{e} = \boldsymbol{0}$*, then the permutation* $\boldsymbol{E} = \mathsf{vec}^{-1}(\boldsymbol{e}) \in \mathsf{Perm}_n(\mathbb{F}_2)$ *is such that* $\boldsymbol{HEV} = \boldsymbol{0}$*.*

Proof. The proof follows by inverting the above procedure. First, we have that $\bar{\boldsymbol{H}}\boldsymbol{e} = \boldsymbol{0} \Rightarrow \varphi_i(\tilde{\boldsymbol{H}})\boldsymbol{e} = \boldsymbol{0}$, for $i = 0, \ldots, \nu - 1$, and then $\tilde{\boldsymbol{H}}\boldsymbol{e} = \boldsymbol{0}$. Recalling that $\tilde{\boldsymbol{H}} = \boldsymbol{V}^\top \otimes \boldsymbol{H}$, we have $\boldsymbol{0} = \tilde{\boldsymbol{H}}\boldsymbol{e} = [\boldsymbol{V}^\top \otimes \boldsymbol{H}]\boldsymbol{e} = \boldsymbol{H}\mathsf{vec}^{-1}(\boldsymbol{e})\boldsymbol{V} = \boldsymbol{HEV}$. □

Extension to PEP. The above transformation naturally extends to PEP. Let $\boldsymbol{H}' \in \mathbb{F}_{2^\nu}^{(n-k)\times n}$ and $\boldsymbol{G} \in \mathbb{F}_{2^\nu}^{k\times n}$ be two public matrices of a $\mathsf{PEP}(n, k, 2^\nu)$ instance. By considering $\ell = k$ in the above transformation, one derives a $\mathsf{PSD}(n^2, n^2 - (n-k)k\nu)$ instance whose parity-check matrix $\bar{\boldsymbol{H}}$ is built analogously as above: first define $\tilde{\boldsymbol{H}} = [\boldsymbol{G} \otimes \boldsymbol{H}'] \in \mathbb{F}_{2^\nu}^{(n-k)k\times n^2}$, then apply to it ν independent $\mathbb{F}_2$-linear projections to obtain $\bar{\boldsymbol{H}} \in \mathbb{F}_2^{(n-k)k\nu\times n^2}$.

4.1 A New Polynomial-time Parameter Regime for PKP

We use the above transformation together with Proposition 3 and the following assumption to derive a new regime of parameters for which PKP is solvable in polynomial time.

Assumption 1 *Let* $\bar{\boldsymbol{H}}$ *be the matrix defined in Equation* (5) *and derived from a random* $\mathsf{PKP}(n, k, \ell, 2^\nu)$ *instance. Let* $\boldsymbol{K} \in \mathbb{F}_2^{2n\times n^2}$ *denote the matrix encoding the permutation equations of order* n *(see Equation* (3)*). Then,*

$$\mathsf{rank}\begin{bmatrix}\bar{\boldsymbol{H}}\\ \boldsymbol{K}_n\end{bmatrix} \geq \min\{n^2,\ (n-k)\ell\nu + 2n - 1\} - \epsilon,$$

for some $\epsilon = \mathcal{O}(\log n)$*, with high probability.*

With the above, we assume that the structure inherited in $\bar{\boldsymbol{H}}$ does not expect any more linear dependence with $\boldsymbol{K}_n$ than a random matrix of the same dimension would do. We verify that this assumption holds experimentally and comment on the results in Appendix A.

Theorem 2. *Let* $\nu, n, k, \ell \in \mathbb{N}$ *be parameters such that* $\ell < k < n$*. Under Assumption 1, if*

$$(n-1)^2 \leq (n-k)\ell\nu, \tag{7}$$

then a random $\mathsf{PKP}(n, k, \ell, 2^\nu)$ *instance is solvable in polynomial time with high probability.*

Proof. We showed how to derive a $\mathsf{PSD}(n', k')$ instance from $\mathsf{PKP}(n, k, \ell, 2^\nu)$, with $n' = n^2$ and $k' = n^2 - (n-k)\ell\nu$. Then, Proposition 3 applies when $2\sqrt{n'} - 1 \geq k'$, that is, when $2n - 1 \geq n^2 - (n-k)\ell\nu$, or equivalently when $(n-k)\ell\nu \geq (n-1)^2$. Under Assumption 1, the augmented matrix obtained by adding the permutation

constraints to the parity-check matrix of the derived PSD instance has rank at least $n' - \varepsilon$, for some $\varepsilon = \mathcal{O}(\log n')$, with high probability. Hence the affine solution space has size at most $2^{\varepsilon} = \text{poly}(n')$. We can therefore enumerate all its solutions and test which ones are permutation regular. Since any permutation regular solution of the derived PSD instance yields a solution to the original PKP instance (Proposition 4), the claim follows. □

Theorem 2 shows that the complexity of PKP over binary field extensions is strictly related to the power ν. Indeed, for n, k, ℓ fixed, the problem falls into this polynomial-time regime when $\nu \geq \frac{(n-1)^2}{(n-k)\ell}$. For example, for PERK 2.1 parameter sets corresponding to NIST level I, that is $(n, k, \ell) = (64, 37, 1)$, the PKP instance becomes polynomial starting from $\nu = 147$. Nevertheless, PERK uses $\nu = 11$ and is therefore far from this regime. We implemented a PKP solving algorithm following the strategy of Theorem 2's proof in SageMath that is available at [23].

4.2 A New Polynomial-time Parameter Regime for PEP

Theorem 2 naturally extends to the case of PEP in the case of random codes over $\mathbb{F}_{2^\nu}$ (i.e., with high probability, with trivial hull). Specifically, setting $\ell = k$, one would get a polynomial-time regime whenever $\nu \geq \frac{(n-1)^2}{(n-k)k}$. Nevertheless, we must consider the cryptographically interesting setting of PEP instantiations with self-dual and self-orthogonal codes. In these settings, as we will see in the following two propositions, the matrix of coefficients $\bar{\boldsymbol{H}}$, constructed as in Equation (5), does not behave as in Assumption 1, that is, as if it was sampled uniformly at random.

Proposition 5. *Let $\boldsymbol{G} \in \mathbb{F}_{2^\nu}^{k \times n}$ and $\boldsymbol{H}' \in \mathbb{F}_{2^\nu}^{(n-k) \times n}$ be the input matrices of a $\mathsf{PEP}(n, k, 2^\nu)$ instance with $k < n/2$ such that the code generated by $\boldsymbol{G}$ is self-orthogonal. Then, the matrix $\bar{\boldsymbol{H}}$ defined as*

$$\bar{\boldsymbol{H}} \coloneqq \begin{bmatrix} \varphi_0(\boldsymbol{G} \otimes \boldsymbol{H}') \\ \vdots \\ \varphi_{\nu-1}(\boldsymbol{G} \otimes \boldsymbol{H}') \end{bmatrix} \in \mathbb{F}_2^{(n-k)k\nu \times n^2}, \tag{8}$$

is such that $\mathsf{rank}(\bar{\boldsymbol{H}}) \leq n^2 - n - 1$, where $\varphi_i : \mathbb{F}_{2^\nu} \to \mathbb{F}_2$, for $i = 0, \ldots, \nu - 1$, are ν linearly independent $\mathbb{F}_2$-linear projections.

Proof. Notice that every codeword of a self-orthogonal code $\mathcal{C}$ is orthogonal to every codeword of $\mathcal{C}$, and, since the field characteristic is 2, they all have coordinate sum equal to zero. Thus, we have that $\boldsymbol{G}\mathbf{1}_n = \mathbf{0}$, where $\mathbf{1}_n$ is the n-column vector of all 1s. Moreover, for any $\boldsymbol{u} \in \mathbb{F}_2^n$, we have that

$$\boldsymbol{H}'(\boldsymbol{u}\mathbf{1}_n^\top)\boldsymbol{G}^\top = \boldsymbol{H}'\boldsymbol{u}(\boldsymbol{G}\mathbf{1}_n)^\top = \mathbf{0}.$$

Therefore, every vector of the form $\mathsf{vec}(\boldsymbol{u}\mathbf{1}_n^\top) \in \mathbb{F}_2^{n^2}$, where $\boldsymbol{u}\mathbf{1}_n^\top \in \mathbb{F}_2^{n \times n}$, belongs to the right-kernel of $\boldsymbol{G} \otimes \boldsymbol{H}' \in \mathbb{F}_{2^\nu}^{k(n-k) \times n^2}$ (see Equation (1)), and, as $\boldsymbol{u}$ ranges

over $\mathbb{F}_2^n$, these form an n-dimensional subspace of $\mathbb{F}_2^{n^2}$. Furthermore, this space does not contain the permutation regular vector derived from the permutation solution to the original PEP instance, also belonging to the right-kernel of $\boldsymbol{G}\otimes\boldsymbol{H}'$. Hence, the matrix $\boldsymbol{G}\otimes\boldsymbol{H}'$ has a right-kernel W of dimension at least $n+1$. When applying the ν $\mathbb{F}_2$-linear projections to construct $\bar{\boldsymbol{H}}$, the vectors in W remain in the kernel since $W \subset \mathbb{F}_2^{n^2}$. Hence, $\mathsf{rank}(\bar{\boldsymbol{H}}) \leq n^2 - n - 1$. □

Proposition 6. *Let $\boldsymbol{G} \in \mathbb{F}_{2^\nu}^{k\times n}$ and $\boldsymbol{H}' \in \mathbb{F}_{2^\nu}^{(n-k)\times n}$ be the input matrices of a $\mathsf{PEP}(n,k,2^\nu)$ instance with $k = n/2$ such that the code generated by $\boldsymbol{G}$ is self-dual. Then, the matrix $\bar{\boldsymbol{H}} \in \mathbb{F}_2^{(n-k)k\nu\times n^2}$ defined as in Equation (8) is such that $\mathsf{rank}(\bar{\boldsymbol{H}}) \leq n^2 - 2n$.*

Proof. We follow a similar argument as in the proof of Proposition 5. If the code generated by $\boldsymbol{G}$ is self-dual, then also the code generated by $\boldsymbol{H}'$ is self-dual as it is the parity-check of an equivalent that is also self-dual. Hence $\boldsymbol{G}\mathbf{1}_n = \mathbf{0}$ and $\boldsymbol{H}'\mathbf{1}_n = \mathbf{0}$. Let $\boldsymbol{u}, \boldsymbol{v} \in \mathbb{F}_2^n$. Then

$$\boldsymbol{H}'(\mathbf{1}_n\boldsymbol{v}^\top + \boldsymbol{u}\mathbf{1}_n^\top)\boldsymbol{G}^\top = (\boldsymbol{H}'\mathbf{1}_n)(\boldsymbol{v}^\top\boldsymbol{G}^\top) + (\boldsymbol{H}'\boldsymbol{u})(\boldsymbol{G}\mathbf{1}_n)^\top = \mathbf{0}.$$

Therefore, every vector $\mathsf{vec}(\mathbf{1}_n\boldsymbol{v}^\top + \boldsymbol{u}\mathbf{1}_n^\top) \in \mathbb{F}_2^{n^2}$ belongs to the right-kernel of $\boldsymbol{G}\otimes\boldsymbol{H}'$. These vectors generate a space of dimension $2n-1$, since the matrix $\mathbf{1}_n\mathbf{1}_n^\top$ is generated both by $\mathbf{1}_n\boldsymbol{v}^\top$ and by $\boldsymbol{u}\mathbf{1}_n^\top$. Moreover, the right-kernel of $\boldsymbol{G}\otimes\boldsymbol{H}'$ also contains the secret permutation vector coming from the original PEP instance. For $n > 2$, this vector is not of the form $\mathsf{vec}(\mathbf{1}_n\boldsymbol{v}^\top + \boldsymbol{u}\mathbf{1}_n^\top)$, and therefore the right-kernel contains a subspace of dimension $2n$. With analogous arguments to those in the proof of Proposition 5, we get that $\mathsf{rank}(\bar{\boldsymbol{H}}) \leq n^2 - 2n$. □

With the following assumption, we claim that, whenever $\bar{\boldsymbol{H}}$ has more rows than columns, there are no other structural vectors in the right-kernel of $\bar{\boldsymbol{H}}$ other than those uncovered in Propositions 5 and 6. We explicitly set $\nu > 4$ to avoid degenerate cases.

Assumption 2 *Let $\nu > 4$ and let $\bar{\boldsymbol{H}}$ be the matrix defined as in Equation (8) and derived from a random $\mathsf{PEP}(n,k,2^\nu)$ instance with input code self-orthogonal and $k \leq n/2$. Assume that $(n-k)k\nu > n^2$. Then, with high probability,*

$$\mathsf{rank}(\bar{\boldsymbol{H}}) \geq \min\{(n-k)k\nu,\ n^2 - d\} - \epsilon,$$

for some $\epsilon = \mathcal{O}(\log n)$, where $d = n+1$ if the input code is self-orthogonal with $k < n/2$, and $d = 2n$ if the input code is self-dual, i.e., $k = n/2$.

Experimentally, we validate that Assumption 2 holds and reports the results in Appendix A.

Theorem 3. *Let $\nu > 4, n, k \in \mathbb{N}$ be parameters, and let d be as in Assumption 2. Assume that*

$$(n-k)k\nu \geq n^2 - d. \tag{9}$$

Then, under Assumption 2, $\mathsf{PEP}(n,k,2^\nu)$ with either self-orthogonal or self dual random input codes can be solved in polynomial time with high probability.

Proof. The solution to the PEP input instance is a permutation corresponding to a permutation regular vector that is itself solution to the system $\bar{\boldsymbol{H}}\boldsymbol{x} = \boldsymbol{0}$.

If the input code is self-orthogonal with $k < n/2$, then $d = n+1$. By Assumption 2, the kernel of $\bar{\boldsymbol{H}}$ has dimension at most $n + 1 + \epsilon$, for some $\epsilon = \mathcal{O}(\log n)$. The structural part of this kernel is the $(n + 1)$-dimensional space described in Proposition 5. We fix one row of the secret permutation and try all n possible positions for its nonzero entry. For the correct guess, the corresponding row and column are fully determined, and we obtain a reduced non-homogeneous system $\hat{\boldsymbol{H}}\boldsymbol{x} = \boldsymbol{b}$ with $n^2 - 2n + 1$ variables. Without loss of generality, assume that the secret permutation $\boldsymbol{P} = \{p_{i,j}\}$ is such that $p_{0,0} = 1$, and therefore that the correct guess is the entry $(0, 0)$. The solutions described in Proposition 5 have the form $X = \alpha\boldsymbol{P} + \boldsymbol{u}\boldsymbol{1}_n^\top$, where we view solution vectors in matrix form via vec^{-1}. The guess imposes $x_{0,0} = 1$, $x_{0,c} = 0$ for $c > 0$, and $x_{r,0} = 0$ for $r > 0$. Since $p_{0,c} = 0$ for $c > 0$ and $p_{r,0} = 0$ for $r > 0$, we have $x_{0,c} = u_0$ for $c > 0$ and $x_{r,0} = u_r$ for $r > 0$. Hence, we must have $\boldsymbol{u} = \boldsymbol{0}$ and $\alpha = 1$, and the only solution of this form compatible with the correct guess is the secret permutation $\boldsymbol{P}$ itself. Thus, after the correct guess, the corresponding homogenized solution space has dimension at most $1 + \epsilon$.

Finally, if the input code is self-dual, then $d = 2n$. By Assumption 2, the kernel has dimension at most $2n+\epsilon$, where a subspace of dimension $2n-1$ is given by the structural solutions described in the proof of Proposition 6. As above, we fix one row and try all n possible guesses for a non-zero entry. For the correct guess, the corresponding row and column are fully determined, and we obtain a reduced non-homogeneous system $\hat{\boldsymbol{H}}\boldsymbol{x} = \boldsymbol{b}$ with $n^2 - 2n + 1$ variables. Again, without loss of generality, we assume that $p_{0,0} = 1$. The solutions described in Proposition 6 have matrix form via vec^{-1} as $\boldsymbol{X} = \alpha\boldsymbol{P} + \boldsymbol{1}_n\boldsymbol{v}^\top + \boldsymbol{u}\boldsymbol{1}_n^\top$. The guess imposes $x_{0,0} = 1$, $x_{0,c} = 0$ for $c > 0$, and $x_{r,0} = 0$ for $r > 0$. Since $p_{0,c} = 0$ for $c > 0$ and $p_{r,0} = 0$ for $r > 0$, we have $x_{0,c} = v_c + u_0$ for $c > 0$ and $x_{r,0} = v_0 + u_r$ for $r > 0$. Hence, $v_c = u_0$ for $c > 0$, $u_r = v_0$ for $r > 0$, and the condition $x_{0,0} = 1$ gives $\alpha + v_0 + u_0 = 1$. Therefore, for $r, c > 0$, we have $x_{r,c} = \alpha p_{r,c} + 1 + \alpha$. If $\alpha = 1$, this is the restriction of the secret permutation $\boldsymbol{P}$ to the remaining variables, while if $\alpha = 0$, this is the all-one vector on the remaining variables. Thus, after the correct guess, the corresponding homogenized solution space has dimension at most $2 + \epsilon$.

Overall, the time complexity of solving PEP in both above cases is polynomial in n. In fact, this is dominated by making n guesses, solving a linear system, and then enumerating a solution space of dimension at most $2 + \epsilon$, for each guess, which takes $2^{2+\epsilon}$, with $\epsilon = \mathcal{O}(\log n)$. □

In particular, Equation (9) always holds for self-dual codes and $\nu > 4$. Hence, Theorem 3 implies that PEP with self-dual codes over $\mathbb{F}_{2^\nu}$ is solvable in polynomial-time whenever $\nu \geq 5$. This is, to the best of our knowledge, a novel result in the literature of PEP cryptanalysis. We implemented and successfully tested a solver for PEP over binary extension fields in SageMath following the strategy of Theorem 3 that is available at [23].

5 Permutation-Based Regular-ISD for PSD and Applications to PKP

In this section, we adapt the permutation-based regular-ISD algorithm of Esser and Santini [28], itself a variant of Prange's original ISD algorithm for solving RSD, to the PSD setting. We focus on this algorithm because its simplicity makes it particularly suitable for isolating and studying the effect of the additional permutation structure of the solution vector. Adapting the other more advanced algorithms proposed by Esser and Santini is left for future work. We also investigated an adaptation of the algebraic attack of [19] for RSD to PSD, building on the techniques of [48]. Since this approach did not yield competitive results, we omit it for reasons of space and simplicity.

5.1 Algorithm

The core of our adaptation is a new family of permutations tailored to the bijective structure of permutation-regular solutions, which increases the probability of sampling an error-free information set. We further show that the chosen family maximizes this probability among all regular permutation families.

Setup. Let $(\boldsymbol{H}, \boldsymbol{s})$ be a $\mathsf{PSD}(n, k)$ instance and write $w \coloneqq \sqrt{n}$ and $[w] \coloneqq \{1, \dots, w\}$. We append the permutation equations to the syndrome equation as in Equation (4) and discard one redundant row of $\boldsymbol{K}$ (recall $\mathsf{rank}(\boldsymbol{K}) = 2w - 1$), obtaining a reduced system $\boldsymbol{H}'\boldsymbol{x} = \boldsymbol{s}'$, with $\boldsymbol{H}' \in \mathbb{F}_2^{(n-k')\times n}$ and $\boldsymbol{s}' \in \mathbb{F}_2^{n-k'}$, where

$$k' \coloneqq k - 2w + 1$$

is the effective code dimension of the augmented instance.

Recall that a solution $\boldsymbol{e} \in \mathbb{F}_2^n$ is permutation regular, i.e., $\boldsymbol{E} \coloneqq \mathsf{vec}^{-1}(\boldsymbol{e}) \in \mathsf{Perm}_w(\mathbb{F}_2)$. We write $\boldsymbol{e} = (\boldsymbol{e}_1, \dots, \boldsymbol{e}_w)$ with blocks $\boldsymbol{e}_i \in \mathbb{F}_2^w$ being the columns of $\boldsymbol{E}$, and, for $j \in [w]$, we call the j-th coordinate of $\boldsymbol{e}_i$, i.e., the global coordinate $(i-1)w+j$ of $\boldsymbol{e}$, *position j within block i*. Let $\pi(i) \in [w]$ denote the within-block position of the unique nonzero entry of $\boldsymbol{e}_i$. Since $\boldsymbol{E}$ is a permutation matrix, the map $i \mapsto \pi(i)$ is a permutation of $[w]$, which is the essential difference to standard RSD.

Basic idea. As in any Prange-style ISD algorithm, we search for an information set that contains only zero entries of $\boldsymbol{e}$. For a permutation matrix $\boldsymbol{P} \in \mathsf{Perm}_n(\mathbb{F}_2)$ write

$$\boldsymbol{H}'\boldsymbol{P}^{-1} = (\boldsymbol{H}_1 \mid \boldsymbol{H}_2) \qquad \text{and} \qquad \boldsymbol{P}\boldsymbol{e} = (\bar{\boldsymbol{e}}, \boldsymbol{e}_J)^\top,$$

with $\boldsymbol{H}_1 \in \mathbb{F}_2^{(n-k')\times(n-k')}$ and $\boldsymbol{e}_J \in \mathbb{F}_2^{k'}$ being the part of $\boldsymbol{e}$ moved to the last k' coordinates, the candidate information set. If $\boldsymbol{e}_J = \boldsymbol{0}$, the syndrome equation collapses to $\boldsymbol{s}' = \boldsymbol{H}'\boldsymbol{e} = \boldsymbol{H}_1\bar{\boldsymbol{e}}$, and if additionally $\boldsymbol{H}_1$ is invertible, we recover

$$\bar{\boldsymbol{e}} = \boldsymbol{H}_1^{-1}\boldsymbol{s}' \qquad \text{and} \qquad \boldsymbol{e} = \boldsymbol{P}^{-1}(\bar{\boldsymbol{e}}, \boldsymbol{0}_{k'})^\top.$$

For now let us assume $\boldsymbol{H}_1$ is non-singular. We now restrict to permutations that select from *every* block the same number $v \coloneqq k'/w$ of coordinates for the information set, which requires $w \mid k'$. Each block then keeps $r \coloneqq w - v$ coordinates outside the information set, and in total $wr = n - wv = n - k'$ coordinates remain, matching the number of rows of $\boldsymbol{H}'$. In the terminology of [28, Def. 4.1], these permutations are called v-regular. We now further restrict sampling to a subset of all v-regular permutations.

A permutation family with improved success probability. The condition $\boldsymbol{e}_J = \mathbf{0}$ holds if and only if, for every block i, the r positions kept outside the information set per block contain the position $\pi(i)$. We now let blocks share their kept positions: we partition the blocks into groups of size r and assign to each group a set of r positions that remain outside the information set. For now we fix the block partition as the canonical consecutive grouping of r blocks, i.e.,

$$G_a \coloneqq \{(a-1)r + 1, \ldots, ar\}, \qquad a \in [m].$$

Further, for now we assume $r \mid w$ and set $m \coloneqq w/r$ to obtain evenly sized partitions. With this grouping, $\boldsymbol{e}_J = \mathbf{0}$ means that, for each group G_a, the one-positions $\{\pi(i) : i \in G_a\}$ of the blocks in that group are included in the group's assigned kept-position set, i.e., the positions outside the information set, which we call Q_a in the following. As we will see in the following, since $|Q_a| = |G_a|$, we obtain $\boldsymbol{e}_J = \mathbf{0} \quad \Leftrightarrow \quad Q_a = \pi(G_a)$.

Note that in standard Prange a fixed error admits many good information sets, namely any set avoiding the error's w nonzero positions leads to success. Restricting to permutations built from a block partition $G_1, \ldots, G_m$ and a position partition $Q_1, \ldots, Q_m$ narrows this down sharply: a fixed solution is compatible with a unique choice of position partition $(Q_1, \ldots, Q_m) = \pi(G_1), \ldots, \pi(G_m)$, so the algorithm is effectively guessing the images $\pi(G_i)$ for $i \in [m]$. While such a collapse to an enumeration strategy might sound counterintuitive, we show in the following that the corresponding reduction of the search space still in total leads to an exponential time improvement.

Let us formally define the permutation family we restrict to.

Definition 6 (partition-induced permutation). *Let $r \mid w$ and $m = w/r$. We denote by $\mathcal{Q}_{w,r}$ the set of all ordered partitions $Q = (Q_1, \ldots, Q_m)$ of $[w]$ into m sets of size r. For $Q \in \mathcal{Q}_{w,r}$, the* permutation induced by Q *is the unique permutation matrix $\boldsymbol{P}_Q \in \mathsf{Perm}_n(\mathbb{F}_2)$ that maps the coordinates*

$$T_Q \coloneqq \{(i-1)w + j \ : \ a \in [m],\ i \in G_a,\ j \in Q_a\}$$

to the first $n - k'$ positions and all remaining coordinates to the last k' positions, in both cases preserving their relative order.

In words, every block i in group G_a keeps exactly the within-block positions Q_a outside the information set, while its remaining v positions are moved into the information set.

The following lemma formalizes this relation between PSD solutions and partition induced permutations and contains the exact counting on which our analysis rests.

Lemma 1. *Let $\boldsymbol{e} \in \mathbb{F}_2^n$ be permutation regular with associated permutation π, and let $Q \in \mathcal{Q}_{w,r}$. Then Q is* compatible *with $\boldsymbol{e}$, i.e., $\boldsymbol{P}_Q \boldsymbol{e} = (\bar{\boldsymbol{e}}, \mathbf{0}_{k'})^\top$ for some $\bar{\boldsymbol{e}} \in \mathbb{F}_2^{n-k'}$, if and only if $\pi(G_a) = Q_a$ for all $a \in [m]$. In particular,*

1. *every permutation regular vector $\boldsymbol{e}$ is compatible with exactly one $Q \in \mathcal{Q}_{w,r}$, namely $Q_a = \pi(G_a)$,*
2. *every fixed $Q \in \mathcal{Q}_{w,r}$ is compatible with exactly $(r!)^m$ permutation regular vectors, and*
3. *$|\mathcal{Q}_{w,r}| = \binom{w}{r,\ldots,r} = \frac{w!}{(r!)^m}$.*

Proof. Fix $a \in [m]$ and $i \in G_a$. By Definition 6, the coordinates of block i moved to the last k' positions are those at within-block positions $[w] \setminus Q_a$. Hence the last k' coordinates of $\boldsymbol{P}_Q \boldsymbol{e}$ receive no nonzero entry from block i if and only if $\pi(i) \in Q_a$, and consequently $\boldsymbol{P}_Q \boldsymbol{e} = (\bar{\boldsymbol{e}}, \mathbf{0}_{k'})^\top$ if and only if $\pi(G_a) \subseteq Q_a$ for every $a \in [m]$. Since the G_a partition $[w]$ and π is a bijection, the sets $\pi(G_a)$ are pairwise disjoint of size r; as also $|Q_a| = r$, the inclusions are equivalent to the equalities $\pi(G_a) = Q_a$ for all $a \in [m]$. The first claim then follows. For the second claim, the permutations π with $\pi(G_a) = Q_a$ for all a, Q fixed, correspond exactly to the tuples of bijections from G_a to Q_a, of which there are $(r!)^m$. The third claim is the count of ordered partitions of $[w]$ into m labeled sets of size r, given by the multinomial coefficient. □

Put differently, the family $\{\boldsymbol{P}_Q\}_{Q \in \mathcal{Q}_{w,r}}$ partitions the set of secrets: each of the $w!$ permutation regular candidates is compatible with exactly one of the $w!/(r!)^m$ partitions, and each partition is compatible with exactly $(r!)^m$ of those candidates. Consequently, sampling Q uniformly from $\mathcal{Q}_{w,r}$ yields an error-free information set with probability exactly $(r!)^m/w!$ for *every* fixed secret.

For comparison, sampling v-regular permutations uniformly, as done in the RSD setting, produces an error-free information set with probability exactly $(r/w)^w$ per iteration [28, Eq. (8)]. By Stirling's approximation, $r! \approx \sqrt{2\pi r}\,(r/e)^r$ and $w! \approx \sqrt{2\pi w}\,(w/e)^w$, so

$$\frac{(r!)^m/w!}{(r/w)^w} \approx \frac{(2\pi r)^{m/2}}{\sqrt{2\pi w}},$$

i.e., the partition family for PSD improves on independent per-block sampling in the standard RSD case by a factor of at least $\Theta\left(r^{m/2}\right) = \Theta\left(r^{\frac{w}{2r}}\right)$. Note that for the PSD instances obtained from PKP instances in Section 3 this corresponds to an exponential factor of order $2^{\Theta(w)}$, when assuming the PKP instance uses constant rate as well as ℓ and ν are constant, which is commonly the case. We specify the full algorithm in Algorithm 1.

Algorithm 1 Permutation-Based Regular-ISD for $\mathsf{PSD}(n, k)$

Input: Parity-check matrix $\boldsymbol{H} \in \mathbb{F}_2^{(n-k)\times n}$ and syndrome $\boldsymbol{s} \in \mathbb{F}_2^{n-k}$, where $w \mid k'$ and $(w - k'/w) \mid w$ for $w = \sqrt{n}$ and $k' = k - 2w + 1$
Output: Permutation regular $\boldsymbol{e} \in \mathbb{F}_2^n$ such that $\boldsymbol{He} = \boldsymbol{s}$

Construct $(\boldsymbol{H}', \boldsymbol{s}') \in \mathbb{F}_2^{(n-k')\times n} \times \mathbb{F}_2^{n-k'}$ from $(\boldsymbol{H}, \boldsymbol{s})$ as in Equation (4), discarding one redundant row
Set $v \leftarrow k'/w, \quad r \leftarrow w - v, \quad m \leftarrow w/r$
while true **do**
    Sample $Q = (Q_1, \ldots, Q_m)$ uniformly at random from $\mathcal{Q}_{w,r}$
    Compute $\boldsymbol{H}'\boldsymbol{P}_Q^{-1} = (\boldsymbol{H}_1 \mid \boldsymbol{H}_2)$
    Compute $L_Q \leftarrow \{\bar{\boldsymbol{x}} \in \mathbb{F}_2^{n-k'} : \boldsymbol{H}_1\bar{\boldsymbol{x}} = \boldsymbol{s}'\}$ ▷ empty or of size $2^{\dim\ker(\boldsymbol{H}_1)}$
    for $\bar{\boldsymbol{e}} \in L_Q$ **do**
        $\boldsymbol{e} \leftarrow \boldsymbol{P}_Q^{-1}(\bar{\boldsymbol{e}}, \boldsymbol{0}_{k'})^\top$
        if $\boldsymbol{e}$ is permutation regular **then**
            Return $\boldsymbol{e}$
        end if
    end for
end while

5.2 Complexity Analysis

The cost of the linear-algebra step. So far we assumed that $\boldsymbol{H}_1$ is non-singular, which is not necessarily the case. Since every solution only has a single corresponding compatible permutation discarding any non-singular $\boldsymbol{H}_1$ leads to PSD instances that cannot be solved. We therefore instead enumerate all existing solutions in those cases. Enumerating L_Q costs time proportional to $2^{d(Q)}$, where $d(Q) \coloneqq \dim\ker(\boldsymbol{H}_1)$ is the corank of the matrix $\boldsymbol{H}_1$ associated with Q. For a uniformly random $\boldsymbol{H}'$ this cost is bounded on average: $\boldsymbol{H}_1$ is invertible with probability $\prod_{i\geq 1}(1 - 2^{-i}) \approx 0.2888$, giving $\mathbb{E}[2^d] = O(1)$ for any fixed Q. Here, however, $\boldsymbol{H}'$ contains the $2w - 1$ deterministic rows of $\boldsymbol{K}$. However, a direct adaptation of [28, Remark 3.1] shows that the probability is not meaningfully changed by those equations.

The following theorem states the runtime of algorithm 1.

Theorem 4. *Let n be a perfect square, $w \coloneqq \sqrt{n}$ and $k' \coloneqq k - 2w + 1$, and assume $w \mid k'$ and that $r \coloneqq w - k'/w$ divides w. Let $(\boldsymbol{H}, \boldsymbol{s})$ be a random $\mathsf{PSD}(n, k)$ instance and let D denote the number of distinct partitions in $\mathcal{Q}_{w,r}$ compatible with any solution. Then, with high probability over the choice of the instance, Algorithm 1 solves $\mathsf{PSD}(n, k)$ in expected time*

$$T = \tilde{\mathcal{O}}\left(\frac{w!}{(r!)^{w/r} \cdot D}\right)$$

using memory polynomial in n, where the expectation is over the internal randomness of the algorithm.

Proof. The correctness follows from the previous discussion: The algorithm succeeds whenever a permutation is chosen such that $\boldsymbol{P}_Q\boldsymbol{e} = (\bar{\boldsymbol{e}}, \boldsymbol{0})$, and we have

shown that there exists exactly one partition-induced permutation for which this holds.

The running time of the algorithm is the expected number of repetitions of the loop multiplied by the time per iteration. Note that the time per iteration is bounded by a polynomial in n. As shown previously, every solution is compatible with exactly one permutation $\boldsymbol{P}_Q$ for some $Q \in \mathcal{Q}_{w,r}$. Then if the number of distinct partitions that is compatible with any solutions is $1 \leq D \leq S$, we sample one of those partitions with probability

$$D \cdot \frac{(r!)^m}{w!},$$

from which the statement follows. □

Note that, as long as $|S| \leq (r!)^{m/2}$ we generally expect $D = S$ when assuming that the compatible partitions for different solutions distributed independently and uniformly at random among all $(r!)^m$ partitions. In other words, by the birthday paradox, we do not expect compatible partitions to collide before S is larger than the square root of the total number of partitions.

Remark 1 (Optimality of the partition family). Any v-regular permutation is described by sets $R_1, \ldots, R_w \subseteq [w]$ of size r, where R_i contains the within-block positions of block i kept outside the information set. It is compatible with a permutation-regular secret defined by the permutation π precisely when

$$\pi(i) \in R_i \qquad \text{for all } i \in [w].$$

Let $\boldsymbol{A} \in \{0,1\}^{w \times w}$ be defined by $A_{i,j} = 1$ if and only if $j \in R_i$. For each permutation π, the product

$$\prod_{i=1}^{w} A_{i,\pi(i)}$$

equals 1 exactly when π is compatible and equals 0 otherwise. Hence, the permanent of $\boldsymbol{A}$ defined as

$$\operatorname{per}(\boldsymbol{A}) = \sum_{\pi \in S_w} \prod_{i=1}^{w} A_{i,\pi(i)}$$

counts exactly the number of compatible permutation-regular secrets. The Brégman–Minc inequality, conjectured by Minc [43] and proven by Brégman [18], bounds this permanent by

$$\operatorname{per}(\boldsymbol{A}) \leq (r!)^{w/r}.$$

By Lemma 1, the partition-induced choice $R_i = Q_a$ for $i \in G_a$ attains this bound whenever $r \mid w$. Hence, for a uniformly random secret, no per-block selection strategy achieves a higher probability of producing an error-free information set, and in this respect the choice underlying Algorithm 1 is optimal within the class of v-regular permutations.

Removing the divisibility conditions. The analysis above assumes that both $w \mid k'$ and $r \mid w$. These conditions need not hold for concrete parameters, and replacing r by a non-integral value would not define an actual family of permutations. We remove the conditions by allowing two retained sizes. Write

$$k' = aw + b, \qquad 0 \le b < w,$$

and fix $w - b$ blocks in which the information set has size a, and b blocks in which it has size $a+1$. The corresponding numbers of positions retained outside the information set are

$$r_0 = w - a, \qquad r_1 = w - a - 1.$$

In the parameter range considered here these quantities are positive; an empty block family is simply omitted. For $i \in \{0, 1\}$, write

$$w - b = q_0 r_0 + s_0, \qquad b = q_1 r_1 + s_1, \qquad 0 \le s_i < r_i.$$

We partition the first block family into q_0 groups of size r_0 and, when $s_0 > 0$, one remainder group of size s_0; we partition the second family analogously.

We sample an ordered partition of $[w]$ whose cell sizes are the sizes of these block groups. For a full group of size r_i, every block retains the corresponding cell. For a remainder group of size s_i, every block retains its cell together with $r_i - s_i$ padding positions outside that cell, chosen by any fixed deterministic rule. Thus every block retains its prescribed number of positions, and

$$(w - b)r_0 + br_1 = w^2 - k',$$

so the resulting matrix $\boldsymbol{H}_1$ is square.

For a fixed permutation π, the attack succeeds whenever every sampled cell is equal to the image under π of its block group. There is exactly one such ordered partition. Consequently, the success probability is at least

$$\frac{(r_0!)^{q_0} s_0! (r_1!)^{q_1} s_1!}{w!},$$

where a remainder factor is omitted when the corresponding s_i is zero. The inequality may be strict because the padding positions can create additional successful information sets.

If D denotes the number of distinct ordered partitions associated with the solutions, the exponential contribution to the expected running time is such that

$$T \le \frac{w!}{(r_0!)^{q_0} s_0! (r_1!)^{q_1} s_1!\, D}. \tag{10}$$

Moreover, when r_0 and r_1 are bounded, the remainder factors change the cost only by a constant factor, since

$$(r_0!)^{q_0} s_0! (r_1!)^{q_1} s_1! = \Theta\Big((r_0!)^{(w-b)/r_0} (r_1!)^{b/r_1}\Big).$$

We consider the expression in Equation (10) for the asymptotic comparison in the following subsection.

5.3 Application to PKP over Binary Extension Fields

The case of PKP-derived PSD instances. The above approach for solving PSD from Section 5 is stated for a random PSD instance. In contrast, the PSD instance obtained from PKP via the transformation of Section 4 has parity-check matrix $\bar{\boldsymbol{H}}$ defined in Equation (5), whose rows are binary projections of $\boldsymbol{V}^\top \otimes \boldsymbol{H}$. Thus, its entries are correlated and the random matrix estimates used for the corank

$$d(Q) = \dim \ker(\boldsymbol{H}_1)$$

do not apply directly. Nevertheless, we do not need $\bar{\boldsymbol{H}}$ to be indistinguishable from a random binary matrix. It suffices that the linear-algebra step remains polynomial on average. More precisely, for Q sampled uniformly from $\mathcal{Q}_{w,r}$ (defined in Definition 6), we require

$$\mathbb{E}_Q\left[2^{d(Q)}\right] = \mathrm{poly}(n)$$

with high probability over the derived instance. Canonical choices of $G_1, \ldots, G_m$ expose dependencies among the corresponding rows of $\boldsymbol{V}$. Fix an $\mathbb{F}_2$-basis of $\mathbb{F}_{2^\nu}$ and let $\boldsymbol{V}_i^{(2)} \in \mathbb{F}_2^{\ell\nu}$ be the vector obtained by expanding every entry of the i-th row $\boldsymbol{V}_i \in \mathbb{F}_{2^\nu}^{\ell}$ in this basis and concatenating the resulting binary coefficients. We choose the block partition $G_1, \ldots, G_m$ so that

$$\mathsf{rank}_{\mathbb{F}_2}\{\boldsymbol{V}_i^{(2)} : i \in G_a\} = r \qquad \text{for every } a \in [m].$$

Once the groups are fixed, we sample the ordered position partition $Q = (Q_1, \ldots, Q_m)$ uniformly from $\mathcal{Q}_{w,r}$ as in Algorithm 1. This adaptation does not change the success probability $(r!)^{w/r}/w!$, which holds for every fixed block partition.

Experiments. We compare experimentally the empirical probability $\Pr[d(Q) = 0]$ and the average enumeration cost $\mathbb{E}[2^{d(Q)}]$ for the matrices derived from random $\mathsf{PKP}(\bar{n}, \bar{k}, \ell, 2^\nu)$ instances, with $\ell = 1$ fixed and rate $\bar{k}/\bar{n} = 1/2$, with the corresponding quantities for uniformly random square binary matrices. For each $\bar{n}$ we sample 20 independent PKP instances and 100 independent partitions $Q \leftarrow \mathcal{Q}_{w,r}$ per instance. The adapted block partition was found for 99.1% of the instances. Figure 1 presents the result of such an experiment and shows that the PKP-derived matrices and the uniform matrices remain close. In contrast, the canonical consecutive block partition for PKP-derived instances results in a significant increase of the average enumeration cost. For instance, for $(\bar{n}, \bar{k}, \ell, \nu, r) = (64, 32, 1, 13, 8)$, we observe $E[2^{d(Q)}] = 21.09$.

Comparison with KMP. We now compare the asymptotic cost of the attack presented in Section 5 with that of the attack by Koussa, Macario-Rat, and Patarin (KMP) [39]. To isolate the benefit of exploiting the permutation structure of the solution, we also include the plain permutation-based Regular-ISD approach outlined by Esser and Santini [28]. More precisely, by the Esser–Santini

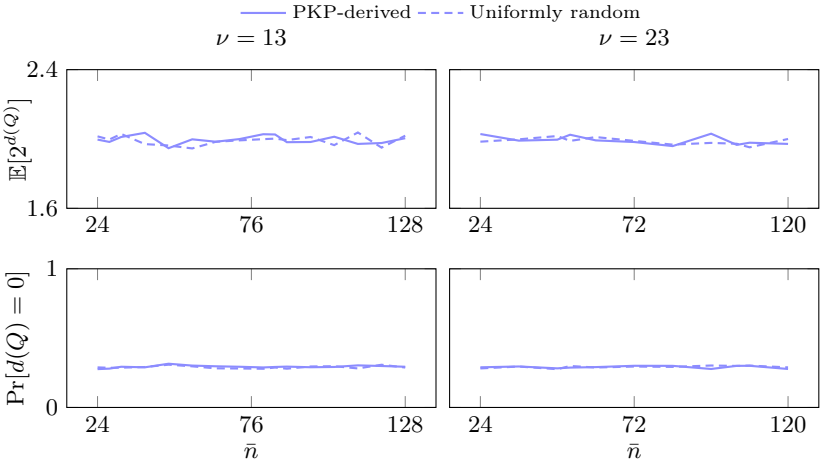

Fig. 1: PKP-derived matrices vs uniformly random.

baseline we mean applying the transformation of Section 4 to the PKP instance and then running their original permutation-based regular-ISD solving algorithm on the resulting PSD instance viewed as an RSD instance; hence, only the $\bar{n}$ regularity equations used by that algorithm are included. Since the comparison is only asymptotic, throughout this subsection, we suppress polynomial factors in the estimated running times.

Figure 2 compares the asymptotic time complexities in two parameter regimes. In Figure 2a, we fix the PERK v2.2.1 parameters corresponding to NIST security level I, namely $\bar{n} = 64$, $\bar{k} = 37$, and $\ell = 1$, and vary ν (note, $\nu = 11$ in PERK). The attack from Section 5 significantly outperforms KMP for small values of ν. Indeed, in this regime, the PKP instance is expected to admit multiple solutions, a scenario from which our attack particularly benefits, whereas KMP does not. Our attack also performs better for large values of ν, since increasing ν yields more equations in the induced PSD instance.

KMP performs better in the intermediate regime that is, pherarps, the most interesting cryptographically speaking. Specifically, this corresponds to the case where ν is close to the smallest value for which the PKP instance has a unique solution in expectation, as is the case for the parameters in PERK [2] and SETH [8]. This behavior is more clearly illustrated in Figure 2b. In this plot, we fix $\ell = 1$ and $\nu = 11$ and, for each value of $\bar{n}$, choose the largest $\bar{k}$ for which the PKP instance is expected to have a unique solution. In this regime, KMP has the lower asymptotic cost. Finally, across both experiments, our attack consistently improves upon the plain permutation-based Regular-ISD approach of Esser and Santini.

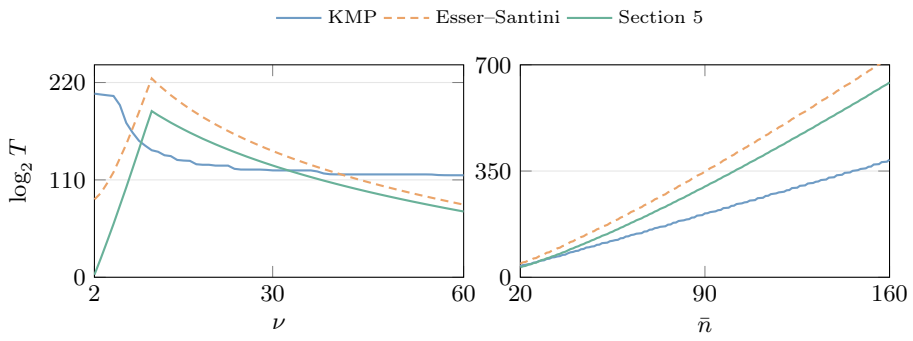

(a) $\mathsf{PKP}(64, 37, 1, 2^{\nu})$. (b) $\mathsf{PKP}(\bar{n}, \bar{k}, 1, 2^{11})$.

Fig. 2: Asymptotic time-complexity comparison.

6 Additional Polynomial Regimes for PEP

We now present an additional reduction from PEP to the graph isomorphism problem that allows for polynomial-time solutions under certain parameter combinations. In particular, we heuristically expect the reduction to work in all cases for random self-orthogonal codes over non-prime fields, as well as for (weakly) self-dual codes in most non-binary scenarios.

Let $\boldsymbol{G}_0, \boldsymbol{G}_1 \in \mathbb{F}_q^{k \times n}$ define a $\mathsf{PEP}(n, k, q)$ instance for $q = p^{\nu}$. The reduction works by exploiting the fact that we can use any endomorphism σ in the Galois group $\mathrm{Gal}(\mathbb{F}_q/\mathbb{F}_p)$ to obtain a second $\mathsf{PEP}(n, k, q)$ instance sharing the secret permutation. Indeed, if $\boldsymbol{G}_1 = \boldsymbol{S}\boldsymbol{G}_0\boldsymbol{P}$ then it follows also that $\boldsymbol{G}_1^{\sigma} = \boldsymbol{S}^{\sigma}\boldsymbol{G}_0^{\sigma}\boldsymbol{P}$, where $\boldsymbol{S}^{\sigma} \coloneqq \sigma(\boldsymbol{S})$, $\boldsymbol{G}_i^{\sigma} \coloneqq \sigma(\boldsymbol{G}_i)$, and the permutation $\boldsymbol{P}$ is fixed by σ since its entries are binary. Therefore, $\boldsymbol{G}_0^{\sigma}, \boldsymbol{G}_1^{\sigma}$ is a second $\mathsf{PEP}(n, k, q)$ instance whose solution shares the same permutation $\boldsymbol{P}$.

We then make use of the binary operator

$$\tau \colon (\boldsymbol{G}_i, \boldsymbol{G}_j) \mapsto \boldsymbol{G}_j^{\top}(\boldsymbol{G}_i\boldsymbol{G}_j^{\top})^{-1}\boldsymbol{G}_i.$$

to reduce to a Graph Isomorphism Problem, using the following proposition.

Proposition 7. *Let $\boldsymbol{G_1} = \boldsymbol{S}\boldsymbol{G_0}\boldsymbol{P}$ and $\boldsymbol{G'_1} = \boldsymbol{S'}\boldsymbol{G'_0}\boldsymbol{P}$ be two $\mathsf{PEP}(n, k, q)$ instances sharing the same permutation $\boldsymbol{P}$. If $\boldsymbol{G}'_1\boldsymbol{G}_1^{\top}$ is invertible, then $\boldsymbol{P}$ is also a solution to the graph isomorphism problem $\bar{\boldsymbol{G}}_1 = \boldsymbol{P}^{\top}\bar{\boldsymbol{G}}_0\boldsymbol{P}$, where $\bar{\boldsymbol{G}}_i \coloneqq \tau(\boldsymbol{G}'_i, \boldsymbol{G}_i)$, $i = 0, 1$.*

Proof. The result follows from the following steps

$$\begin{aligned}
\tau(\boldsymbol{G}_1', \boldsymbol{G}_1) &= \boldsymbol{G}_1^\top (\boldsymbol{G}_1' \boldsymbol{G}_1^\top)^{-1} \boldsymbol{G}_1' \\
&= (\boldsymbol{S}\boldsymbol{G}_0\boldsymbol{P})^\top \left((\boldsymbol{S}'\boldsymbol{G}_0'\boldsymbol{P})(\boldsymbol{S}\boldsymbol{G}_0\boldsymbol{P})^\top \right)^{-1} (\boldsymbol{S}'\boldsymbol{G}_0'\boldsymbol{P}) \\
&= \boldsymbol{P}^\top \boldsymbol{G}_0^\top \boldsymbol{S}^\top ((\boldsymbol{S}'\boldsymbol{G}_0'\boldsymbol{P})(\boldsymbol{P}^\top \boldsymbol{G}_0^\top \boldsymbol{S}^\top))^{-1} \boldsymbol{S}'\boldsymbol{G}_0'\boldsymbol{P} \\
&= \boldsymbol{P}^\top \boldsymbol{G}_0^\top \boldsymbol{S}^\top (\boldsymbol{S}^\top)^{-1} (\boldsymbol{G}_0' \boldsymbol{G}_0^\top)^{-1} (\boldsymbol{S}')^{-1} \boldsymbol{S}'\boldsymbol{G}_0'\boldsymbol{P} \\
&= \boldsymbol{P}^\top \boldsymbol{G}_0^\top (\boldsymbol{G}_0' \boldsymbol{G}_0^\top)^{-1} \boldsymbol{G}_0'\boldsymbol{P} \\
&= \boldsymbol{P}^\top \tau(\boldsymbol{G}_0', \boldsymbol{G}_0) \boldsymbol{P},
\end{aligned}$$

where we note that the invertibility of $\boldsymbol{G}_0'\boldsymbol{G}_0^\top$ is tied to that of $\boldsymbol{G}_1'\boldsymbol{G}_1^\top$, since $\boldsymbol{G}_1'\boldsymbol{G}_1^\top = \boldsymbol{S}(\boldsymbol{G}_0'\boldsymbol{G}_0^\top)\boldsymbol{S}^\top$. □

The crucial requirement is being able to find $\sigma \in \mathrm{Gal}(\mathbb{F}_q/\mathbb{F}_p)$ such that $\boldsymbol{G}_0^\sigma \boldsymbol{G}_0^\top$ has full-rank. For self-dual codes, we have identified two families of cases when this is never possible, as demonstrated in the next subsection: (i) if the field is binary, and (ii) if the field has characteristic $p \equiv 1 \mod 4$ and the length is $n \equiv 2 \mod 4$. In all other cases, we expect to have no orthogonality relationship between G and G^σ, which we state formally through the following assumption.

Assumption 3 *Let $\boldsymbol{G} \in \mathbb{F}_q^{k \times n}$ generate a random self-orthogonal code, with $q = p^\nu$ and p prime. If neither one of the following conditions hold:*

(i) *the code is self-dual and $p = 2$,*
(ii) *the code is self-dual and $p \equiv 1 \mod 4$ and $n \equiv 2 \mod 4$,*

then any $\sigma \in \mathrm{Gal}(\mathbb{F}_q/\mathbb{F}_p)$ that is non-trivial will make $\boldsymbol{G}^\sigma \boldsymbol{G}^\top$ have full rank k most of the time, with lower ranks occurring with probability $\mathcal{O}(1/p^{n(\nu/2-1)+1})$.

The failure probability for Assumption 3 is heuristically explained in the next subsection, and experiments to validate the assumption are presented in Appendix B.

We can use any of the powers of the Frobenius map, given by $\sigma_i : x \mapsto x^{p^i}$, as our endomorphism. Although there are $\nu - 1$ such endomorphisms that are non-trivial, the events of each $\boldsymbol{G}^{\sigma_i}\boldsymbol{G}^\top$ having full rank are not all independent. In particular, since endomorphisms preserve the rank, we have $\mathsf{rank}(\boldsymbol{G}^{\sigma_i}\boldsymbol{G}^\top) = \mathsf{rank}(\sigma_{\nu-i}(\boldsymbol{G}^{\sigma_i}\boldsymbol{G}^\top)) = \mathsf{rank}(\boldsymbol{G}(\boldsymbol{G}^{\sigma_{\nu-i}})^\top) = \mathsf{rank}(\boldsymbol{G}^{\sigma_{\nu-i}}\boldsymbol{G}^\top)$, so the events of σ_i and $\sigma_{\nu-i}$ giving full rank are perfectly correlated. We assume no further correlation between these events by making the following assumption.

Assumption 4 *Under the same conditions as Assumption 3, there exists at least one $\sigma \in \mathrm{Gal}(\mathbb{F}_q/\mathbb{F}_p)$ such that $\boldsymbol{G}^\sigma \boldsymbol{G}^\top$ has full rank, except with probability $\mathcal{O}(1/p^{(n(\nu/2-1)+1)\cdot\lceil\frac{\nu-1}{2}\rceil})$.*

Once we find an endomorphism σ that makes $\boldsymbol{G}^\sigma \boldsymbol{G}^\top$ invertible, we obtain the graph isomorphism problem $\bar{\boldsymbol{G}}_1 = \boldsymbol{P}^\top \bar{\boldsymbol{G}}_0 \boldsymbol{P}$ derived from Proposition 7. Note that $\bar{\boldsymbol{G}}_i \coloneqq \tau(\boldsymbol{G}_i^\sigma, \boldsymbol{G}_i)$ are $n \times n$, but only have rank k by construction.

If we model their rows as determined by k uniformly distributed $\mathbb{F}_q$-degrees of freedom, then there is only an $\mathcal{O}(q^{-k})$ probability for a given pair of rows to collide. This gives a probability $\mathcal{O}(n^2q^{-k})$ for at least one pair to collide, which we formalize as the following assumption.

Assumption 5 *If $\boldsymbol{G} \in \mathbb{F}_q^{k\times n}$ generates a random self-orthogonal code over $\mathbb{F}_q$ and σ is an automorphism such that $\boldsymbol{G}^\sigma\boldsymbol{G}^\top$ is invertible, then the rows of $\tau(\boldsymbol{G}^\sigma, \boldsymbol{G})$ are all unique, except with probability $\mathcal{O}(n^2q^{-k})$.*

Assumption 5 implies in particular that the solution to the graph isomorphism problem derived from Proposition 7 is unique, thus solving it will have solved the original PEP problem. We can now state our main result for this section.

Theorem 5. *Under Assumption 4 and Assumption 5, there exists an algorithm that solves $\mathsf{PEP}(n, k, q)$ for $q = p^\nu$ with p prime in polynomial time, with failure probability*

$$\mathcal{O}(p^{-(n(\nu/2-1)+1)\cdot\lceil\frac{\nu-1}{2}\rceil} + n^2p^{-\nu k}),$$

except in the following scenarios:

(i) $\nu = 1$
(ii) *The code is self-dual and* $p = 2$
(iii) *The code is self-dual and* $p \equiv 1 \mod 4$ *and* $n \equiv 2 \mod 4$

Proof. The result follows immediately by probing the powers of Frobenius σ_i for $i = 1, 2, \ldots \lceil\frac{\nu-1}{2}\rceil$ until $\boldsymbol{G}_1^{\sigma_i}\boldsymbol{G}_1^\top$ has full rank, then solving the Graph Isomorphism Problem derived from Proposition 7. □

Remark 2. For the general Graph Isomorphism Problem no polynomial-time algorithm is known; the best known upper bound is Babai's celebrated quasi-polynomial time algorithm, which runs in time $\exp(O(\log n)^c)$ for some constant c [5]. However, given the (expected) uniqueness of the PEP solution for random instances, there is a simple row-matching solver for GIP instances with unique solution [13], which we describe next. Let $\bar{\boldsymbol{G}}_0, \bar{\boldsymbol{G}}_1 \in \mathbb{F}_q^{n\times n}$ be the (weighted) adjacency matrices of two isomorphic graphs, and assume that every row of $\bar{\boldsymbol{G}}_0$ has a *unique* collision among the rows of $\bar{\boldsymbol{G}}_1$, i.e., the multiset of entries of each row (including the diagonal entry) is distinct. Under this assumption, the search version of the underlying PEP can be solved by the following procedure:

- **Row labeling of $\bar{\boldsymbol{G}}_0$.** For each row index $i \in \{0, \ldots, n-1\}$, extract the entries of the i-th row of $\bar{\boldsymbol{G}}_0$ and sort them, obtaining a canonical representation of the row's entry multiset. Store the pair (sorted multiset, i) in a hash table, using the (string-encoded) sorted multiset as the key.
- **Collision search for $\bar{\boldsymbol{G}}_1$.** For each row index $j \in \{0, \ldots, n-1\}$, compute the sorted entry multiset of the j-th row of $\bar{\boldsymbol{G}}_1$ in the same way, and look up the resulting key in the hash table. By the uniqueness assumption, this yields exactly one matching index i of $\bar{\boldsymbol{G}}_0$.
- **Output.** Return the list of matched pairs (i, j), which encodes the row permutation (and hence the isomorphism) mapping $\bar{\boldsymbol{G}}_0$ to $\bar{\boldsymbol{G}}_1$.

Complexity. The algorithm runs in polynomial time. Sorting the n entries of a single row costs $O(n \log n)$ operations, and this is done for all $2n$ rows of $\bar{\boldsymbol{G}}_0$ and $\boldsymbol{G}_1$, for a total of $O(n^2 \log n)$. Each hash-table insertion and lookup takes expected time linear in the key length, i.e., $O(n)$ per row and $O(n^2)$ overall. Hence the total expected running time is

$$O\big(n^2 \log n\big) = \tilde{O}(n^2),$$

which is polynomial in the input size. We stress that this efficiency crucially relies on the uniqueness of the row multisets: the procedure solves only this restricted ("easy") class of instances.

Remark 3. Proposition 7 allows for the reduction to the graph isomorphism problem when $\boldsymbol{G}_0, \boldsymbol{G}_1$ and $\boldsymbol{G}'_0, \boldsymbol{G}'_1$ are any pair of PEP instances sharing the same permutation, not necessarily related by Frobenius conjugation. Therefore, it also breaks the security of the two sample PEP problem studied in [20] (with the name of 2-PCE), even in the prime field setting, and shows that constructions like the multi-key variant left as an open problem in [4] are not viable.

6.1 Non-Reducible Cases

The following theorem gives a characterization of the cases when Assumption 3 fails to hold.

Theorem 6. *Let $\boldsymbol{G} \in \mathbb{F}_q^{k \times n}$ generate a code $\mathcal{C}$ over a finite field of characteristic p, and let $\sigma \in \mathrm{Gal}(\mathbb{F}_q/\mathbb{F}_p)$. Then, $\boldsymbol{G}^\sigma \boldsymbol{G}^\top$ has full rank if and only if $\mathcal{C}^\perp \cap \mathcal{C}^\sigma = \{\mathbf{0}\}$.*

Proof. Suppose that $\boldsymbol{v} \in \mathcal{C}^\perp \cap \mathcal{C}^\sigma$ is non-zero. Because $\boldsymbol{v} \in \mathcal{C}^\sigma$, there exists $\boldsymbol{a} \in \mathbb{F}_q^k$ such that $\boldsymbol{a}\boldsymbol{G}^\sigma = \boldsymbol{v}$, but because $\boldsymbol{v} \in \mathcal{C}^\perp$ we have $\boldsymbol{a}\boldsymbol{G}^\sigma \boldsymbol{G}^\top = \boldsymbol{v}\boldsymbol{G}^\top = \mathbf{0}$, so $\boldsymbol{a}$ is in the left kernel of $\boldsymbol{G}^\sigma \boldsymbol{G}^\top$.

Conversely, if $\boldsymbol{a}$ is in the left kernel of $\boldsymbol{G}^\sigma \boldsymbol{G}^\top$ then $\boldsymbol{v} \coloneqq \boldsymbol{a}\boldsymbol{G}^\sigma$ is in the left kernel of $\boldsymbol{G}^\top$, so $\boldsymbol{v} \in \mathcal{C}^\perp$. On the other hand, $\boldsymbol{v} = \boldsymbol{a}\boldsymbol{G}^\sigma$ is clearly in $\mathcal{C}^\sigma$, so $\boldsymbol{v} \in \mathcal{C}^\perp \cap \mathcal{C}^\sigma$. □

For self-dual codes, we have that $\mathcal{C}^\perp \cap \mathcal{C}^\sigma \neq \{\mathbf{0}\}$ if and only if $\mathcal{C} \cap \mathbb{F}_p \neq \{\mathbf{0}\}$, so we need to consider the probability of a self-dual code containing a non-zero codeword in the base field. Note that $\boldsymbol{v} \in \mathcal{C}$ if and only if $\boldsymbol{v}\boldsymbol{G}^\top = 0$. Assuming that $\boldsymbol{v}$ is defined over $\mathbb{F}_p$ and unrolling $\boldsymbol{G}$ into its base-field components to obtain a matrix $\tilde{\boldsymbol{G}} \in \mathbb{F}_p^{k\nu \times n}$, we obtain $\boldsymbol{v}\tilde{\boldsymbol{G}}^\top = 0$ which is a system of $k\nu = n\nu/2$ equations in n unknowns over $\mathbb{F}_p$. For $\nu \geq 2$ this system is overdetermined, and in generic cases we expect $\boldsymbol{v} = \mathbf{0}$ to be the only solution. This will be true unless the columns of $\tilde{\boldsymbol{G}}$ are linearly dependent, which modeled as a random matrix in $\mathbb{F}_p^{n\nu/2 \times n}$, happens with probability $\mathcal{O}(1/p^{n(\nu/2-1)+1})$. This heuristic supports the failure probability that was presented in Assumption 3.

Despite the aforementioned heuristics, the special cases of Assumption 3 identify specific settings with additional structure that always forces a non-trivial intersection $\mathcal{C}^\top \cap \mathcal{C}^\sigma$. The next result explains the failure for all self-dual codes over binary fields.

Theorem 7. *If $\mathcal{C}$ is a self-dual code over a binary field $\mathbb{F}_{2^\nu}$ and σ is any endomorphism of $\mathbb{F}_{2^\nu}$, then $\mathcal{C} \cap \mathcal{C}^\sigma$ is always non-trivial.*

Proof. Since $\mathcal{C}$ is self-orthogonal over a binary field, every codeword has zero sum of coordinates. Therefore, $\mathbf{1}_n$ is contained in the dual code, which equals the code itself. Because $\mathbf{1}_n$ is σ-invariant, it is moreover contained in $\mathcal{C}^\sigma$. □

In order to prove the second failure case of Assumption 3, we will make use of the following lemma.

Lemma 2. *For any value of k and $n = 2k$, the set of all (n, k)-codes defined over any finite field that are self-dual are split into two disjoint families, where*

$$dim\,(\mathcal{C}_1 \cap \mathcal{C}_2) \equiv k \mod 2 \iff \mathcal{C}_1 \text{ and } \mathcal{C}_2 \text{ are in the same family.}$$

Proof. This follows from a standard parity theorem for maximal isotropic subspaces with hyperbolic forms, regarding $\mathcal{C}_1$ and $\mathcal{C}_2$ as maximal isotropic with respect to the conventional dot product (see for example [24, Theorem 6.11]). □

Theorem 8. *Let $q = p^\nu$ with $p \equiv 1 \mod 4$ prime, $\sigma \in \mathbb{F}_q/\mathbb{F}_p$, and $n \equiv 2 \mod 4$. Then any self-dual $(n, n/2)$-code $\mathcal{C}$ defined over $\mathbb{F}_q$ satisfies $\mathcal{C} \cap \mathcal{C}^\sigma \neq \{0\}$.*

Proof. We will prove the theorem by showing that $\mathcal{C}$ and $\mathcal{C}^\sigma$ always belong to the same family in the context of Lemma 2, which by that same lemma implies that

$$dim(\mathcal{C} \cap \mathcal{C}^\sigma) \equiv n/2 \equiv 1 \mod 2.$$

and in particular that $\mathcal{C} \cap \mathcal{C}^\sigma \neq \{0\}$.

We begin by proving the existence of a self-dual $(n, n/2)$-code $\mathcal{W}$ that is generated by vectors over $\mathbb{F}_p$. Since $p \equiv 1 \mod 4$, there exists an element $i \in \mathbb{F}_p$ such that $i^2 = -1$. Let $e_j \in \mathbb{F}_q^n$ be the standard unit basis and define the vectors $c_j \coloneqq e_{2j} + ie_{2j+1}$ for $0 \leq j < n/2$. We have that $e_j \cdot e_{j'} = 0$ for $j \neq j'$ because the non-zero coordinates are disjoint between different vectors, and also $e_j \cdot e_j = 1 + i^2 = 0$. Therefore, the code $\mathcal{W} = span\left(\{c_j\}_{j=0}^{n/2-1}\right)$ is self-dual.

Given the existence of $\mathcal{W}$, it then follows that $\mathcal{C}$ and $\mathcal{C}^\sigma$ always belong to the same family in the context of Lemma 2. Indeed, we have $dim(\mathcal{W} \cap \mathcal{C}) = dim((\mathcal{W} \cap \mathcal{C})^\sigma) = dim(\mathcal{W} \cap \mathcal{C}^\sigma)$, so $\mathcal{C}$ and $\mathcal{C}^\sigma$ are either both in the same family as $\mathcal{W}$ or both in the opposite family as $\mathcal{W}$. This completes the proof. □

Remark 4. We emphasize that our results only imply that PEP cannot be reduced to GIP for odd characteristic $p \equiv 1 \bmod 4$ and $n \equiv 2 \bmod 4$ via the Frobenius conjugation. Therefore, our results should be taken as cautionary and avoid field extensions in general. [2]

[2] For example, one approach to tackle the missing scenario could be to extend the results of Section 4.2; however, additional considerations (and assumptions) must be taken and further analysis is required.

Acknowledgments

The authors wish to thank Javier Verbel for insightful discussions during the preparation of this manuscript.

Generative AI agents were utilized to refine the overall quality of the manuscript, as well as to assist in cross-checking and validating the correctness of the results. All final content, calculations, and conclusions were independently reviewed and verified by the authors, who accept full responsibility for the integrity of the work.

References

1. Aaraj, N., Bettaieb, S., Bidoux, L., Budroni, A., Dyseryn, V., Esser, A., Feneuil, T., Gaborit, P., Kulkarni, M., Mateu, V., Palumbi, M., Perin, L., Rivain, M., Tillich, J., Xagawa, K.: PERK. Tech. rep., National Institute of Standards and Technology (2024), available at https://csrc.nist.gov/Projects/pqc-dig-sig/round-2-additional-signatures
2. Aaraj, N., Bettaieb, S., Bidoux, L., Budroni, A., Dyseryn, V., Esser, A., Feneuil, T., Gaborit, P., Kulkarni, M., Mateu, V., Palumbi, M., Perin, L., Rivain, M., Tillich, J.P., Xagawa, K.: PERK: A post-quantum signature scheme (version 2.1.0) (Sep 2025), specification document; available at https://pqc-perk.org/assets/downloads/perk-v2.1.0.pdf
3. Albrecht, M.R., Benčina, B., Lai, R.W.F.: Hollow LWE: A new spin - unbounded updatable encryption from LWE and PCE. In: Fehr, S., Fouque, P.A. (eds.) EUROCRYPT 2025, Part VIII. LNCS, vol. 15608, pp. 363–392. Springer, Cham (May 2025). https://doi.org/10.1007/978-3-031-91101-9_13
4. Albrecht, M.R., Benčina, B., Lai, R.W.F.: Hollow lwe: A new spin, unbounded updatable encryption from lwe and pce. In: Advances in Cryptology – EUROCRYPT 2025. LNCS, vol. 15608, pp. 363–392. Springer (2025). https://doi.org/10.1007/978-3-031-91101-9_13, artifact available at https://artifacts.iacr.org/eurocrypt/2025/a9
5. Babai, L.: Graph isomorphism in quasipolynomial time [extended abstract]. In: Proceedings of the 48th Annual ACM Symposium on Theory of Computing (STOC 2016). pp. 684–697. ACM (2016). https://doi.org/10.1145/2897518.2897542
6. Baldi, M., Barenghi, A., Beckwith, L., Biasse, J., Chou, T., Esser, A., Gaj, K., Karl, P., Mohajerani, K., Pelosi, G., Persichetti, E., Saarinen, M.O., Santini, P., Wallace, R., Zweydinger, F.: LESS — Linear Equivalence Signature Scheme. Tech. rep., National Institute of Standards and Technology (2024), available at https://csrc.nist.gov/Projects/pqc-dig-sig/round-2-additional-signatures
7. Baldi, M., Battagliola, M., Borin, G., Di Crescenzo, G., El Mechri, R., Meneghetti, A., Persichetti, E., Santini, P., Zweydinger, F.: LEAST: Linear Equivalence Action Threshold Signature (Preview Writeup). https://csrc.nist.gov/csrc/media/Projects/threshold-cryptography/documents/TCall-1/LEAST-PW01.pdf (Jan 20 2026), preliminary version submitted to the NIST First Call for Multi-Party Threshold Schemes
8. Baldi, M., Battagliola, M., El Mechri, R., Mattiuz, L., Rinaudo, M., Santini, P., Schiavoni, R., Signorini, E.: SETH: Subcode equivalence in the head. In: Fourteenth International Workshop on Coding and Cryptography (WCC 2026). Paris, France

(Jun 2026), https://wcc2026.inria.fr/assets/final_versions/WCC2026_paper_47.pdf, extended abstract, presented at WCC 2026
9. Baldi, M., Battagliola, M., Mechri, R.E., Santini, P., Schiavoni, R., Zuane, D.D.: SPECK: Signatures from permutation equivalence of codes and kernels. Cryptology ePrint Archive, Report 2025/923 (2025), https://eprint.iacr.org/2025/923
10. Bardet, M., Brion, C., Otmani, A., Saeed, M., Sendrier, N.: Solving the linear equivalence problem from single codeword matching. In: Heninger, N., Rosulek, M. (eds.) Advances in Cryptology – CRYPTO 2026. pp. 256–288. Springer Nature Switzerland, Cham (2026)
11. Bardet, M., Otmani, A., Saeed-Taha, M.: Permutation code equivalence is not harder than graph isomorphism when hulls are trivial. In: 2019 IEEE International Symposium on Information Theory (ISIT). pp. 2464–2468 (2019). https://doi.org/10.1109/ISIT.2019.8849855
12. Battagliola, M., Borin, G., Di Crescenzo, G., Meneghetti, A., Persichetti, E.: Enhancing threshold group action signature schemes: Adaptive security and scalability improvements. In: Niederhagen, R., Saarinen, M.J.O. (eds.) Post-Quantum Cryptography - 16th International Workshop, PQCrypto 2025, Part I. pp. 129–161. Springer, Cham (Apr 2025). https://doi.org/10.1007/978-3-031-86599-2_5
13. Battagliola, M., Mora, R., Santini, P.: Using the schur product to solve the code equivalence problem. IEEE Transactions on Information Theory **72**(7), 4675–4694 (2026). https://doi.org/10.1109/TIT.2026.3694630
14. Berger, T.P., Gueye, C.T., Klamti, J.B.: A np-complete problem in coding theory with application to code based cryptography. In: El Hajji, S., Nitaj, A., Souidi, E.M. (eds.) Codes, Cryptology and Information Security. pp. 230–237. Springer International Publishing, Cham (2017)
15. Beullens, W.: Sigma protocols for MQ, PKP and SIS, and Fishy signature schemes. In: Canteaut, A., Ishai, Y. (eds.) EUROCRYPT 2020, Part III. LNCS, vol. 12107, pp. 183–211. Springer, Cham (May 2020). https://doi.org/10.1007/978-3-030-45727-3_7
16. Biasse, J.F., Micheli, G., Persichetti, E., Santini, P.: LESS is more: Code-based signatures without syndromes. In: Nitaj, A., Youssef, A.M. (eds.) AFRICACRYPT 20. LNCS, vol. 12174, pp. 45–65. Springer, Cham (Jul 2020). https://doi.org/10.1007/978-3-030-51938-4_3
17. Bidoux, L., Gaborit, P.: Compact post-quantum signatures from proofs of knowledge leveraging structure for the pkp, sd and rsd problems. In: El Hajji, S., Mesnager, S., Souidi, E.M. (eds.) Codes, Cryptology and Information Security. pp. 10–42. Springer Nature Switzerland, Cham (2023)
18. Brégman, L.M.: Some properties of nonnegative matrices and their permanents. Soviet Mathematics Doklady **14**(4), 945–949 (1973)
19. Briaud, P., Øygarden, M.: A new algebraic approach to the regular syndrome decoding problem and implications for PCG constructions. In: Hazay, C., Stam, M. (eds.) EUROCRYPT 2023, Part V. LNCS, vol. 14008, pp. 391–422. Springer, Cham (Apr 2023). https://doi.org/10.1007/978-3-031-30589-4_14
20. Budroni, A., Chi-Domínguez, J.J., D'Alconzo, G., Di Scala, A.J., Kulkarni, M.: Don't use it twice! Solving relaxed linear equivalence problems. In: Chung, K.M., Sasaki, Y. (eds.) ASIACRYPT 2024, Part VIII. LNCS, vol. 15491, pp. 35–65. Springer, Singapore (Dec 2024). https://doi.org/10.1007/978-981-96-0944-4_2
21. Budroni, A., Defranceschi, M., Pintore, F.: Recursion enabled: Improved cryptanalysis of the permuted kernel problem. In: Bardet, M., Niederhagen, R. (eds.) Post-Quantum Cryptography – 17th International Workshop, PQCrypto 2026,

Saint-Malo, France, April 14–16, 2026, Proceedings, Part II. Lecture Notes in Computer Science, vol. 16492, pp. 207–241. Springer Nature Switzerland, Cham (2026). https://doi.org/10.1007/978-3-032-22698-3_7
22. Budroni, A., Esser, A.: One pair to rule them all: Towards an optimal algorithm for solving code equivalence via codeword search. In: Heninger, N., Rosulek, M. (eds.) Advances in Cryptology – CRYPTO 2026. pp. 319–346. Springer Nature Switzerland, Cham (2026)
23. Budroni, A., Chi-Domínguez, J.J., Chavez-Saab, J., Esser, A.: Accompanying repository, https://github.com/AlessandroBudroni/extension_fields_pep_pkp/
24. Cameron, P.J.: Projective and Polar Spaces. No. 13 in QMW Maths Notes, Queen Mary and Westfield College, School of Mathematical Sciences, London (1992), https://webspace.maths.qmul.ac.uk/p.j.cameron/pps/
25. Chi, J.J., Oliveira, T.: Attacking a binary GLS elliptic curve with Magma. In: Lauter, K.E., Rodríguez-Henríquez, F. (eds.) LATINCRYPT 2015. LNCS, vol. 9230, pp. 308–326. Springer, Cham (Aug 2015). https://doi.org/10.1007/978-3-319-22174-8_17
26. Chi-Domínguez, J.J., Rodríguez-Henríquez, F., Smith, B.: Extending the gls endomorphism to speed up ghs weil descent using magma. Finite Fields and Their Applications **75**, 101891 (2021). https://doi.org/https://doi.org/10.1016/j.ffa.2021.101891
27. Couvreur, A., Otmani, A., Tillich, J.P.: Polynomial time attack on wild McEliece over quadratic extensions. In: Nguyen, P.Q., Oswald, E. (eds.) EUROCRYPT 2014. LNCS, vol. 8441, pp. 17–39. Springer, Berlin, Heidelberg (May 2014). https://doi.org/10.1007/978-3-642-55220-5_2
28. Esser, A., Santini, P.: Not just regular decoding: Asymptotics and improvements of regular syndrome decoding attacks. In: Reyzin, L., Stebila, D. (eds.) CRYPTO 2024, Part VI. LNCS, vol. 14925, pp. 183–217. Springer, Cham (Aug 2024). https://doi.org/10.1007/978-3-031-68391-6_6
29. Faugère, J.C., Koussa, E., Macario-Rat, G., Patarin, J., Perret, L.: PKP-based signature scheme. Cryptology ePrint Archive, Report 2018/714 (2018), https://eprint.iacr.org/2018/714
30. Feneuil, T.: Building MPCitH-based signatures from MQ, MinRank, rank SD and PKP. Cryptology ePrint Archive, Report 2022/1512 (2022), https://eprint.iacr.org/2022/1512
31. Frey, G.: How to disguise an elliptic curve (1998), alk at ECC'98, Waterloo, public version available at https://cr.yp.to/bib/1998/frey-disguise.ps (1998)
32. Galbraith, S.D., Hess, F., Smart, N.P.: Extending the GHS Weil descent attack. In: Knudsen, L.R. (ed.) EUROCRYPT 2002. LNCS, vol. 2332, pp. 29–44. Springer, Berlin, Heidelberg (Apr / May 2002). https://doi.org/10.1007/3-540-46035-7_3
33. Gaudry, P., Hess, F., Smart, N.P.: Constructive and destructive facets of Weil descent on elliptic curves. Journal of Cryptology **15**(1), 19–46 (Jan 2002). https://doi.org/10.1007/s00145-001-0011-x
34. Hankerson, D., Karabina, K., Menezes, A.: Analyzing the galbraith-lin-scott point multiplication method for elliptic curves over binary fields. IEEE Transactions on Computers **58**(10), 1411–1420 (2009)
35. Hanzlik, L., Lai, Y.F., Mula, M., Paracucchi, E., Slamanig, D., Tang, G.: Tanuki: New frameworks for (concurrently secure) blind signatures from post-quantum group actions. In: Hanaoka, G., Yang, B.Y. (eds.) ASIACRYPT 2025, Part IV. LNCS, vol. 16248, pp. 35–69. Springer, Singapore (Dec 2025). https://doi.org/10.1007/978-981-95-5113-2_2

36. Hanzlik, L., Lai, Y.F., Paracucchi, E., Persichetti, E.: Post-quantum blind signature from standard group action assumptions and more. In: Daemen, J., Thomé, E. (eds.) EUROCRYPT 2026, Part I. LNCS, vol. 16541, pp. 277–307. Springer, Cham (May 2026). https://doi.org/10.1007/978-3-032-25291-3_10
37. Hess, F.: Weil descent attacks. In: Blake, I.F., Seroussi, G., Smart, N.P. (eds.) Advances in Elliptic Curve Cryptography, p. 151–180. London Mathematical Society Lecture Note Series, Cambridge University Press (2005)
38. Hess, F.: The GHS attack revisited. In: Biham, E. (ed.) EUROCRYPT 2003. LNCS, vol. 2656, pp. 374–387. Springer, Berlin, Heidelberg (May 2003). https://doi.org/10.1007/3-540-39200-9_23
39. Koussa, E., Macario-Rat, G., Patarin, J.: On the complexity of the permuted kernel problem. Cryptology ePrint Archive, Report 2019/412 (2019), https://eprint.iacr.org/2019/412
40. Leon, J.: Computing automorphism groups of error-correcting codes. IEEE Transactions on Information Theory **28**(3), 496–511 (1982), https://doi.org/10.1109/TIT.1982.1056498
41. Maurer, M., Menezes, A., Teske, E.: Analysis of the GHS Weil descent attack on the ECDLP over characteristic two finite fields of composite degree. In: Rangan, C.P., Ding, C. (eds.) INDOCRYPT 2001. LNCS, vol. 2247, pp. 195–213. Springer, Berlin, Heidelberg (Dec 2001). https://doi.org/10.1007/3-540-45311-3_19
42. Menezes, A., Qu, M.: Analysis of the Weil descent attack of Gaudry, Hess and Smart. In: Naccache, D. (ed.) CT-RSA 2001. LNCS, vol. 2020, pp. 308–318. Springer, Berlin, Heidelberg (Apr 2001). https://doi.org/10.1007/3-540-45353-9_23
43. Minc, H.: Upper bounds for permanents of $(0, 1)$-matrices. Bulletin of the American Mathematical Society **69**(6), 789–791 (1963)
44. National Institute of Standards and Technology: Post-quantum cryptography: Digital signature schemes. Round 2 Additional Signatures (2023), https://csrc.nist.gov/projects/pqc-dig-sig/round-2-additional-signatures
45. National Institute of Standards and Technology: Multi-Party Threshold Cryptography Project and NIST First Call for Multi-Party Threshold Schemes. https://csrc.nist.gov/projects/threshold-cryptography (2026), call for Multi-Party Threshold Schemes (NIST IR 8214C)
46. Overbeck, R.: Structural attacks for public key cryptosystems based on Gabidulin codes. Journal of Cryptology **21**(2), 280–301 (Apr 2008). https://doi.org/10.1007/s00145-007-9003-9
47. Petrank, E., Roth, R.M.: Is code equivalence easy to decide? IEEE Transactions on Information Theory **43**(5), 1602–1604 (1997), https://doi.org/10.1109/18.623157
48. Rhoades, B.: Increasing subsequences, matrix loci and viennot shadows. Forum of Mathematics, Sigma **12**, e97 (2024). https://doi.org/10.1017/fms.2024.75
49. Saeed, M.A.: Algebraic Approach for Code Equivalence. Ph.D. thesis, Normandie Université, University of Khartoum, (2017), Available at https://theses.hal.science/tel-01678829v2
50. Santini, P., Baldi, M., Chiaraluce, F.: Computational hardness of the permuted kernel and subcode equivalence problems. IEEE Trans. Inf. Theor. **70**(3), 2254–2270 (Mar 2024). https://doi.org/10.1109/TIT.2023.3323068, https://doi.org/10.1109/TIT.2023.3323068
51. Sendrier, N.: Finding the permutation between equivalent linear codes: the support splitting algorithm. IEEE Transactions on Information Theory **46**(4), 1193–1203 (2000). https://doi.org/10.1109/18.850662

52. Shamir, A.: An efficient identification scheme based on permuted kernels (extended abstract) (rump session). In: Brassard, G. (ed.) CRYPTO'89. LNCS, vol. 435, pp. 606–609. Springer, New York (Aug 1990). https://doi.org/10.1007/0-387-34805-0_54

A Experiments on Assumptions 1 and 2

We experimentally test the rank predictions in Assumptions 1 and 2 by measuring the deviation of the observed rank from the corresponding prediction. For each parameter set, we sample 100 independent instances.

For Assumption 1, we set the rate $R = k/n = 1/2$ and $1/3$, take $\ell = 1$, and choose $\nu = \lceil (n-1)^2/((n-k)\ell) \rceil$, corresponding to the boundary of the polynomial-time regime. Similarly, for Assumption 2, we test rates $R = 1/4, 1/3$ and $1/2$, and set ν as the minimal value satisfying the boundary for the polynomial-time regime, that is either $\nu = 6$ or $\nu = 5$. For each of the above settings, we test five values of n, with $20 \leq n \leq 64$. Notice that resulting the matrix $\bar{\boldsymbol{H}}$ has dimension between 400 and 4096.

For all experiments, the observed rank matched the predictions of Assumptions 1 and 2 exactly, i.e., with $\epsilon = 0$. We omit tables and plots because they would convey no information beyond this summary. We give the scripts used to run the experiments together with the raw data at [23]. Overall, the experiments provide strong empirical support on the validity of the assumptions.

B Experiments on PEP-to-GIP reduction

For the binary case, we validate Assumption 3 by sampling (n, k)-codes over $\mathbb{F}_{2^\nu}$ that are weakly self-orthogonal, using $k = n/2 - 1$ with various combinations of n and ν. The codes were generated with the algorithm of Albrecht *et al.* [4] (adapted to the field extension scenario), and in each case we calculated the fraction of experiments in which $\boldsymbol{G}^\sigma \boldsymbol{G}^\top$ is invertible. Figure 3 shows the clearly exponential decay of the failure probability. As for the odd characteristic setting, we present another series of experiments, this time using self-dual codes, over $\mathbb{F}_{3^\nu}$ for increasing ν, also shown in Figure 3.

To validate the assumption also for larger field characteristics, we used self-orthogonal codes over $\mathbb{F}_{p^2}$ and performed two series of experiments to support the claim that failures only occur when both $p \equiv 1 \bmod 4$ and $n \equiv 2 \bmod 4$ hold simultaneously. In the first series, we used only values $n \equiv 0 \bmod 4$ while letting p take values both $1 \bmod 4$ and $3 \bmod 4$. In the second series, we used only primes $p \equiv 3 \bmod 4$ and let n take values both $0 \bmod 4$ and $2 \bmod 4$. In all cases, we found that the probability of $\boldsymbol{G}^\sigma \boldsymbol{G}^\top$ not having full rank drops exponentially to zero as shown in Figure 4.

Finally, we conducted a last series of experiments to validate the claim that solving the GIP problem derived from Proposition 7 is likely to solve the underlying PEP problem as well. Selecting only the samples where $\boldsymbol{G}^\sigma \boldsymbol{G}^\top$ was already found to be invertible, we generated a random PEP instance and used

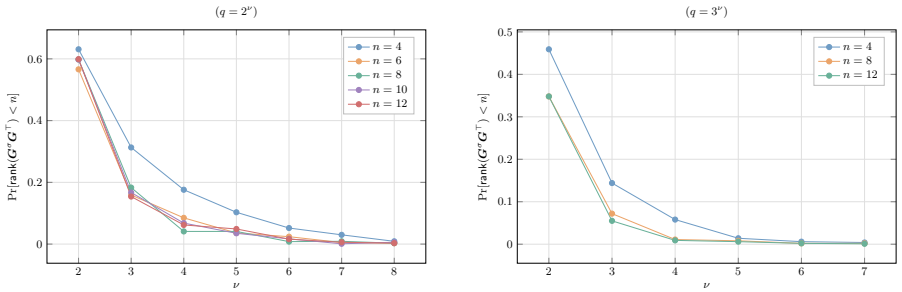

Fig. 3: Probability of $\boldsymbol{G}^\sigma \boldsymbol{G}^\top$ not having full rank for self-orthogonal codes in $\mathbb{F}_{2^\nu}^{k \times n}$ with $k = n/2 - 1$ (left) and for self-dual codes in $\mathbb{F}_{3^\nu}^{n/2 \times n}$ (right).

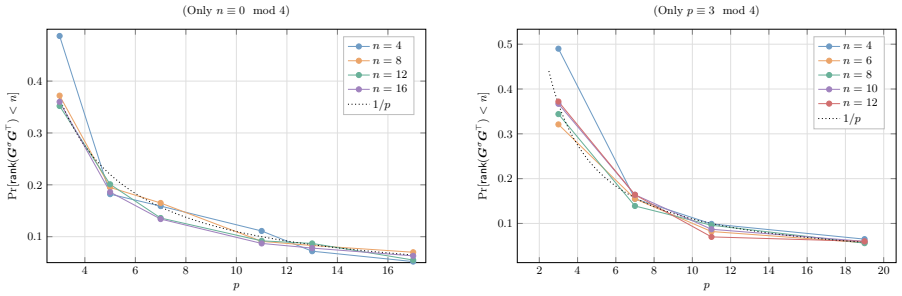

Fig. 4: Probability of $\boldsymbol{G}^\sigma \boldsymbol{G}^\top$ not having full rank for self-dual codes in $\mathbb{F}_{p^2}^{n/2 \times n}$. The experiments in the left-hand side figure only use $n \equiv 0 \mod 4$ and any value of p, while the experiments in the right-hand side figure only use $p \equiv 3 \mod 4$ and any even value of n. The dotted line represents the predicted $1/p$ complexity.

the algorithm described in Remark 2 to solve the associated GIP problem, then measured the probability that this solution also solves the original PEP problem. The results from Figure 5 show the reduction failing only in very small instances (due to the GIP problem not having a unique solution), but quickly achieving 100% success rate as parameters grow.

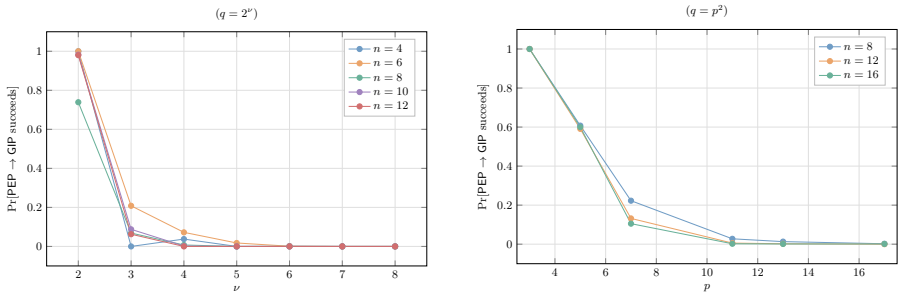

Fig. 5: Failure probability of the PEP to GIP reduction for self-orthogonal codes in $\mathbb{F}_{2^\nu}^{k\times n}$ with $k = n/2 - 1$ (left) and for self-dual codes in $\mathbb{F}_{p^2}^{n/2\times n}$ (right).