A Geometric Approach to Quantum Distinguishers

Zhili Wu[1] and Zhenzhen Bao[1,2] ✉

[1] Institute for Network Sciences and Cyberspace, Tsinghua University, Beijing, China
wuzl24@mails.tsinghua.edu.cn,zzbao@tsinghua.edu.cn
[2] State Key Laboratory of Cryptography and Digital Economy Security, Tsinghua University, Beijing, 100084, China

Abstract. This paper introduces a geometric framework for Q2 quantum distinguishers by combining the geometric approach to classical symmetric-key cryptanalysis with the generalized correlation extraction algorithm. Our main technical tool shows that one superposition query, followed by appropriate (unitary) change-of-basis operations, prepares a "correlation state" whose amplitudes are the entries of the geometric correlation matrix in the chosen basis. This yields a unified preparation-measurement template that recovers several known quantum distinguishers: (1) hidden structure detection via support constraints in Fourier-type bases (e.g., Simon, Bernstein-Vazirani, Deutsch-Jozsa), and (2) event probability deviation tests via amplitude estimation (covering standard quantizations of linear and differential distinguishers). We analyze when relevant distinguishing mass is diluted across many basis states, identify it as a cause of poor query efficiency in several recent distinguishers, and provide basis-specific mechanisms to concentrate the signal (phase-oracle row restriction in the Fourier setting; chosen-plaintext subset-state restriction in the quasidifferential setting) to restore quadratic advantage. We illustrate the framework on Fourier and quasidifferential instantiations and discuss obstacles for non-unitary integral bases.

Keywords: Quantum attacks · Geometric approach · Correlation extraction algorithm · Symmetric cryptanalysis · Block ciphers

1 Introduction

The advent of quantum computing poses a significant threat to the security of existing cryptographic primitives. A primary concern is the quantum speed-up, wherein quantum algorithms can compromise ciphers using far fewer queries than their classical counterparts. Leveraging such speed-ups, several foundational algorithms have demonstrated the vulnerability of classical cryptography: Grover's algorithm [33] achieves a quadratic speed-up for generic search problems; Shor's algorithm [56] solves the discrete logarithm problem, thereby challenging the foundations of public-key cryptography; Simon's algorithm [57] has been used to break multiple symmetric-key constructions, such as 3-round Feistel (in [45]),

4-round Feistel (in [41]), Even-Mansour [30] (in [46]), FX construction [44] (in [47]) and many others following the seminal work [42].

In this paper, we focus on quantum distinguishers for block ciphers. Rather than attempting to resolve all open questions, we begin by articulating three questions that we find particularly persistent:

1. *Which quantum algorithm should one use?* Quantum cryptanalysis employs particular algorithm primitives as subroutines, which is the power source of their speed-up. A growing toolkit of algorithm primitives has emerged in the past decades, forming a nontrivial task for cryptanalysts to choose one from them. Notable examples include Grover's algorithm [33] for unstructured dataset search, Simon's algorithm [57] for the hidden period problem, Shor's algorithm [56] for the discrete logarithm problem, BHT algorithm [21] for collision search, Deutsch-Josza algorithm [26] for distinguishing balanced or constant Boolean functions, Bernstein-Vazirani algorithm [4] for distinguishing linear Boolean functions, Ambainis' algorithm [1] for element distinctness, quantum counting and amplitude amplification and estimation [20,19], quantum random walks [60,49], quantum rejection sampling [52], the Grover-meet-Simon algorithm [47] and offline Simon's algorithm [13]. We recommend [17] for a recent survey of quantum cryptanalysis of block ciphers, which lists more useful algorithms.
2. *How should a quantum attack be designed?* Designing an attack is a challenge for cryptanalysts on both classical and quantum computers. Apart from Simon-like attacks, for which many inspiring theories and useful tools have been proposed in recent years [27,22,62,48], we want to note the quantum adaptions of classical attack: [43] quantizes differential attacks by Grover's algorithm, truncated differential attacks by Grover's algorithm and Ambainis' algorithm, and linear attacks by quantum counting. It depends on the feature of cryptanalytic techniques to choose a quantum algorithm, making it an ad-hoc procedure; [31] speeds up boomerang and mixing boomerang attacks using Grover's algorithm, QAA and Ambainis' algorithm (needs qRAM); [37] quantizes multidimensional (zero-correlation) linear attacks and integral attacks based on balanced properties using QAA and a dedicated correlation extraction algorithm.
   - Previous works demonstrate that cryptanalysis is more than quantum algorithm choices. For example, linear attacks are quantized by [43] via directly counting parities and by [37] via a Simon-like subroutine called the correlation extraction algorithm (CEA). The two approaches share the same asymptotic query complexity but differ in non-asymptotic query complexity, hinting their difference in nature.
   - Compared to Simon's algorithm, several algorithms have only recently been considered in cryptanalysis, and a systematic understanding would be helpful. For example: [63] applies Bernstein-Vazirani algorithm combined with a Simon-like subroutine; [38] applies Shor's algorithm to speed up computation of empirical correlations in fast correlation attacks; [14] applies Bernstein-Vazirani, Deutsch-Josza, Simon's and Shor's algorithm

in quantum linearization attacks. This calls for general insights across classical and quantum algorithms.

3. *What speed-up can actually be achieved?* It is often said that Grover's algorithm yields a quadratic speed-up and Simon's algorithm an exponential one. But when these quantum algorithms are deployed against real ciphers, the picture becomes far more nuanced.
   - Speed-ups of Grover's algorithm vary from none to more-than-quadratic: (1) [43] achieves a quadratic speed-up for differential and linear cryptanalysis, but the speed-up for truncated differential attacks are usually smaller. (2) [16] attacks 2XOR-Cascade construction using the offline Simon's algorithm, gaining a 2.5 speed-up. (3) The basic distinguishers ($\mathcal{A}_1$ and $\mathcal{B}_1$) in [37] have strictly no speed-up, while its advanced distinguishers achieve a more- or less-than-quadratic speed-up, depending on attack parameters. Thus, we need more cryptanalytic insights than merely doubling the key size for mitigation.
   - Simon's algorithm is often held up as a textbook example of exponential quantum speed-up, yet achieving linear query complexity requires considerably more ingenuity than a direct invocation of the algorithm: (1) [63] identifies hidden linear structures with complexity $\Theta(n^2)$ using a Simon-like subroutine with phase oracles; (2) [35] classifies strong, semi-strong and weak separability properties for general Feistel networks and show that the strong and semi-strong separability is suitable for Simon's algorithm. The weak separability is initially believed to require $O(n) \cdot 2^n$ queries; this is later improved to linear by [48] via a general model that sums up results of two related calls.

While quantum cryptanalysis has produced many powerful individual algorithms, a unifying framework that connects them remains underdeveloped. In contrast, the classical cryptanalysis community has recently seen the emergence of the geometric approach [6] as a unifying framework for multiple types of attacks. Empowered by algebraic geometry, this theoretical framework models a cipher as a linear operator on high-dimensional free vector spaces and studies its transition matrix under carefully chosen bases, providing an intrinsic algebraic description that unifies classical techniques such as linear [5], differential [8], and integral [9] cryptanalysis. The geometric approach not only reveals deep connections among different attacks but has also led to novel theories for attack evaluation in the fixed-key model, encouraged new analytical techniques (e.g., ultrametric integral cryptanalysis [10]) and facilitated the development of automated search tools. Recently, the geometric approach has witnessed improvements, such as d-differential [61], mixed-basis [40], and other extensions [11,7,24,18]. Yet, the potential connections between the geometric approach and quantum cryptanalysis have not been fully explored, leaving a significant gap in our understanding of how this powerful unifying classical framework could inform or systematize the study of quantum distinguishers.

Based on this observation, we attempt to introduce the geometric approach into quantum cryptanalysis to provide a systematic perspective for understand-

ing and designing quantum distinguishers. Specifically, we focus on three key points: (1) the compatibility of geometric bases with quantum computing, i.e., when the number field is chosen to be the complex numbers (field $\mathbb{C}$), the geometric bases are orthogonal and naturally correspond to quantum gate operations; (2) the generalization of the correlation extraction algorithm introduced in [37], which can prepare the correlation state of a function in a single quantum query, providing a unified entry point for subsequent structure detection or probability deviation hypothesis testing; and (3) the cryptographic interpretation of geometric bases and their quantum counterparts, including Fourier bases, quasidifferential bases, and ultrametric integral bases.

Through this geometric perspective, we contribute partial but concrete answers to the three questions posed above. (1) For the choice of quantum algorithms, we recommend the quantized geometric approach, observing that its Fourier, quasidifferential and integral bases each correspond to a distinct family of quantum algorithms; the CEA makes their common structure explicit and reveals that many seemingly unrelated attacks (e.g. Simon, Bernstein-Vazirani, Deutsch-Jozsa, quantum linear/differential distinguishers) are in fact instances of the same preparation-measurement template. (2) For attack design, this template yields a modular strategy: first prepare the correlation state of the function (using one query), then probe it either for zero-correlation supports or for biased amplitudes. The geometric CEA (Theorem 7) unifies the two hitherto separate approaches of [43] and [37]. (3) For the achievable speed-up, the framework pinpoints a major cause of poor query complexity in several recent distinguishers, namely, the *dilution* of the discriminating signal across many irrelevant basis states. By analyzing how each geometric basis distributes this signal, we formalize two principled countermeasures (phase oracles for the Fourier basis, subset-state preparation for the quasidifferential basis) that recover the expected quadratic advantage and, in some non-asymptotic regimes, even improve upon previous works.

Nevertheless, the present work is primarily a theoretical proposal; many concrete instantiations and optimizations remain to be explored. For instance, the geometric CEA has only been demonstrated for a limited number of bases, each corresponding to a class of attacks. Extending it to new attacks would both test the classical geometric framework's generality and potentially yield new quantum distinguishers. Moreover, the interplay between geometric bases and advanced quantum algorithms (e.g. quantum random walks) is still largely uncharted. We hope that the clarity offered by the geometric viewpoint will facilitate the future discovery of such extensions.

Related works. [5,8,9,10,6] introduce the geometric approach to classical symmetric-key cryptanalysis. In the quantum setting, Kaplan et al. [43] pioneer the use of Grover search and quantum counting to accelerate linear and differential distinguishers, while Hosoyamada [37] introduces the Correlation Extraction Algorithm (CEA) for multidimensional linear distinguishers. Simon's algorithm [57] and its cryptographic applications (e.g. [42,48]) demonstrate exponential speed-ups via hidden period detection. Besides Simon's, this paper discusses

some other quantum algorithms [33,20,19,26,4,52]. Several cryptanalytic works [14,15,39,40,63] notably encourage the paper's development in various aspects.

Contributions. The main contributions of this paper are as follows:

1. We propose a geometric framework for quantum cryptanalysis by generalizing the *Correlation Extraction Algorithm* (CEA, Theorem 7), introduced by [37], in the light of the classical geometric approach with number field $\mathbb{C}$, providing a unified foundation for constructing quantum distinguishers. With a single superposition query, the CEA prepares a correlation state $|\Psi_F\rangle$ with support from the geometric correlation matrix C^F of function F. Building on a unified preparation-measurement template for $|\Psi_F\rangle$, we then systematize quantum distinguishers into two primary categories (Subsection 3.1): *hidden algebraic structure detection* (e.g., Simon's algorithm) for sequentially testing zero-correlation properties on $|\Psi_F\rangle$, whose speed-up depends on the specific property, and *event probability deviation hypothesis testing* using amplitude estimation on $|\Psi_F\rangle$, which typically achieves a quadratic speed-up (Theorem 9). It thereby offers a systematic algebraic toolkit for analyzing the source of quantum speed-ups and designing new distinguishers within a coherent geometric perspective.
2. We investigate three (possible) instantiations of the CEA under different bases, discussing the geometric origin and limitations of quantum speed-ups. For the Fourier basis (Section 4), we show that Simon's, Bernstein-Vazirani's, and Deutsch-Josza's algorithms correspond to testing "zero-correlation" hypotheses by sequential sampling from $|\Psi_F\rangle$. For the quasidifferential basis (Subsection 5.1), we reinterpret the quantum differential distinguisher of [43] as sampling from the chosen-plaintext (Theorem 8) quasidifferential correlation matrix, using a new property of the quasidifferential basis (Theorem 16), and suggest a barrier on the achievable speed-up for its applications. For the ultrametric integral basis (Subsection 5.2), we note its incompatibility with quantum mechanics due to its non-Euclidean metric. These analyses reveal the intrinsic connection between quantum acceleration and the algebraic structure of the chosen basis.
3. We observe the "dilution" problem in $|\Psi_F\rangle$ when using the CEA (Theorem 10). Typically only a tiny fraction of its amplitude contributes to the distinguisher, while the rest are non-informative. Due to the unitary nature of quantum mechanics, the non-informative states cannot be abandoned like classical computing. This leads to poor query complexity (e.g. distinguishers $\mathcal{A}_1, \mathcal{B}_1$ of [37]). We propose two generic strategies to suppress non-informative components, each tailored to a geometric basis. (1) For Fourier basis, the (generalized) phase oracles restrict the output mask subspace, eliminating unwanted rows of the correlation matrix (Theorem 12). This also explains Bernstein-Vazirani and Deutsch-Josza algorithms. Additionally, we prove that the phase oracle technique is applicable for Fourier bases only (Theorem 14), demonstrating its uniqueness. (2) For quasidifferential basis, chosen-plaintext techniques (Theorem 8) replace the uniform superposition

by a "subset state". The resulting CEA then directly amplifies the differential probability.

Organization of the paper. Section 2 introduces preliminaries about quantum cryptanalysis and revisit the classical geometric approach. Section 3 proposed the geometric correlation extraction algorithm. Section 4 discusses the Fourier basis and its related quantum algorithms. Section 5 discusses the quasidifferential and ultrametric integral bases. Section 6 concludes the paper.

2 Preliminaries

2.1 Notations

Equality symbol $:=$ represents definition. $\mathbb{F}_2$ denotes the Galois field from elements $\{0,1\}$, with modulo-two addition (exclusive-or, XOR) and multiplication. $\mathbb{F}_2^n$ denotes the n-dimensional vector space over $\mathbb{F}_2$. The addition in $\mathbb{F}_2^n$ is bitwise XOR, denoted by $\oplus$. The inner product for two vectors a, b in $\mathbb{F}_2^n$ is defined by $a \cdot b := \bigoplus_{i=1}^n a_i b_i$, where a_i is the i-th component (1-indexed) of vector a. Vector Boolean functions (or functions in short by abuse) are functions of the form $F : \mathbb{F}_2^n \to \mathbb{F}_2^m$, while Boolean functions have $m = 1$. $a||b$ denotes concatenation of vectors a, b. For functions F_1, F_2, their bricklayer concatenation $F_1 \otimes F_2$ is a function mapping $x_1||x_2$ to $F_1(x_1)||F_2(x_2)$. F^{-1} denotes the inverse of an invertible function F. Preimages of $y \in \mathbb{F}_2^m$ for $F : \mathbb{F}_2^n \to \mathbb{F}_2^m$ are denoted by $F^{-1}(y) := \{x \in \mathbb{F}_2^n : F(x) = y\}$.

$\mathbb{C}$ denotes the complex number field. For a complex number a, its conjugate is a^*, squared norm $|a|^2 := a^* a$. For a vector or matrix a over $\mathbb{C}$, its conjugate is a^*, transpose $a^\top$, Hermitian conjugate $a^\dagger$. H denotes the Hadamard matrix (Hadamard gate). Pauli gates are denoted by I, X, Y, Z. For matrices M_1, M_2, their Kronecker tensor product is denoted by $M_1 \otimes M_2$.

With bra-ket notation, $|a\rangle$ represents a standard basis vector with index a in a free vector space, and its dual basis vector is $\langle a|$. The inner product $\langle a \mid b\rangle$ equals 1 if $a = b$ and 0 otherwise. Kronecker tensor products of vectors are denoted by $|a\rangle \otimes |b\rangle$ or $|a\rangle |b\rangle$ for short, while their adjoint is $\langle b| \langle a|$. For a standard basis vector $|x\rangle |y\rangle$, we sometimes abbreviate it to $|x, y\rangle$ and its adjoint to $\langle x, y|$. For simplicity, we define the subset state $|S\rangle$ given $S \subseteq \mathbb{F}_2^n$ by $|S\rangle := |S|^{-1/2} \sum_{x \in S} |x\rangle$.

$\Pr[\cdot]$ denotes probability. $\mathbb{E}[\cdot]$ stands for expectation. $\mathrm{Bern}(\cdot)$ stands for a Bernoulli distribution. $D_{\mathrm{KL}}(\cdot||\cdot)$ denotes the Kullback-Leibler divergence for two distributions. $O(\cdot), \Theta(\cdot), \Omega(\cdot)$ are standard asymptotic complexity symbols.

In this paper, matrices use capital letters (e.g. A, B, C); vectors use small letters (e.g. x, y, z) and greek letters (e.g. λ, Δ); change-of-basis operators use calligraphic letters (e.g. $\mathcal{F}, \mathcal{Q}, \mathcal{U}$); functions use capital letters (e.g. F, G, H).

For function $F : \mathbb{F}_2^n \to \mathbb{F}_2^m$, the linear approximation table (LAT) is defined by

$$\mathrm{LAT}_F(\lambda_I, \lambda_O) := 2^{-n} \sum_{x \in \mathbb{F}_2^n} (-1)^{\lambda_I \cdot x \oplus \lambda_O \cdot F(x)}. \tag{1}$$

The differential distribution table (DDT) is defined by

$$\mathrm{DDT}_F(\Delta_I, \Delta_O) := 2^{-n} \sum_{x \in \mathbb{F}_2^n} \langle \Delta_O \mid F(x) \oplus F(x \oplus \Delta_I) \rangle \,. \tag{2}$$

2.2 Quantum Computing

In this subsection, we introduce important settings of quantum cryptanalysis and recite several known quantum algorithms for later use. Readers are assumed to have elementary knowledge of quantum computing. We refer readers to textbooks like [51] for fundamental concepts including qubits, gates and measurements.

Quantum oracles. Quantum oracles allow superposition queries to a function. We are concerned of two types of oracles: the standard XOR oracle and the phase oracle. Here we follow the interpretation of [15].

Definition 1 (Standard oracle). *The standard oracle for function $F : \mathbb{F}_2^n \to \mathbb{F}_2^m$ is an operator O_F that maps $|x\rangle |y\rangle$ to $|x\rangle |y \oplus F(x)\rangle$.*

Definition 2 (Phase oracle). *The phase oracle for function $F : \mathbb{F}_2^n \to \mathbb{F}_2^m$ is an operator $O_F^{\pm}$ that maps $|x\rangle |\lambda\rangle$ to $(-1)^{\lambda \cdot F(x)} |x\rangle |\lambda\rangle$.*

[15] points out that both oracles are equivalent by composing with a Hadamard transform (quantum Fourier transform):

$$O_F^{\pm} = (I^{\otimes n} \otimes H^{\otimes m}) \cdot O_F \cdot (I^{\otimes n} \otimes H^{\otimes m}). \tag{3}$$

Note that Definition 2 already implies vector Boolean functions, although its applications for "hidden shift problems" in [15] have $m = 1$. The phase oracle will be revisited with a linear attack perspective in Subsection 4.1.

Q1/Q2 adversaries. Quantum attacks fall into two categories according to [43]: Q1 and Q2. A Q1 adversary collects data from classical oracles and then computes on a quantum computer, while a Q2 adversary has access to quantum oracles besides quantum computers. The Q2 adversary model, while requiring stronger assumptions, provides a powerful framework for exposing cryptographic vulnerabilities. It enables exponential speed-ups via quantum algorithms such as Simon's algorithm. This paper focuses on the Q2 setting, since the CEA requires one superposition query per call.

Quantum memory. In the design of quantum algorithms, quantum memory, particularly quantum random-access memory (qRAM), constitutes a key computational model that allows algorithms to efficiently access data stored in memory in superposition. Two primary models are commonly discussed: the quantum random access to classical memory (QRACM), where classical memory supports superposition-state addressing, and the quantum random access to quantum memory (QRAQM), which enables superposition-state read and write operations (typically through a "data swap" operator) on quantum memory itself. It

is notable that many algorithms (e.g. Ambainis's [1]) based on quantum random walks require efficient implementation of QRAQM to reduce the time complexity. Therefore, the qRAM reliance is a critical factor in evaluating the practicality of quantum algorithms. As for this paper, our algorithms do not rely on qRAM.

Attack complexities. As with classical cryptanalysis, quantum attacks are characterized by three main complexity measures: (1) time complexity: the number of quantum gates in the circuit. in practice, circuit depth and Toffoli gate counts also serve as key metrics. (2) memory complexity: the number of qubits needed. (3) query complexity: the number of (superposition) queries to the oracle. Additionally, the cost of any non-negligible classical operations should be incorporated into these complexities. For Q2 attacks, which is the focus of this paper, the time and memory complexities are typically dominated by quantum oracle calls. Thus, in most cases, it suffices to analyze the query complexity alone.

2.3 Quantum Algorithms

This subsection presents a brief introduction to several well-known quantum algorithms, namely Grover's algorithm [33], amplitude amplification [20,19], Simon's algorithm [57], quantum counting and amplitude estimation [20,19]. Other algorithms like Deutsch-Josza algorithm [26] and Bernstein-Vazirani algorithm [4] will be discussed in Section 4 due to their reliance on Fourier basis (Hadamard gates).

Grover's algorithm. Grover's algorithm is a quantum search method for unstructured databases, achieving a quadratic speed-up over classical approaches when given access to a quantum oracle. Formally, consider a database of size 2^n whose entries are indexed by bitstrings of length n. Let $F : \mathbb{F}_2^n \to \mathbb{F}_2$ be a Boolean function that marks the sought items, i.e., $F(x) = 1$ if and only if x corresponds to a marked entry. The goal is to find any x such that $F(x) = 1$, given oracle O_F. In other words, suppose the states are partitioned into "good" states (where $F(x) = 1$) and "bad" states (where $F(x) = 0$), our goal is to find a good state from measurement. Classically, one must query the oracle an average of $\Theta(1/p)$ times to succeed with constant probability, where $p = |F^{-1}(1)|/2^n$. By contrast, Grover's algorithm needs only $O(1/\sqrt{p})$ quantum queries, as formally stated in Theorem 1.

Theorem 1 (Grover's algorithm [33]). *Given a quantum oracle of a Boolean function $F : \mathbb{F}_2^n \to \mathbb{F}_2$, there exists an algorithm that finds a preimage of 1 with high probability, with query complexity $O(1/\sqrt{p})$, where $p := |F^{-1}(1)|/2^n$.*

The procedure of Grover's algorithm is as follows:

1. Prepare the initial state $|0\rangle^{\otimes n}$.
2. Apply the n-qubit Hadamard transform $H^{\otimes n}$ and obtain state $|\varphi_0\rangle := |+\rangle^{\otimes n}$.
3. Repeat the following subroutine (Grover operator) for r times:
   - Apply phase oracle $O_F^{\pm}$ on the state.

- Reflect the state about $|\varphi_0\rangle$. This is done by three subsequent operations: $H^{\otimes n}$, a reflection about $|0\rangle^{\otimes n}$, and $H^{\otimes n}$ again.

4. Measure the state by the computational basis.

The Grover iteration in Step 3 is the core of the algorithm. Each iteration consistently *increases* the amplitudes of all good states and *decreases* those of all bad states. Precisely, the optimal number of iterations $r \approx \pi/(4\sqrt{p})$, with success probability nearly 1. See [33] for details.

Quantum amplitude amplification. Quantum amplitude amplification (QAA) [20,19] generalizes Grover's search. The procedure is identical to Grover's algorithm, except that the Hadamard transform $H^{\otimes n}$ is replaced by an arbitrary unitary U, producing the initial state $|\varphi_0\rangle := U|0\rangle$. As in Grover's algorithm, marked states are amplified through iterative applications of a reflection about the initial state and a phase oracle.

Theorem 2 (Amplitude amplification [20,19]). *Let $F : \mathbb{F}_2^n \to \mathbb{F}_2$ be a Boolean function, U be a unitary operator and $|\varphi_0\rangle := U|0\rangle$. Then amplitude amplification finds a preimage of 1 with high probability using $O(1/\sqrt{p})$ queries to the oracle for F and to U and $U^\dagger$, where $p := \sum_{x \in F^{-1}(1)} |\langle x|\varphi_0\rangle|^2$.*

The algorithm succeeds with high probability when the number of iterations r is chosen appropriately. The optimal choice of r depends on whether the value of p is known. Following [37,38], we distinguish two cases: (1) if p is known, set $r = \lfloor \pi/(4 \arcsin(\sqrt{p})) \rfloor$; (2) if p is unknown, it suffice to take $r \leq (9/2)/\sqrt{p}$ for $p \leq 3/4$.

Simon's algorithm. Simon's algorithm [57] provides a polynomial-query quantum solution to the hidden period problem. A formal revisit based on the geometric approach will be presented in Subsection 4.2.

Definition 3 (Hidden period problem, HPP). *Given an oracle of a function $F : \mathbb{F}_2^n \to \mathbb{F}_2^m$ that satisfies condition*

$$F(x) = F(y) \Leftrightarrow x \oplus y \in \{0, s\} \tag{4}$$

for a secret vector $s \in \mathbb{F}_2^n$, find s.

Classically, solving this problem requires querying the function until a collision is found, which takes $\Theta(2^{m/2})$ queries. In contrast, Simon's algorithm requires only a linear number of quantum queries.

Theorem 3 (Simon's algorithm [57]). *There exists an algorithm to solve the hidden period problem using $O(n)$ queries to the quantum oracle of F.*

The procedure of Simon's algorithm is described as follows:

- Prepare state $|0\rangle^{\otimes n}|0\rangle^{\otimes m}$ on $n+m$ qubits.
- Apply $H^{\otimes n} \otimes I^{\otimes m}$.

– Apply oracle O_F . (Optional: measure the last m registers.)
– Apply $H^{\otimes n} \otimes I^{\otimes m}$ again.
– Measure the first n registers.

A significant practical constraint of Simon's algorithm is that the function must have no unwanted collision (x, y) such that $x \oplus y \notin \{0, s\}$. To address this issue, [42] shows that a bounded probability (not too close to 1) of unwanted collisions can be tolerated, while [41] further relaxes this requirement in distinguishing settings by analyzing the underlying vector-space structure of candidates of the secret s.

Quantum counting and amplitude estimation. Quantum counting and amplitude estimation [20,19] estimates the proportion of marked items without necessarily retrieving them. These techniques rely on a controlled version of the Grover operator (or the generalized amplitude amplification operator), which can be implemented efficiently given the underlying oracles.

Theorem 4 (Quantum counting [20,19]). *Given a quantum oracle for a Boolean function $F : \mathbb{F}_2^n \to \mathbb{F}_2$, there exists an algorithm that outputs an estimate $\tilde{p}$ of $p := |F^{-1}(1)|/2^n$ such that*

$$|p - \tilde{p}| \leq \frac{2\pi\sqrt{p(1-p)}}{N} + \frac{\pi^2}{N^2} \tag{5}$$

with probability at least $8/\pi^2$, using N queries to the oracle.

The earlier work [20] does not have the factor $\sqrt{1-p}$ in the bound, and it turns out to be sufficient for most discussions within this paper. Note that the success probability can be amplified via a constant number of repetitions. Besides the original literature, [38] shows that the circuit can be implemented with depth at most N times the circuit depth of O_F and n auxiliary qubits, settling further issues for implementation.

Quantum amplitude estimation (QAE) generalizes quantum counting in the same way that amplitude amplification generalizes Grover's search: it estimates the probability of obtaining a "good" outcome from an arbitrary quantum circuit, not just from the uniform superposition.

Theorem 5 (Amplitude estimation [20,19]). *Let $F : \mathbb{F}_2^n \to \mathbb{F}_2$ be a Boolean function and let U be a unitary operator. Define the initial state $|\varphi_0\rangle := U|0\rangle$ and let $p := \sum_{x \in F^{-1}(1)} |\langle x|\varphi_0\rangle|^2$. Then there exists an algorithm that outputs an estimate $\tilde{p}$ of p such that*

$$|p - \tilde{p}| \leq \frac{2\pi\sqrt{p(1-p)}}{N} + \frac{\pi^2}{N^2} \tag{6}$$

with probability at least $8/\pi^2$, using N queries to unitaries O_F, U and $U^\dagger$.

When $|p - \tilde{p}|$ is dominated by the first term $2\pi\sqrt{p(1-p)}/N$, a quadratic speed-up is achieved compared to the classical Monte-Carlo approach. However, this is merely a rough intuition because p and the error $|p - \tilde{p}|$ control whether N is large, so they need to be considered as well. We present a rigorous evaluation in Theorem 9.

2.4 The Geometric Approach

In this subsection, we introduce the geometric approach to classical cryptanalysis over the number field $\mathbb{C}$. While the geometric framework accommodates arbitrary algebraically closed fields, we restrict to $\mathbb{C}$ in this work to align with the Hilbert space formalism of quantum mechanics. We follow the theoretical developments of [5,8,9,10,6] with minor adaptations, including vector normalization and bra-ket notation. These works not only provide a unified perspective on a wide range of classical distinguishing attacks, but also rigorously analyze trail correlations and approximations in the fixed-key model, providing a novel and powerful framework for attack evaluation. This paper focuses primarily on the former aspect; for the latter, we refer readers to the original literature and recent advances [61,11,7,24,18,40], since the latter aspect draws another seminal line of research.

Definition 4 (Pushforward operator [6]). *Let $F : \mathbb{F}_2^n \to \mathbb{F}_2^m$ be a function. The pushforward operator along F is a linear map from $\mathbb{C}^{2^n}$ to $\mathbb{C}^{2^m}$:*

$$T^F := \sum_{x \in \mathbb{F}_2^n} |F(x)\rangle \langle x| . \tag{7}$$

Its adjoint is the pullback operator, a linear map from $\mathbb{C}^{2^m}$ to $\mathbb{C}^{2^n}$:

$$T^{F\vee} := \sum_{x \in \mathbb{F}_2^n} |x\rangle \langle F(x)| . \tag{8}$$

Note that the term "operator" in Definition 4 means algebraic pushforward matrices but not unitary operators in quantum mechanics.

The pushforward (and pullback) operators mirror the structure of function composition and parallel execution:

Theorem 6 (Properties of pushforward operators [6]).

1. (Composition) *If $F = F_1 \circ F_2$, then $T^F = T^{F_1} T^{F_2}$.*
2. (Bricklayer) *If $F = F_1 \otimes F_2$, then $T^F = T^{F_1} \otimes T^{F_2}$.*

This rule is vital for trail search and further applications (see [6, Chapter 2] for details). We note that representing Boolean functions via transition matrices also appears in the study of Boolean control networks [25].

From a representation-theoretic viewpoint, the set of functions $F : \mathbb{F}_2^n \to \mathbb{F}_2^n$ forms a monoid, and $F \mapsto T^F$ constitutes its associated representation (see textbooks like [32,58]). Specifically, [6, Chapter 2] studies the left action of a monoid on itself, termed as the regular representation. This algebraic perspective underlies the geometric approach.

Since T^F is a linear operator, we may change the bases of its domain and codomain to obtain a different matrix representation.

Definition 5 (Correlation matrix [6]). *Let $F : \mathbb{F}_2^n \to \mathbb{F}_2^m$ be a function, and let P_I and P_O be invertible matrices representing changes of basis for the input and output spaces, respectively. The correlation matrix of F with respect to these bases is*

$$C^F := P_O T^F P_I^{-1}. \tag{9}$$

Definition 5 permits distinct input and output bases, which is a technique termed "mixed-basis", first introduced by [40]. In contrast, the original "same-basis" model in [6] assumes $n = m$ and $P_I = P_O$.

The choice of bases is not arbitrary; we seek bases that simultaneously simplify the correlation matrices for a family of functions. In the geometric approach, the preferred bases (in the same-basis model) are constructed from characters of a commutative (inverse) monoid acting on the state space [6, Section 2.4]. Over $\mathbb{C}$, these character bases are orthonormal with respect to the Euclidean inner product, ensuring that the change-of-basis matrices P_I, P_O are unitary.

Henceforth, we restrict to cases where P_I and P_O are unitary, i.e., the bases are orthonormal. This ensures that the correlation matrix C^F inherits the unitarily invariant quantities of T^F (e.g. the Frobenius norm of T^F and C^F is $2^{n/2}$) and facilitates the quantum analogies developed later in this paper. Concretely, Section 4 and Subsection 5.1 show that P_I and P_O can be efficiently implemented as quantum gates for the bases we study in this paper. Thus, Equation (9) can be rewritten as

$$C^F = P_O T^F P_I^\dagger. \tag{10}$$

Definition 6 (One-dimensional approximation [6]). *Given a function $F : \mathbb{F}_2^n \to \mathbb{F}_2^m$, an approximation is an ordered pair $(|u\rangle, |v\rangle)$, where $|u\rangle \in \mathbb{C}^{2^n}, |v\rangle \in \mathbb{C}^{2^m}$. The correlation of approximation $(|u\rangle, |v\rangle)$ is defined as $\langle v \mid T^F \mid u\rangle$.*

In this paper, we usually choose approximation $(|u\rangle, |v\rangle)$ from basis vectors in bases $P_I^\dagger$ and $P_O^\dagger$, i.e., $(|u\rangle, |v\rangle) = (P_I^\dagger |\lambda_I\rangle, P_O^\dagger |\lambda_O\rangle)$ for some $\lambda_I \in \mathbb{F}_2^n, \lambda_O \in \mathbb{F}_2^m$. The correlation is then given by

$$\langle v \mid T^F \mid u\rangle = \langle \lambda_O \mid C^F \mid \lambda_I\rangle = \sum_{x \in \mathbb{F}_2^n} \langle \lambda_O \mid P_O \mid F(x)\rangle \langle x \mid P_I^\dagger \mid \lambda_I\rangle . \tag{11}$$

Example 1. We obtain LAT_F by setting $P_I = H^{\otimes n}, P_O = H^{\otimes m}$:

$$\langle \lambda_O \mid C^F \mid \lambda_I\rangle = \sum_{x \in \mathbb{F}_2^n} 2^{-(n+m)/2} (-1)^{\lambda_I \cdot x \oplus \lambda_O \cdot F(x)} := 2^{(n-m)/2}\, \mathrm{LAT}_F(\lambda_I, \lambda_O). \tag{12}$$

The index for columns and rows of C^F are indeed characters, which is denoted by χ in tradition. However, since this paper does not intend to dive deep into representation theory, we denoted them by λ to make it easier for common readers. We refer interested readers to [6].

3 Correlation Extraction Algorithm

This section presents the unified preparation-measurement template based on (generalized) correlation extraction algorithm (CEA) and explains its application in constructing quantum distinguishers. We observe and analyze the "dilution" problem in Subsection 3.3.

The CEA is first introduced by [37], as a quantum algorithm for multidimensional (zero-correlation) properties, inspired by Simon's algorithm. Formally, the CEA corresponds to Simon's algorithm without the final measurement, followed by the application of Hadamard gates $H^{\otimes m}$ on the last m registers. This modification is far from trivial: while Simon's algorithm relies heavily on quantum interference to suppress states λ that do not satisfy $\lambda \cdot s = 0$, the CEA effectively prepares the linear approximation table (LAT) of a vector Boolean function, thereby revealing substantially more information than Simon's algorithm. Although prior work [22] had already noted connections between Simon's algorithm and linear properties of functions, its scope was limited to Boolean functions, i.e. functions with single-bit outputs. Since LATs are special cases of the geometric approach (Example 1), the CEA is indeed more general than a multidimensional linear distinguisher. In this section, we propose a generalized CEA framework that unifies the understanding of quantum attacks, including Simon's algorithm, within the geometric approach to classical symmetric-key cryptanalysis. This framework naturally applies to vector Boolean functions.

Theorem 7 (Correlation extraction algorithm, generalized with the geometric approach). *Given an oracle O_F of a function $F : \mathbb{F}_2^n \to \mathbb{F}_2^m$, computing $(P_I^* \otimes P_O)O_F \left|+\right\rangle^{\otimes n} \left|0\right\rangle^{\otimes m}$ yields*

$$\left|\Psi_F\right\rangle := \sum_{\lambda_I \in \mathbb{F}_2^n} \sum_{\lambda_O \in \mathbb{F}_2^m} 2^{-n/2} \left\langle \lambda_O \mid C^F \mid \lambda_I \right\rangle \left|\lambda_I\right\rangle \left|\lambda_O\right\rangle, \tag{13}$$

at the cost of one superposition query to O_F. The final state is called the correlation state.

Proof. We calculate it step by step.

1. At the beginning, we have state $\left|+\right\rangle^{\otimes n} \left|0\right\rangle^{\otimes m}$. This state can be computed from $\left|0\right\rangle^{\otimes n} \left|0\right\rangle^{\otimes m}$ by applying an n-qubit Hadamard gate on the first n registers.
2. Apply oracle O_F. The state is now

$$\sum_{x \in \mathbb{F}_2^n} 2^{-n/2} \left|x\right\rangle \left|F(x)\right\rangle. \tag{14}$$

3. Apply $P_I^* \otimes P_O$:

$$
\begin{aligned}
&(P_I^* \otimes P_O) \sum_{x \in \mathbb{F}_2^n} 2^{-n/2} |x\rangle |F(x)\rangle \\
&= \sum_{\lambda_I \in \mathbb{F}_2^n} \sum_{\lambda_O \in \mathbb{F}_2^m} \left(\sum_{x \in \mathbb{F}_2^n} 2^{-n/2} \langle \lambda_I \mid P_I^* \mid x \rangle \langle \lambda_O \mid P_O \mid F(x) \rangle \right) |\lambda_I\rangle |\lambda_O\rangle \\
&= \sum_{\lambda_I \in \mathbb{F}_2^n} \sum_{\lambda_O \in \mathbb{F}_2^m} 2^{-n/2} \left(\sum_{x \in \mathbb{F}_2^n} \langle \lambda_O \mid P_O \mid F(x) \rangle \langle x \mid P_I^\dagger \mid \lambda_I \rangle \right) |\lambda_I\rangle |\lambda_O\rangle \\
&= \sum_{\lambda_I \in \mathbb{F}_2^n} \sum_{\lambda_O \in \mathbb{F}_2^m} 2^{-n/2} \langle \lambda_O \mid C^F \mid \lambda_I \rangle |\lambda_I\rangle |\lambda_O\rangle .
\end{aligned}
\tag{15}
$$

The last equation is due to Equation (11).

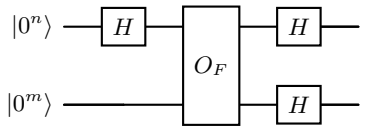


Fig. 1. the CEA in [37]

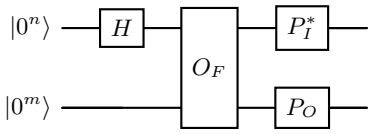


Fig. 2. the geometric CEA

The circuit for the CEA as described in Theorem 7 (and depicted in Figure 2) consists of one n-qubit Hadamard gate, one call to oracle O_F, and change-of-basis gates P_I^* and P_O. In the applications considered in this paper, P_I and P_O can be implemented efficiently, so the cost is dominated by O_F, i.e., the cost of one call to CEA approximates the cost of one call to O_F. To uncompute the CEA, one must implement $P_I^\top$ and $P_O^\dagger$. By [64], transposes, conjugates and their controlled version of a unitary are assumed feasible if the unitary itself is efficiently implemented (including the overall global phase). Therefore, assuming P_I and P_O can be efficiently implemented, so are $P_I^\top$ and $P_O^\dagger$. Additionally, note that we do not need the decryption oracle of F (when F is invertible) to uncompute the CEA.

Example 2. With Fourier basis, the CEA produces

$$
|\Psi_F\rangle = \sum_{\lambda_I \in \mathbb{F}_2^n} \sum_{\lambda_O \in \mathbb{F}_2^m} 2^{-m/2} \, \mathrm{LAT}_F(\lambda_I, \lambda_O) \, |\lambda_I\rangle \, |\lambda_O\rangle \, , \tag{16}
$$

which recovers the CEA of [37]. The CEA circuit is depicted in Figure 1.

3.1 Distinguish oracles by CEA

Quantum cryptanalysis is a special case of quantum hypothesis testing. Analysts formulate hypotheses regarding the cryptographic properties of an encryption

oracle. Since encryption function E is characterized by $|\Psi_E\rangle$, hypotheses about E translate directly into hypotheses about $|\Psi_E\rangle$. This perspective aligns with the classical geometric approach to cryptanalysis [6]. For example, assume $E : \mathbb{F}_2^n \to \mathbb{F}_2^n$ has a highly biased linear approximation $\lambda_I \cdot x \oplus \lambda_O \cdot E(x) = 0$ where $\lambda_I, \lambda_O \in \mathbb{F}_2^n$. This implies that $\langle \lambda_O \mid C^E \mid \lambda_I \rangle$ $(= \mathrm{LAT}_E(\lambda_I, \lambda_O))$ is close to ± 1. Hence, $\langle \Psi_E \mid \lambda_I, \lambda_O \rangle$ is close to $\pm 2^{-n/2}$ due to Example 2. In contrast, a random oracle expects $|\langle \Psi_{RO} \mid \lambda_I, \lambda_O \rangle| \approx 2^{-n}$. Thus, a significant deviation in amplitude exists between $|\Psi_E\rangle$ and $|\Psi_{RO}\rangle$. If we can distinguish $|\Psi_E\rangle$ from $|\Psi_{RO}\rangle$, then we can distinguish E from RO.

Apparently, this task is analogous to quantum state discrimination (QSD, see [3]). In the symmetric binary case, QSD identifies an unknown state drawn from two known states by uniformly random. However, our cryptographic distinguishing problem is different from QSD in adversaries' knowledge, ability and power: (1) QSD assumes full knowledge of $|\Psi_E\rangle$ and $|\Psi_{RO}\rangle$, while in cryptanalysis attackers typically possess only partial information about $|\Psi_E\rangle$. (2) QSD focuses on optimal measurement strategies, while cryptanalysts have extra ability to query the oracle O_F, enabling powerful quantum algorithms. (3) The cryptographic distinguishing problem can often be solved more efficiently than by direct application of QSD as a subroutine, as discussed in Appendix A.

Based on the unified preparation-measurement template, we categorize quantum distinguishing attacks into two main classes (not exhaustive):

1. *Hidden algebraic structure detection.* The encryption function has certain algebraic structures, such as a hidden period. Known algorithms detecting such structures can be viewed as distinguishers based on zero-correlation properties, implemented via sequential sampling. This category includes Simon's algorithm (and its variant for the hidden linear structure problem discussed in Subsection 4.3), Deutsch-Josza algorithm, Bernstein-Vazirani algorithm, and their generalizations or applications. This type of attacks somehow resembles QSD since they do not bother cascading multiple O_F's, but they are in fact more elegant as they leverage algebraic structures. Because these attacks rely on Fourier bases, we discuss them in Subsection 4.2.
2. *Event probability deviation hypothesis testing.* The probability of certain outcomes $(\lambda_I, \lambda_O) \in \mathbb{F}_2^n \times \mathbb{F}_2^m$ when measuring the correlation state differs noticeably between two hypotheses. We apply amplitude estimation to capture the probability deviation. The testing typically offers an approximately quadratic speed-up. Examples of known quantization include quantum differential, linear, multidimensional linear, zero-correlation attacks. Details are discussed in Subsection 3.3.

We note that the above classification is not exhaustive. For example, [43] quantizes truncated differential attacks by applying Ambainis' algorithm to plaintext structures besides an outer-loop Grover's search. Analyses of such quantum-walk-based hybrid approaches is left for future work.

A crucial question is whether a classical adversary can efficiently sample from the correlation matrix C^F. In some cases, such as the quasidifferential basis (Subsection 5.1), efficient classical sampling is possible. In other cases, like

the Fourier basis (Section 4), a classical adversary cannot obtain the exact C^F without querying the entire codebook. This classical limitation enables quantum algorithms, such as Simon's algorithm, to achieve a speed-up beyond quadratic.

3.2 Chosen-plaintext attacks

The CEA is by default a known-plaintext attack, i.e. any possible plaintext-ciphertext pair $(x, F(x))$ contributes to $|\Psi_F\rangle$ equally. However, some cryptanalysis techniques are chosen-plaintext, such as differential attacks. In this subsection, we characterize chosen-plaintext properties in the geometric approach and quantum algorithms.

Definition 7 (Chosen-plaintext pushforward operator). *Let $F : \mathbb{F}_2^n \to \mathbb{F}_2^m$ be a function and $S_I \subseteq \mathbb{F}_2^n$. The pushforward operator along $F|_{S_I}$ is a linear map from $\mathbb{C}^{2^n}$ to $\mathbb{C}^{2^m}$:*

$$T^{F|_{S_I}} := \sum_{x \in S_I} |F(x)\rangle \langle x| . \tag{17}$$

Definitions of correlation matrices and approximations follow Definitions 5 and 6. The element at row λ_O and column λ_I in $C^{F|_{S_I}}$ is

$$\langle \lambda_O \mid C^{F|_{S_I}} \mid \lambda_I \rangle = \sum_{x \in S_I} \langle \lambda_O \mid P_O \mid F(x) \rangle \langle x \mid P_I^\dagger \mid \lambda_I \rangle . \tag{18}$$

Additionally, the Frobenius norm of $C^{F|_{S_I}}$ is $\sqrt{|S_I|}$. The CEA is also adapted to chosen plaintexts, replacing the initial state $|+\rangle^{\otimes n} |0\rangle^{\otimes m}$ by $|S_I\rangle |0\rangle^{\otimes m}$.

Theorem 8 (Chosen-plaintext CEA). *Given an oracle O_F of a function $F : \mathbb{F}_2^n \to \mathbb{F}_2^m$, computing $(P_I^* \otimes P_O) O_F |S_I\rangle |0\rangle^{\otimes m}$ yields*

$$|\Psi_{F|_{S_I}}\rangle := \sum_{\lambda_I \in \mathbb{F}_2^n, \lambda_O \in \mathbb{F}_2^m} |S_I|^{-1/2} \langle \lambda_O \mid C^{F|_{S_I}} \mid \lambda_I \rangle |\lambda_I\rangle |\lambda_O\rangle , \tag{19}$$

using one superposition query to O_F.

Subset state preparation. Preparing $|S\rangle$ (where $S \subseteq \mathbb{F}_2^n$) is a special case of the *quantum resampling problem*: given an initial state (e.g., the uniform superposition) that is easy to prepare, we wish to transform it into the target state (here $|S\rangle$). Depending on the algebraic structure of S, different preparation strategies can be employed. On the one hand, the cases with clear structure admit preparation in polynomial time:

1. Singleton set. If $S = \{a\}$, then $|S\rangle$ can be prepared by applying $\bigotimes_{i=1}^n X^{a_i}$.
2. Affine linear subspace. If $S = a + L$ where L is a linear subspace, then we apply a three-step strategy: first prepare $|a\rangle$ via pauli X gates, then apply Hadamard gates on the $\dim(L)$ free qubits corresponding to a basis of L, and finally implement a linear transform via CNOT gates.

On the other hand, if S lacks any discernible property, one resorts to QAA. Given a quantum oracle O_S that recognizes S, the preparation requires $O(\sqrt{2^n/|S|})$ queries to O_S. Note that O_S itself may be expensive to implement.

When S is a subset of an affine subspace T, quantum rejection sampling [52] offers an efficient method, which specializes QAA to the resampling setting with known amplitude ratios. Assume that we can prepare $|T\rangle$ via oracle O_T, i.e. $O_T|0\rangle = |T\rangle$, then $|S\rangle$ can be prepared exactly in $\Theta(\sqrt{|T|/|S|})$ queries to O_T and $O_T^\dagger$; a detailed derivation is given in Appendix B.

3.3 Event Probability Deviation Hypothesis Testing

In this subsection, we describe a hypothesis testing with two hypotheses under which the probability of a certain event differs noticeably and demonstrate that quantum algorithms can achieve a quadratic speed-up in query complexity for two typical scenarios. We begin by defining the oracles that provide access to the underlying random variable under different hypotheses.

Definition 8 (Random variable oracle). *Let D be a probability distribution over a finite set S. We define two oracles:*

1. *Classical oracle: each query returns a single classical sample $x \sim D$.*
2. *Quantum oracle: a unitary operator O_D acting on a Hilbert space H_S with basis $\{|x\rangle\}_{x\in S}$ and an ancilla register, such that for some fixed initial state $|0\rangle_S \in H_S$,*

$$O_D|0\rangle_S|0\rangle = \sum_{x\in S} \alpha_x |x\rangle|\psi_x\rangle, \tag{20}$$

where $\alpha_x \in \mathbb{C}$, $|\alpha_x|^2 = D(x)$ for all $x \in S$ and $|\psi_x\rangle$ are arbitrary unit vectors. Querying O_D or its inverse counts as one query.

Definition 9 (Event Probability Deviation Hypothesis Testing, EPDHT). *Let S be a finite set, and let $G \subseteq S$ be a fixed subset (the "event"). Consider two known probability distributions D_1 and D_2 over S, with associated probabilities $p_i := \Pr_{x\sim D_i}[x \in G]$ for $i = 1, 2$, where $p_1 < p_2$ (without loss of generality).*

Let O be an oracle that is promised to be either for D_1 or for D_2 by uniform random. The EPDHT problem is to determine whether O implements D_1 or D_2.

Cryptanalysis often observes small deviations, i.e. $|p_1 - p_2| \ll 1$, and EPDHT in such scenario requires a large number of queries. To analyze its asymptotic complexity, a theorem is presented to demonstrate the quantum quadratic speed-up. Note that the presence of quantum speed-ups is not limited to the two assume cases in Theorem 9. Nevertheless, the two cases are essential and worth highlights. See Appendix C for the proof.

Theorem 9. *For EPDHT with knowledge no other than p_1, p_2 and G, there exists a quantum algorithm that succeeds with high probability and with query complexity*

$$N_Q = \Theta\left(\frac{1}{|\sqrt{p_1} - \sqrt{p_2}|}\right). \tag{21}$$

Meanwhile, any classical algorithm that succeeds with high probability requires a complexity

$$N_C = \Omega\left(\frac{1}{D_{\mathrm{KL}}(\mathrm{Bern}(p_1)||\mathrm{Bern}(p_2))}\right). \tag{22}$$

Specifically, we name two regimes:

1. *If $p_1 \ll p_2 \ll 1$, then $N_Q = \Theta(1/\sqrt{p_2}), N_C = \Omega(1/p_1)$;*
2. *If $\varepsilon := |p_1 - p_2| \ll \min\{p_1, 1-p_1\} = \Theta(1)$, then $N_Q = \Theta(1/\varepsilon), N_C = \Omega(1/\varepsilon^2)$.*

Here we provide some remarks on the two regimes of Theorem 9.

- Regime 1 is a rare-event distinguisher. The quantum complexity of $2\pi/\sqrt{p_2}$ resembles the quadratic speed-up of Grover's search (or QAA) for hypothesis 2, where the query complexity $(\pi/(2\sqrt{p_2}))$ is essentially determined by the larger probability p_2 while p_1 does not have asymptotic impact. This behavior is exactly what underlies the quantum speed-up in differential cryptanalysis: the speed-up can be obtained either by Grover-search for a differential right pair (equivalent to QAA on a difference distribution table entry) [43] or by QAE on a quasidifferential correlation state, as detailed in Subsection 5.1.
- Regime 2 recovers quantum counting. We will demonstrate this through a concrete application that reproduces the linear distinguisher of [43]. Later in Subsection 4.1 another example reproduces the linear distinguisher of [37], alongside a non-asymptotic comparison.

Example 3 (linear distinguisher [43] by quantum counting). Computing $O_F \cdot |0^n\rangle |0^m\rangle$ gives $\sum_{x\in\mathbb{F}_2^n} |x\rangle |F(x)\rangle$. Given a highly-biased linear approximation $\lambda_I \to \lambda_O$, consider $G := \{(x,y) \in \mathbb{F}_2^n \times \mathbb{F}_2^m : \lambda_I \cdot x \oplus \lambda_O \cdot y = 0\}$. For EPDHT, $p_1 = 1/2, p_2 = (1+c)/2$ (assuming $c > 0$ and $c \ll 1$ without loss of generality). The quantum query complexity is then $4\sqrt{2}\pi/c = \Theta(1/c)$ according to Regime 2. Actually, the complexity is $1/2$ times as Regime 2 because the algorithm employs O_F in the marking oracle instead; however, computing the parity doubles the complexity and cancels with this advantage.

It follows directly from the CEA (Theorem 7) and the analysis in Theorem 9 that one can derive a CEA-based version of the EPDHT.

Theorem 10 (CEA-based EPDHT). *Let $F : \mathbb{F}_2^n \to \mathbb{F}_2^m$ be a function accessible with quantum oracle O_F and $G \subseteq \mathbb{F}_2^n \times \mathbb{F}_2^m$. Assume we have two hypotheses 1 and 2. Denote $\sum_{(\lambda_I,\lambda_O)\in G} |\langle \lambda_O \mid C^F \mid \lambda_I \rangle|^2$ by "capacities" Cap_1, Cap_2 under hypotheses 1 and 2 respectively. Then, the EPDHT between the two hypotheses succeeds with high probability using N queries to O_F, where $N = \Theta(2^{n/2}/|\sqrt{Cap_1} - \sqrt{Cap_2}|)$.*

Example 4 (Distinguishers $\mathcal{A}_1, \mathcal{B}_1$ of [37]). Applying Theorem 10 to a linear attack that uses a single linear approximation $\lambda_I \to \lambda_O$ for a function $F : \mathbb{F}_2^n \to \mathbb{F}_2^n$ yields a query complexity of $\Theta(2^{n/2}/c)$, where c denotes the absolute correlation of the approximation.

Example 4 reveals a central limitation of the CEA-based EPDHT approach: the problem of *dilution*. Only a specific subset of states contributes to discriminating between the hypotheses, while the remaining states are not constrained by the hypotheses. Consequently, distinguishers such as $\mathcal{A}_1$ and $\mathcal{B}_1$, while representing an important step toward advanced quantum multidimensional linear distinguishers, perform strictly worse than the generic lower bound of $2^{n/2}$. It is therefore crucial to suppress irrelevant states when employing EPDHT; the appropriate technique depends on the geometric basis in use, which explains the improvement of $\mathcal{A}_2, \mathcal{A}_3, \mathcal{B}_2, \mathcal{B}_3$ of [37] (see Appendix D for details). For instance, phase oracles are effective for the Fourier basis (see Subsection 4.1), while chosen-plaintext techniques are suitable for the quasidifferential basis (see Subsection 5.1).

When p_1 and p_2 are not fixed but drawn from distributions, the previous analysis can be generalized using quantiles. The following theorem provides a distinguisher when the quantiles of the two distributions are sufficiently separated.

Theorem 11. *Let $\mathcal{D}_1, \mathcal{D}_2$ be two distributions over functions $F : \mathbb{F}_2^n \to \mathbb{F}_2^m$. For a fixed set $G \subseteq \mathbb{F}_2^n \times \mathbb{F}_2^m$, define the random variable*

$$p_i(F) := \sum_{g \in G} \left| \langle \Psi_F \mid g \rangle \right|^2, \quad F \sim \mathcal{D}_i, \tag{23}$$

where $|\Psi_F\rangle$ is the correlation state of F. Let p_L be the 75%-quantile of p_1 and p_R the 25%-quantile of p_2. Define $\varepsilon := p_R - p_L > 0$ and $\sigma := \sqrt{p_L} + \sqrt{p_R}$. Then, given quantum oracle access to F, one can distinguish between $\mathcal{D}_1$ and $\mathcal{D}_2$ with advantage at least 0.2 using at most $2N$ superposition queries to O_F, where

$$N = \frac{\pi}{\varepsilon} \left(\sigma + \sqrt{\sigma^2 + 2\varepsilon} \right) \tag{24}$$

Proof. We construct a QAE-based distinguisher. Let $t \in \mathbb{R}$ be a threshold to be determined. The algorithm proceeds as follows:

1. Run QAE with N queries (requiring $2N$ calls to O_F) to obtain an estimate $\tilde{p}$ of $p(F)$.
2. If $\tilde{p} < t$, output "$\mathcal{D}_1$"; otherwise output "$\mathcal{D}_2$".

Choose t and N such that

$$p_L + \frac{2\pi\sqrt{p_L}}{N} + \frac{\pi^2}{N^2} \leq t, \tag{25}$$

$$p_R - \frac{2\pi\sqrt{p_R}}{N} - \frac{\pi^2}{N^2} \geq t. \tag{26}$$

Under hypothesis $\mathcal{D}_1$, with probability at least 0.75 we have $p_1 \leq p_L$. Conditioned on this event, Equation (25) together with QAE (Theorem 5) guarantees $\tilde{p} \leq t$ with probability at least $8/\pi^2$. Hence, the probability of correctly outputting "$\mathcal{D}_1$" is at least $0.75 \times 8/\pi^2 > 0.6$. Similarly, under $\mathcal{D}_2$, with probability

at least 0.75 we have $p_2 \geq p_R$, and Equation (26) ensures $\tilde{p} \geq t$ with probability at least $8/\pi^2$, giving a correct-output probability also exceeding 0.6. The overall advantage is therefore at least 0.2. Solving Equations (25) and (26) for N yields the result.

The choice of the 75% and 25% quantiles in Theorem 11 is flexible; other quantile pairs can be used, trading off the confidence level against the required gap ε. In addition, the QAE success probability ($8/\pi^2$) can be amplified by repeating for constant times.

The conditions for Theorem 11 are fairly loose: analysts often assume that the underlying distribution is relatively dense. For example, the seminal work by Selçuk [54] introduces a model for linear key recovery, in which the exact correlation of the linear distinguisher under the two hypotheses (using either the right or a wrong subkey in partial decryption) is assumed to be concentrated around two distinct points with a noticeable deviation, referred to as *the right-key and wrong-key assumptions*, respectively. In practice, these assumptions can be inaccurate; consequently, some previous works turn to fixed-key models, which are precisely exemplified by the geometric approach (e.g., [8]) and are well-suited for quantization via Theorem 10.

4 Fourier Basis

Choosing the Fourier basis [5], as demonstrated in Examples 1 and 2, corresponds to change-of-basis matrices $P_I = H^{\otimes n}$ and $P_O = H^{\otimes m}$, which diagonalizes the addition group $(\mathbb{F}_2^n, \oplus)$ and reveals the linear approximation table (LAT) of the function.

4.1 Phase Oracles

This subsection investigates the general phase oracles for vector Boolean functions (Definition 2). They efficiently suppresses any designated non-informative C^F rows in $|\Psi_F\rangle$, providing a strong countermeasure to the dilution. For Boolean functions ($m = 1$), a phase oracle extracts its Fourier spectrum, which underlies Bernstein-Vazirani and Deutsch-Josza algorithm and save half of their superposition queries. To elaborate, [22] observes that for an arbitrary Boolean function $F : \mathbb{F}_2^n \to \mathbb{F}_2$, executing the circuit of Bernstein-Vazirani algorithm without final measurement yields the Fourier spectrum of F; this spectrum coincides (up to a factor of $\sqrt{2}$) with the amplitudes of states $|\lambda_I\rangle |\lambda_O\rangle$ where $\lambda_O = 1 \in \mathbb{F}_2$ in $|\Psi_F\rangle$. In effect, the algorithm suppresses the non-informative row $\lambda_O = 0$. The theorems below clarify that this capability stems from the use of phase oracles (Definition 2) for vector Boolean functions.

Theorem 12. *Given an oracle O_F for a function $F : \mathbb{F}_2^n \to \mathbb{F}_2^m$ and an arbitrary state $|\psi\rangle := \sum_{\lambda_O \in \mathbb{F}_2^m} \psi_{\lambda_O} |\lambda_O\rangle \in \mathbb{C}^{2^m}$, computing*

$$(H^{\otimes n} \otimes I^{\otimes m}) \cdot O_F^{\pm} \cdot |+\rangle^{\otimes n} |\psi\rangle \tag{27}$$

yields

$$\sum_{\lambda_I \in \mathbb{F}_2^n} \sum_{\lambda_O \in \mathbb{F}_2^m} \psi_{\lambda_O} \cdot \mathrm{LAT}_F(\lambda_I, \lambda_O) \cdot |\lambda_I\rangle |\lambda_O\rangle . \tag{28}$$

Proof. Compute step in step.

1. The initial state is $|+\rangle^{\otimes n} |\psi\rangle$.
2. Apply $O_F^{\pm}$:

$$\sum_{x \in \mathbb{F}_2^n} \sum_{\lambda_O \in \mathbb{F}_2^m} 2^{-n/2} (-1)^{\lambda_O \cdot F(x)} \langle \lambda_O \mid \psi \rangle |x\rangle |\lambda_O\rangle . \tag{29}$$

3. Apply an n-qubit Hadamard transform on the first n registers:

$$\sum_{\lambda_I \in \mathbb{F}_2^n} \sum_{x \in \mathbb{F}_2^n} \sum_{\lambda_O \in \mathbb{F}_2^m} 2^{-n} (-1)^{\lambda_I \cdot x \oplus \lambda_O \cdot F(x)} \langle \lambda_O \mid \psi \rangle |\lambda_I\rangle |\lambda_O\rangle , \tag{30}$$

   which equals the desired result.

Specifically, by choosing $|\psi\rangle$ as the subset state $|S_O\rangle$ for some $S_O \subseteq \mathbb{F}_2^m$, one obtains a reduced version of $|\Psi_F\rangle$. This is formalized in the following theorem.

Theorem 13 (Reduced correlation state). *Given an oracle O_F of a function $F : \mathbb{F}_2^n \to \mathbb{F}_2^m$ and a subset $S_O \subseteq \mathbb{F}_2^m$, computing $H^{\otimes(n+m)} \cdot O_F \cdot H^{\otimes(n+m)} \cdot |0^n\rangle |S_O\rangle$ yields the reduced correlation state*

$$|\Psi'_F\rangle := \sum_{\lambda_I \in \mathbb{F}_2^n} \sum_{\lambda_O \in S_O} |S_O|^{-1/2} \mathrm{LAT}_F(\lambda_I, \lambda_O) |\lambda_I\rangle |\lambda_O\rangle . \tag{31}$$

Please see Subsection 3.2 for subset state preparation techniques. Besides, three special cases of S_O are of particular interest:

1. $S_O = \mathbb{F}_2^m$: the plain CEA, which retains all rows of C^F.
2. $S_O = \{\lambda_O\}, \lambda_O \neq 0$: the phase oracle for function $\lambda_O \cdot F$. This saves a factor 2 for complexity brought by computing function $\lambda_O \cdot F$. In addition, replacing XOR oracles by phase oracles improves query complexity by 2 times if the algorithm directly measures $|\Psi'_F\rangle$, and by $\sqrt{2}$ times if the algorithm applies QAA on $|\Psi'_F\rangle$.
3. $S_O = \{\lambda || 0^{m-m'} : \lambda \in \mathbb{F}_2^{m'}\}$: the CEA in [37] (distinguishers $\mathcal{A}_1, \mathcal{A}_2, \mathcal{B}_1, \mathcal{B}_2$). Intuitively, this restricts attention to linear approximations whose output masks are inactive on the last $m - m'$ bits. Such approximations capture the full information of the truncated function $\mathrm{msb}_m[F]$, where msb_m denotes the truncation to the first m output bits. This recovers the technique of [39], which truncates a quantum oracle with a single superposition query, in contrast to the two queries required by naive CNOT copying.

Example 5. Consider linear cryptanalysis using a linear approximation $\lambda_I \to \lambda_O$ with correlation $c > 0$. Applying case 2 yields a state $|\Psi'_F\rangle$ whose amplitude on $|\lambda_I\rangle |\lambda_O\rangle$ is proportional to c. QAA on this state requires $\pi/(2c) = \Theta(1/c)$ queries.

In terms of non-asymptotic complexity, Example 5 achieves a factor of $8\sqrt{2}$ improvement over the linear distinguisher of [43] (Example 3) and a factor of $4\sqrt{2}$ improvement over the approach in [37] (when augmented with the truncation technique of [39]). Nevertheless, claims of such non-asymptotic improvements should be made with caution, as they depend on concrete circuit implementations. Current quantum devices operate in the noisy intermediate-scale quantum (NISQ) regime, where modifications to an algorithm can inadvertently introduce additional noise, potentially reducing the success probability and thus requiring extra queries to maintain it.

Surprisingly, the ability to reduce C^F by selecting specific rows is a unique feature of the Fourier basis among all geometric bases, unless we employ complicated constructions involving more operators and possibly ancilla qubits. In other words, phase oracles cannot be constructed using other geometric bases via the conjugate relation (cf. Equation (3)). This explains why quantum algorithms based on phase oracles are naturally discussed within the Fourier basis framework. To understand this, consider constructing a new phase oracle O'_F by composing the standard oracle O_F with general geometric change-of-basis operators other than Hadamard matrices. Given the linearity of both quantum mechanics and the geometric approach, Theorem 12 holds if and only if such a newly-defined O'_F maps each basis state $|x\rangle|\lambda_O\rangle$ to a product state $|\psi\rangle|\lambda_O\rangle$ for some $|\psi\rangle$. The following theorem formalizes this constraint (see Appendix E for its proof).

Theorem 14. *Let H_1, H_2 be unitaries of dimension 2^m. Define the operator*

$$O'_F := (I^{\otimes n} \otimes H_2) \cdot O_F \cdot (I^{\otimes n} \otimes H_1). \tag{32}$$

Suppose that for every function $F : \mathbb{F}_2^n \to \mathbb{F}_2^m$, every $x \in \mathbb{F}_2^n$, and every $\lambda_O \in \mathbb{F}_2^m$, operator O'_F maps each basis state $|x\rangle|\lambda_O\rangle$ to a product state $|\psi\rangle|\lambda_O\rangle$ for some $|\psi\rangle$. Then H_1 and H_2 must satisfy

$$H_1 = H^{\otimes m} D_1, \qquad H_2 = D_2 H^{\otimes m}, \tag{33}$$

where D_1, D_2 are diagonal unitaries.

4.2 Hidden Algebraic Structure Detection

In this subsection, we reinterpret three foundational quantum algorithms (Simon's, Deutsch-Jozsa, and Bernstein-Vazirani) with the geometric approach, corresponding to the *hidden algebraic structure detection* in Subsection 3.1. We show how each algorithm efficiently detects certain hidden algebraic structure and corresponds to a specific zero-correlation hypothesis-testing scenario on $|\Psi_F\rangle$ for a function F. Concretely, these algorithms follow the same scheme: first prepare $|\Psi_F\rangle$ via the CEA, then measure it, and finally carry out specific classical postprocess. The key idea is that the measurement outcomes always come from a specific hidden proper subset of the computational basis. Therefore, these well-known algorithms, which are based on the quantum Fourier transform (QFT), emerge not as isolated primitives, but as natural instances of testing zero-correlation hypotheses within the Fourier basis of the geometric approach.

Simon's algorithm. Simon's algorithm (Theorem 3, [57]) speeds up solving HPP (Definition 3) by essentially suppressing states $|\lambda_I\rangle\,|y\rangle$ in its output state that does not satisfy $\lambda_I \cdot s = 0$. In the geometric approach, this speed-up stems from a zero-correlation property encoded in C^F (or say LAT_F). Simon's premise imposes a strong constraint on the LAT:

$$\mathrm{LAT}_F(\lambda_I, \lambda_O) = (-1)^{\lambda_I \cdot s}\, \mathrm{LAT}_F(\lambda_I, \lambda_O) \tag{34}$$

which follows directly from the definition of LATs. Consequently, for every λ_I we have either $\lambda_I \cdot s = 0$ or $\mathrm{LAT}_F(\lambda_I, \lambda_O) = 0$. If we perform the CEA with Fourier basis, then measuring $|\Psi_F\rangle$ yields a pair (λ_I, λ_O) with $\lambda_I \cdot s = 0$ with certainty. After collecting $n-1$ such linearly independent samples, the secret period s can be recovered by solving a system of linear equations over $\mathbb{F}_2$ on classical computers. This process is precisely the standard Simon procedure, now reinterpreted as *sampling from the Fourier correlation matrix and detecting its support restricted to the subspace* $\{\lambda_I : \lambda_I \cdot s = 0\}$.

The practical applicability of Simon's algorithm crucially depends on the structural fidelity of $|\Psi_F\rangle$. As highlighted in prior works, two ideal conditions must approximately hold: (1) All basis states satisfying $\lambda_I \cdot s = 0$ should possess equal amplitude in $|\Psi_F\rangle$. The seminal work [42] suggest ensuring this condition by bounding the maximum probability of unwanted collisions, i.e., a pair (x, y) satisfying $F(x) = F(y)$ and $x \oplus y \notin \{0, s\}$. For random functions or permutations, such a bound is generally attainable. (2) All basis states satisfying $\lambda_I \cdot s = 1$ should have exactly zero amplitude. On realistic NISQ devices, however, this condition cannot be met. Device noise introduces decoherence and gate errors that populate the otherwise forbidden subspace, effectively reducing Simon's problem to a *learning parities with noise* (LPN) problem, as shown in [50,29]. The problem can become computationally hard even for quantum algorithms, especially when the circuit noise exceeds a certain threshold.

Bernstein-Vazirani algorithm. The Bernstein-Vazirani algorithm identifies the hidden mask $s \in \mathbb{F}_2^n$ of an n-bit linear Boolean function $F(x) := s \cdot x$ using one superposition query, while the classical counterpart requires n queries, one for each basis vector of $\mathbb{F}_2^n$. Compared with Simon's algorithm, the algorithm has seen limited use in cryptanalysis to the best of our knowledge. As noted in [22], the state before the final measurement corresponds to the Fourier spectrum of the Boolean function. [63] proposes attacks using Bernstein-Vazirani algorithm by combining it with Simon's algorithm. In fact, it deals with a generalized variant of Simon's problem which is to find hidden linear structures, and it benefits from Bernstein-Vazirani by borrowing its idea of using Boolean function phase oracles. [63] proposes attacks that combine Bernstein-Vazirani with Simon's algorithm, addressing a generalized Simon problem that identifies hidden linear structures; their approach benefits from the use of Boolean function phase oracles inspired by Bernstein-Vazirani. Meanwhile, [14] applies the algorithm to parallelizable MACs constructed as XOR sums of keyed permutations (one per message block). Their method encodes two choices per message block into a Boolean value, thereby

converting the XOR sum into a linear Boolean function. In summary, these works decompose a target function into Boolean components amenable to the Bernstein–Vazirani routine.

Its procedure is as follows:

1. Initialize $|0\rangle^{\otimes n}$.
2. Apply n-qubit Hadamard gate on all qubits and obtain $|+\rangle^{\otimes n}$.
3. Apply phase oracle $O_F^{\pm}$ and obtain $\sum_{x\in\mathbb{F}_2^n}(-1)^{s\cdot x}|x\rangle$.
4. Apply n-qubit Hadamard gate on all qubits again and obtain

$$\sum_{\lambda\in\mathbb{F}_2^n}\sum_{x\in\mathbb{F}_2^n} 2^{-n/2}(-1)^{(\lambda\oplus s)\cdot x}|\lambda\rangle = |s\rangle\,. \tag{35}$$

We now reinterpret the Bernstein-Vazirani algorithm within the geometric framework of Section 3 while extending it naturally to vector Boolean functions. For a linear function $F:\mathbb{F}_2^n\to\mathbb{F}_2^m$ defined by $F(x)=Mx$ with an $m\times n$ Boolean matrix M, its LAT satisfies $\mathrm{LAT}_F(\lambda_I,\lambda_O)=\langle\lambda_I \mid M^\top\lambda_O\rangle$. Following a strategy analogous to Simon's algorithm, one collects samples (λ_I,λ_O) and solves the linear system $\lambda_I = M^\top\lambda_O$ for the unknown matrix $M^\top$. Once sufficiently many linearly independent output masks λ_O have been obtained, i.e., a basis of $\mathbb{F}_2^m$, the matrix $M^\top$ can be recovered. This process requires $O(m)$ quantum queries in expectation, though the exact non-asymptotic cost may exceed m due to the possibility of obtaining linearly dependent samples.

Compared to the classical algorithm that queries each canonical basis vector of $\mathbb{F}_2^n$ individually, a polynomial quantum speed-up occurs only when $n > m$. For large n the quantum circuit becomes deep and may pose practical challenges; an example is the MAC-forgery attack of [14, Subsection 3.3], where n is the number of message blocks and m the block size. With a typical block size $m = 128$, realizing such a speed-up would require on the order of 10^4 qubits, placing it beyond the reach of near-term quantum devices.

The original Bernstein-Vazirani algorithm for Boolean functions ($m = 1$) illustrates the geometric idea clearly. In that case the correlation state becomes $|\Psi_F\rangle = (|0^n\rangle|0\rangle + |s\rangle|0\rangle)/\sqrt{2}$. Measuring the first n qubits yields either 0 or the secret string s, each with probability $1/2$, giving an expected query count of 2. However, by employing a phase oracle, which suppresses the non-informative row $\lambda_O = 0$ as detailed in Subsection 4.1, the algorithm deterministically outputs s after a single superposition query.

Deutsch-Josza algorithm. The Deutsch-Jozsa algorithm distinguishes a balanced Boolean function from a constant one with a single superposition query. Using a phase oracle, it prepares the Fourier spectrum of the function. For a balanced function F, $\mathrm{LAT}_F(0,1) = 0$; for a constant function $F(x) = a \in \mathbb{F}_2$, $\mathrm{LAT}_F(0,1) = (-1)^a$. Measuring the reduced correlation state $|\Psi'_F\rangle$ therefore decides the hypothesis with certainty: outcome 0 indicates a constant function, any other outcome a balanced one.

In the view of the geometric approach, the Deutsch-Jozsa problem corresponds to a zero-correlation hypothesis on the partition of the column index set

$\mathbb{F}_2^n$ into the two disjoint subsets $\{0\}$ and $\mathbb{F}_2^n \setminus \{0\}$. If F is balanced, the reduced correlation state has support only in $\mathbb{F}_2^n \setminus \{0\}$; if constant, only in $\{0\}$. Hence a single measurement of $|\Psi_F\rangle$ suffices to distinguish the two cases.

This idea extends to vector Boolean functions: one may partition $\mathbb{F}_2^n \times (\mathbb{F}_2^m \setminus \{0\})$ into disjoint subsets offline, with the promise that $|\Psi_F'\rangle$ is supported on exactly one subset. After running the CEA and measuring, membership of the outcome in a particular subset decides the hypothesis with one superposition query. In the Boolean case, partitioning $\mathbb{F}_2^n$ into 2^n singletons and requiring the support to be a single element recovers the Bernstein-Vazirani algorithm; indeed for $n = 1$ the two algorithms coincide, as already noted in [14].

Both Deutsch-Jozsa and Simon's algorithms fit the same hypothesis-testing pattern: each hypothesis restricts the support of the correlation state to a prescribed subset of $\mathbb{F}_2^n \times \mathbb{F}_2^m$. The key difference is that Deutsch-Jozsa uses disjoint subsets, allowing a decision after one query, whereas Simon's algorithm employs a structured family of subsets that requires $O(n)$ queries. A general characterization of such hypothesis-testing problems, potentially employing tools like multi-hypothesis sequential probability ratio tests (MSPRT), is left for future work.

4.3 Hidden Linear Structure Problem

In this subsection, we revisit the hidden linear structure problem, a variant of Simon's problem (cf. [63,35,48]).

Definition 10 (Hidden linear structure problem, HLSP). *Given an oracle of a function $F : \mathbb{F}_2^n \to \mathbb{F}_2^m$ that satisfies condition*

$$x \oplus y = \Delta_I \;\Rightarrow\; F(x) \oplus F(y) = \Delta_O \tag{36}$$

for some secret vectors $\Delta_I \in \mathbb{F}_2^n, \Delta_O \in \mathbb{F}_2^m$, find Δ_I and Δ_O.

In the geometric framework with the Fourier basis, a probability-one linear structure $\Delta_I \to \Delta_O$ implies that for every $(\lambda_I, \lambda_O) \in \mathbb{F}_2^n \times \mathbb{F}_2^m$,

$$\lambda_I \cdot \Delta_I \oplus \lambda_O \cdot \Delta_O = 0 \quad \text{or} \quad \langle \lambda_O | C^F | \lambda_I \rangle = 0, \tag{37}$$

as shown via Fourier analysis in [28]. Hence, by collecting pairs (λ_I, λ_O) from measurements, one can construct linear equations in $\Delta_I \| \Delta_O$ and solve the system. Formally, we repeatedly prepare the correlation state $|\Psi_F\rangle$ via the CEA and measure in the computational basis, obtaining samples (λ_I, λ_O). Candidates for $\Delta_I \| \Delta_O$ are filtered by the condition $\lambda_I \cdot \Delta_I \oplus \lambda_O \cdot \Delta_O = 0$.

Theoretically, $O(n+m)$ samples suffice to determine (Δ_I, Δ_O), provided the sampling distribution is sufficiently uniform. However, realistic functions may introduce "unwanted collisions" analogous to those in Simon's problem. Following the idea of [42], we therefore bound the false-positive probability (i.e., misidentify a wrong (Δ_I, Δ_O)) in the following theorem (see Appendix F).

Theorem 15. *Let an outcome $\lambda_I \| \lambda_O$ be obtained from the sampling approach described above. Then*

$$\Pr[\lambda_I \cdot \Delta_I \oplus \lambda_O \cdot \Delta_O = 0] = \frac{1}{2}\left(1 + \Pr_x[F(x) \oplus F(x \oplus \Delta_I) = \Delta_O]\right). \tag{38}$$

By discussions from [42, Theorem 1], $\Theta(n+m)$ superposition queries suffice to solve the HLSP.

In this way, the geometric framework recovers a recent work [48], which systematically investigates Simon's algorithm and introduces a unified scheme for constructing periodic functions, accompanied by automated tools. A valuable insight of that work is to differentiate selected input branches in a generalized Feistel network (GFN), e.g., choosing distinct y_1, y_2 and compute $G(x, y_1) \oplus G(x, y_2)$ for a GFN output branch $G(\cdot, \cdot)$. By computing differences, unknown constant terms in the output are canceled. This modification grants it a significant advantage over traditional approaches that rely on a single query and the concatenation technique (cf. [35]). Certainly, the technique can be adapted to the abstract HLSP by applying a derivative $D_\Delta F$ with a nonzero input difference Δ. However, it should be noted that computing differentials doubles the circuit depth, as illustrated in Figure 3 where $X^\Delta := \bigotimes_{i=1}^n X^{\Delta_i}$. In contrast, the geometric CEA (Figure 2) requires only a single call to O_F. This provides a complementary approach, achieving the same asymptotic query complexity with half the circuit depth, which may be beneficial when depth is a constrained resource.

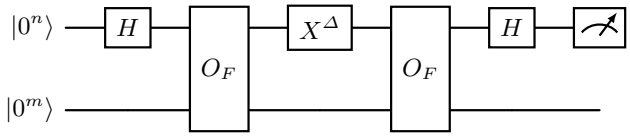


Fig. 3. Typical circuit for Simon's algorithm on a derivative

The HLSP is first tackled in [63] using a clever combination of Simon-like subroutine and a phase oracle to extract the Fourier spectrum of each output bit individually, requiring $\Theta(n^2)$ queries. While this technique resonates closely with the geometric approach (cf. Subsection 4.1), their method decomposes the vectorial Boolean function into its Boolean components and performs a case analysis on Δ_O. However, from the geometric perspective, this decomposition can be avoided by directly exploiting the structure of $|\Psi_F\rangle$, as it directly encodes the joint constraints on Δ_I and Δ_O, allowing both to be recovered without case distinction. This leads to a more unified procedure and reduces the query complexity to linear, offering a complementary alternative to their $\Theta(n^2)$ approach.

5 Other Bases

5.1 Quasidifferential Basis

In this subsection, we present the quantum algorithmic treatment of the quasidifferential basis within the geometric framework. We demonstrate its equivalence to the quantum differential distinguisher of Kaplan et al. [43] and explain why attacks based on this basis are probably limited to a quadratic quantum speed-up in some scenarios.

Samples in high-order spaces. Differential attacks operate on samples from Cartesian products of $\mathbb{F}_2^n$ as they involve two queries per sample. Following [40], we define the order of an attack.

Definition 11 (The order of an attack [40]). *An attack on a cipher $E : \mathbb{F}_2^n \to \mathbb{F}_2^m$ that uses samples from the space $(\mathbb{F}_2^n)^d$ is called a d-th-order attack.*

The d-th-order pushforward operator is defined as $T^{F^{\otimes d}} := (T^F)^{\otimes d}$, which corresponds to the pushforward of d parallel applications of F (see Theorem 6). For instance, the second-order pushforward operator is a linear map from $\mathbb{C}^{2^{2^n}}$ to $\mathbb{C}^{2^{2^m}}$:

$$T^{F^{\otimes 2}} := \sum_{x_1,x_2 \in \mathbb{F}_2^n} |F(x_1)\rangle\, |F(x_2)\rangle\, \langle x_2|\, \langle x_1|\,. \tag{39}$$

The change-of-basis matrices P_I and P_O then have dimensions 2^{nd} and 2^{md}, respectively. Calculations proceed analogously to the first-order case.

The quasidifferential basis. As established in [8], differential distinguishers are second-order attacks. Their geometric description employs the quasidifferential change-of-basis matrix:

$$P = \mathcal{Q} := \sum_{\Delta,\lambda \in \mathbb{F}_2^n} \sum_{x_1,x_2 \in \mathbb{F}_2^n} \langle \Delta \mid x_1 \oplus x_2 \rangle \cdot 2^{-n/2} (-1)^{\lambda \cdot x_1}\, |\Delta, \lambda\rangle\, \langle x_1, x_2|\,. \tag{40}$$

The corresponding correlation matrix satisfies

$$\langle \Delta_O, \lambda_O|\, C^{F^{\otimes 2}}\, |\Delta_I, \lambda_I\rangle = 2^{-(n+m)/2} \sum_{\substack{x \in \mathbb{F}_2^n: \\ F(x) \oplus F(x \oplus \Delta_I) = \Delta_O}} (-1)^{\lambda_I \cdot x \oplus \lambda_O \cdot F(x)}. \tag{41}$$

The quasidifferential correlation matrix $C^{F^{\otimes 2}}$ is not directly accessible to a quantum adversary. To construct an efficient CEA-based distinguisher, we must relate it to the classical differential probability. [8] shows that

$$\langle \Delta_O, 0|\, C^{F^{\otimes 2}}\, |\Delta_I, 0\rangle = 2^{(n-m)/2} \cdot \mathrm{DDT}_F(\Delta_I, \Delta_O). \tag{42}$$

However, directly using this entry with Theorem 10 and the event $\{(\Delta_I, 0, \Delta_O, 0)\}$ would lead to an inefficient distinguisher requiring $O(2^{(n+m)/2}/p)$ queries, exceeding the quantum generic bound.

Instead, the following property of the quasidifferential correlation matrix enables an efficient quantum attack.

Theorem 16.

$$\sum_{\lambda_I \in \mathbb{F}_2^n} \sum_{\lambda_O \in \mathbb{F}_2^m} \left| \langle \Delta_O, \lambda_O | \, C^{F \otimes 2} \, | \Delta_I, \lambda_I \rangle \right|^2 = 2^n \cdot \mathrm{DDT}_F(\Delta_I, \Delta_O). \tag{43}$$

Proof. The left-hand side expands to

$$\sum_{\lambda_I \in \mathbb{F}_2^n} \sum_{\lambda_O \in \mathbb{F}_2^m} \sum_{x, x' \in \mathbb{F}_2^n} 2^{-(n+m)} \times (-1)^{\lambda_I \cdot (x \oplus x') \oplus \lambda_O \cdot (F(x) \oplus F(x'))} \\ \times \langle \Delta_O \mid F(x) \oplus F(x \oplus \Delta_I) \rangle \langle \Delta_O \mid F(x') \oplus F(x' \oplus \Delta_I) \rangle . \tag{44}$$

Summation over λ_I yields a Kronecker delta $\langle x | x' \rangle$, simplifying the expression to $\sum_{x \in \mathbb{F}_2^n} \langle \Delta_O \mid F(x) \oplus F(x \oplus \Delta_I) \rangle$, which equals $2^n \cdot \mathrm{DDT}_F(\Delta_I, \Delta_O)$.

Combining Theorem 16 with the CEA via event $\{(\Delta_I, \lambda_I, \Delta_O, \lambda_O) : \lambda_I \in \mathbb{F}_2^n, \lambda_O \in \mathbb{F}_2^m\}$ allows us to prepare a state supported by the DDT. To alleviate dilution, we employ the chosen-plaintext technique from Subsection 3.2 by setting $S_I = \{(x_1, x_2) \in \mathbb{F}_2^n \times \mathbb{F}_2^n : x_1 \oplus x_2 = \Delta_I\}$. This yields a distinguisher with query complexity $O(1/\sqrt{p})$, achieving a quadratic speed-up over the classical differential attack.

Equivalence to Kaplan et al.'s approach. The geometric formulation elucidates the quantum differential distinguisher of Kaplan et al. [43]. Their algorithm performs a Grover search for a differential right pair $(x, x \oplus \Delta_I)$ satisfying $F(x) \oplus F(x \oplus \Delta_I) = \Delta_O$. To show the equivalence, we inspect the implementation of change-of-basis operator P, which is a tensor product of n identical two-qubit gates, each implementable with a CNOT, a SWAP, and a Hadamard gate, as illustrated in Figure 4. When combined with the CEA EPDHT (Theorem 10) and the event choice (Theorem 16), the Hadamard gates cancel, rendering the circuit equivalent to the one in [43].

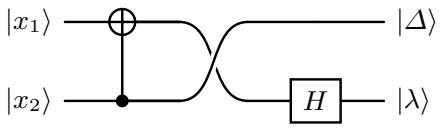


Fig. 4. Quasidifferential change-of-basis operator

A natural question is whether an exponential speed-up is possible for DDT-based attacks. Our geometric analysis suggests that such an improvement would require overcoming a fundamental barrier: the DDT, unlike the LAT, can be efficiently sampled classically, so the exponential speed-up enjoyed by Simon's algorithm does not automatically transfer. Formally, in classical, drawing two random plaintexts x_1, x_2 and obtaining their ciphertexts yields a sample $(\Delta_I, \Delta_O) = (x_1 \oplus x_2, F(x_1) \oplus F(x_2))$ from the DDT. This contrasts with LATs, whose classical sampling is unavailable, so quantum algorithms like Simon's can achieve

exponential speed-ups by exploiting hidden algebraic structures in the LAT. For example, one may consider solving the (vector-output) Bernstein-Vazirani problem using DDT sampling: directly adapt Simon's algorithm to sample (Δ_I, Δ_O) and generate equations $\lambda_I \cdot \Delta_I \oplus \lambda_O \cdot \Delta_O = 0$, since the DDT and LAT are related by a Fourier transform [23]. However, the classical counterpart already has polynomial complexity, so an exponential speed-up is impossible. This observation does not diminish the importance of [43], but rather clarifies the limits of DDT-based quantum attacks.

5.2 Ultrametric Integral Basis

While linear cryptanalysis relies on the *Fourier basis* to diagonalize addition (XOR), integral cryptanalysis is naturally linked to multiplication (bitwise AND). The *ultrametric integral basis* [10] provides a systematic framework for integral properties over fields of characteristic zero, such as $\mathbb{Q}$ or $\mathbb{Q}_2$ (with corresponding non-Archimedean metrics), by capturing divisibility structures in a form analogous to correlation matrices in linear cryptanalysis.

Formally, for the state space $\mathbb{F}_2^n$, the ultrametric integral change-of-basis matrix $\mathcal{U}_n$ has a simple tensor-product structure:

$$\mathcal{U}_n = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}^{\otimes n}, \qquad \mathcal{U}_n^{-1} = \begin{bmatrix} 1 & -1 \\ 0 & 1 \end{bmatrix}^{\otimes n}. \tag{45}$$

The *ultrametric integral transition matrix* of a function $F : \mathbb{F}_2^n \to \mathbb{F}_2^m$ is defined as $C^F = \mathcal{U}_m T^F \mathcal{U}_n^{-1}$. Its entries reflect the divisibility property of the output sum.

Challenges for quantum adaptation. A fundamental obstacle in quantizing the ultrametric integral basis is the mismatch between the underlying number fields. Quantum mechanics operates over the complex Hilbert space $(\mathbb{C}, |\cdot|)$, whereas the ultrametric framework is built on fields with a 2-adic metric (e.g., $\mathbb{Q}_2$). This discrepancy makes it nontrivial to implement the basis transformation $\mathcal{U}_n$ as a unitary quantum gate, which is required for the correlation-extraction algorithm (CEA) to prepare and sample from the correlation state.

Existing quantum methods for integral properties exhibit notable limitations: (1) The algorithm of [53] computes high-order differentials by alternating phase-oracle and CNOT gates, though it requires 2^d superposition queries for a d-th-order differential and thus does not surpass classical methods, leaving room for further improvement. (2) Meanwhile, Hosoyamada [37] successfully speeds up integral attacks based on balanced properties by converting them into zero-correlation linear distinguishers, leveraging results from [12,59]. While this technique is ingenious, it is specifically tailored to balanced properties and does not readily generalize to other divisibility properties. These limitations motivate the search for a more systematic quantization framework.

The geometric approach, though not ready to address it, suggests a deeper reason: integral properties are inherently tied to the 2-adic metric, which is

non-Euclidean and thus not naturally compatible with the unitary formalism of quantum computation. We need to overcome the obstacle of implementing $\mathcal{U}_n$ as a change-of-basis gate and compute under the 2-adic metric, which may require additional machinery. Whether a genuine quantum speed-up can be achieved for any nontrivial integral property remains an open question, inviting further investigation into advanced quantum cryptanalytic techniques.

6 Conclusion

In this work, we have introduced a geometric framework for quantum cryptanalysis, connecting the classical geometric approach and quantum distinguisher design. By extending CEA to operate within a geometric setting over the complex field $\mathbb{C}$, we provide a unified foundation for constructing quantum distinguishers. We recover and generalize many known algorithms, clarifying their geometric origins and the limitations of quantum speed-ups under Fourier, quasidifferential, and ultrametric integral bases.

Since our work establishes a bridge between the classical geometric approach and quantum cryptanalysis, future research can naturally build upon advances in the geometric framework itself. This includes addressing its open challenges, such as the systematic discovery of new bases following the line of exploration initiated in [40]. Any newly proposed basis that operates over $\mathbb{C}$ and admits an efficiently implementable change-of-basis operator can, in principle, be translated into the quantum setting via hidden algebraic structure detection or EPDHT, potentially revealing undetectable properties in concrete block ciphers. Furthermore, an intriguing open question is whether quantum computers could aid the theoretical analysis of classical attacks using the geometric approach, for instance by revealing new cryptanalytic techniques [55].

On the quantum side, several promising avenues warrant further investigation. Many quantum algorithms used in cryptanalysis (e.g., Grover's algorithm and QAA) can be viewed as special instances of quantum random walks. A deeper understanding of quantum random walks might lead to more powerful attack constructions. Moreover, collective measurement techniques and problem-specific Hamiltonian design merit further exploration. We also look forward to future work that extends the framework to other primitives (hash functions, stream ciphers), integrates key recovery mechanisms, or develops automated tools for quantum attack search.

Acknowledgments. We would like to thank all anonymous reviewers for their helpful comments. This research is supported by the National Cyber Security-National Science and Technology Major Project (No. 2026ZD1500700), the National Natural Science Foundation of China (No. 62372262), the National Key R&D Program of China (No. 2023YFA1011200), the Fundamental and Interdisciplinary Disciplines Breakthrough Plan of the Ministry of Education of China (No. JYB2025XDXM114), and the High Performance Computing Center of Tsinghua University.

References

1. Ambainis, A.: Quantum walk algorithm for element distinctness (Apr 2014). https://doi.org/10.48550/arXiv.quant-ph/0311001, http://arxiv.org/abs/quant-ph/0311001, arXiv:quant-ph/0311001
2. Audenaert, K.M.R., Calsamiglia, J., Masanes, L., Munoz-Tapia, R., Acin, A., Bagan, E., Verstraete, F.: The Quantum Chernoff Bound. Physical Review Letters **98**(16), 160501 (Apr 2007). https://doi.org/10.1103/PhysRevLett.98.160501, http://arxiv.org/abs/quant-ph/0610027, arXiv:quant-ph/0610027
3. Bae, J., Kwek, L.C.: Quantum state discrimination and its applications. Journal of Physics A: Mathematical and Theoretical **48**(8), 083001 (Feb 2015). https://doi.org/10.1088/1751-8113/48/8/083001, http://arxiv.org/abs/1707.02571, arXiv:1707.02571 [quant-ph]
4. Bernstein, E., Vazirani, U.: Quantum Complexity Theory. SIAM Journal on Computing **26**(5), 1411–1473 (Oct 1997). https://doi.org/10.1137/S0097539796300921, https://epubs.siam.org/doi/10.1137/S0097539796300921
5. Beyne, T.: A geometric approach to linear cryptanalysis. In: Tibouchi, M., Wang, H. (eds.) Advances in Cryptology – ASIACRYPT 2021, Part I. Lecture Notes in Computer Science, vol. 13090, pp. 36–66. Springer, Cham, Switzerland, Singapore (Dec 6–10, 2021). https://doi.org/10.1007/978-3-030-92062-3_2
6. Beyne, T.: A geometric approach to symmetric-key cryptanalysis. Ph.D. thesis, KU Leuven (Jun 2023), https://cosicdatabase.esat.kuleuven.be/backend/publications/files/these/475
7. Beyne, T., Neyt, A.: Improved differential cryptanalysis of SPEEDY. IACR Transactions on Symmetric Cryptology **2025**(3), 444–474 (2025). https://doi.org/10.46586/tosc.v2025.i3.444-474
8. Beyne, T., Rijmen, V.: Differential cryptanalysis in the fixed-key model. In: Dodis, Y., Shrimpton, T. (eds.) Advances in Cryptology – CRYPTO 2022, Part III. Lecture Notes in Computer Science, vol. 13509, pp. 687–716. Springer, Cham, Switzerland, Santa Barbara, CA, USA (Aug 15–18, 2022). https://doi.org/10.1007/978-3-031-15982-4_23
9. Beyne, T., Verbauwhede, M.: Integral cryptanalysis using algebraic transition matrices. IACR Transactions on Symmetric Cryptology **2023**(4), 244–269 (2023). https://doi.org/10.46586/tosc.v2023.i4.244-269
10. Beyne, T., Verbauwhede, M.: Ultrametric integral cryptanalysis. In: Chung, K.M., Sasaki, Y. (eds.) Advances in Cryptology – ASIACRYPT 2024, Part VII. Lecture Notes in Computer Science, vol. 15490, pp. 392–423. Springer, Singapore, Singapore, Kolkata, India (Dec 9–13, 2024). https://doi.org/10.1007/978-981-96-0941-3_13
11. Beyne, T., Verbauwhede, M.: Integral cryptanalysis in characteristic p. In: Hanaoka, G., Yang, B.Y. (eds.) Advances in Cryptology – ASIACRYPT 2025, Part I. Lecture Notes in Computer Science, vol. 16245, pp. 66–96. Springer, Singapore, Singapore, Melbourne, VIC, Australia (Dec 8–12, 2025). https://doi.org/10.1007/978-981-95-5018-0_3
12. Bogdanov, A., Leander, G., Nyberg, K., Wang, M.: Integral and multidimensional linear distinguishers with correlation zero. In: Wang, X., Sako, K. (eds.) Advances in Cryptology – ASIACRYPT 2012. Lecture Notes in Computer Science, vol. 7658, pp. 244–261. Springer Berlin Heidelberg, Germany, Beijing, China (Dec 2–6, 2012). https://doi.org/10.1007/978-3-642-34961-4_16

13. Bonnetain, X., Hosoyamada, A., Naya-Plasencia, M., Sasaki, Y., Schrottenloher, A.: Quantum attacks without superposition queries: The offline Simon's algorithm. In: Galbraith, S.D., Moriai, S. (eds.) Advances in Cryptology – ASIACRYPT 2019, Part I. Lecture Notes in Computer Science, vol. 11921, pp. 552–583. Springer, Cham, Switzerland, Kobe, Japan (Dec 8–12, 2019). https://doi.org/10.1007/978-3-030-34578-5_20
14. Bonnetain, X., Leurent, G., Naya-Plasencia, M., Schrottenloher, A.: Quantum linearization attacks. In: Tibouchi, M., Wang, H. (eds.) Advances in Cryptology – ASIACRYPT 2021, Part I. Lecture Notes in Computer Science, vol. 13090, pp. 422–452. Springer, Cham, Switzerland, Singapore (Dec 6–10, 2021). https://doi.org/10.1007/978-3-030-92062-3_15
15. Bonnetain, X., Schrottenloher, A.: Single-query quantum hidden shift attacks. IACR Transactions on Symmetric Cryptology **2024**(3), 266–297 (2024). https://doi.org/10.46586/tosc.v2024.i3.266-297
16. Bonnetain, X., Schrottenloher, A., Sibleyras, F.: Beyond quadratic speedups in quantum attacks on symmetric schemes. In: Dunkelman, O., Dziembowski, S. (eds.) Advances in Cryptology – EUROCRYPT 2022, Part III. Lecture Notes in Computer Science, vol. 13277, pp. 315–344. Springer, Cham, Switzerland, Trondheim, Norway (May 30 – Jun 3, 2022). https://doi.org/10.1007/978-3-031-07082-2_12
17. Bootsma, S.E., De Vries, M.: A Survey on the Quantum Security of Block Cipher-Based Cryptography. IEEE Access **12**, 194711–194727 (2024). https://doi.org/10.1109/ACCESS.2024.3520364, https://ieeexplore.ieee.org/document/10807207
18. Boura, C., Derbez, P., Germon, B.: Extending the quasidifferential framework: From fixed-key to expected differential probability. IACR Transactions on Symmetric Cryptology **2025**(1), 515–541 (2025). https://doi.org/10.46586/tosc.v2025.i1.515-541
19. Brassard, G., Hoyer, P., Mosca, M., Tapp, A.: Quantum Amplitude Amplification and Estimation. vol. 305, pp. 53–74 (2002). https://doi.org/10.1090/conm/305/05215, http://arxiv.org/abs/quant-ph/0005055, arXiv:quant-ph/0005055
20. Brassard, G., Hoyer, P., Tapp, A.: Quantum Counting. vol. 1443, pp. 820–831 (1998). https://doi.org/10.1007/BFb0055105, http://arxiv.org/abs/quant-ph/9805082, arXiv:quant-ph/9805082
21. Brassard, G., HØyer, P., Tapp, A.: Quantum cryptanalysis of hash and claw-free functions. In: Lucchesi, C.L., Moura, A.V. (eds.) LATIN'98: Theoretical Informatics. pp. 163–169. Springer, Berlin, Heidelberg (1998). https://doi.org/10.1007/BFb0054319
22. Canale, F., Leander, G., Stennes, L.: Simon's algorithm and symmetric crypto: Generalizations and automatized applications. In: Dodis, Y., Shrimpton, T. (eds.) Advances in Cryptology – CRYPTO 2022, Part III. Lecture Notes in Computer Science, vol. 13509, pp. 779–808. Springer, Cham, Switzerland, Santa Barbara, CA, USA (Aug 15–18, 2022). https://doi.org/10.1007/978-3-031-15982-4_26
23. Chabaud, F., Vaudenay, S.: Links between differential and linear cryptanalysis. In: De Santis, A. (ed.) Advances in Cryptology – EUROCRYPT'94. Lecture Notes in Computer Science, vol. 950, pp. 356–365. Springer Berlin Heidelberg, Germany, Perugia, Italy (May 9–12, 1995). https://doi.org/10.1007/BFb0053450
24. Chang, C., Wang, M., Wang, W., Hu, K.: Quasidifferential saves infeasible differential - improved weak-key key-recovery attacks on round-reduced GIFT. In: Patra, A. (ed.) Topics in Cryptology – CT-RSA 2025. Lecture Notes in Computer Science, vol. 15598, pp. 27–50. Springer, Cham, Switzerland, San Francisco, CA, USA (Apr 28 – May 1, 2025). https://doi.org/10.1007/978-3-031-88661-4_2

25. Cheng, D., Qi, H., Li, Z.: Analysis and Control of Boolean Networks. Communications and Control Engineering, Springer, London (2011). https://doi.org/10.1007/978-0-85729-097-7, http://link.springer.com/10.1007/978-0-85729-097-7
26. Deutsch, D., Jozsa, R.: Rapid solution of problems by quantum computation. Proceedings of the Royal Society of London. Series A: Mathematical and Physical Sciences **439**(1907), 553–558 (Dec 1992). https://doi.org/10.1098/rspa.1992.0167, https://doi.org/10.1098/rspa.1992.0167
27. Dong, X., Dong, B., Wang, X.: Quantum attacks on some feistel block ciphers. Designs, Codes and Cryptography **88**(6), 1179–1203 (2020). https://doi.org/10.1007/s10623-020-00741-y
28. Dubuc, S.: Characterization of linear structures. Designs, Codes and Cryptography **22**(1), 33–45 (2001). https://doi.org/10.1023/A:1008399109102
29. Dutta, S., Jaiswal, A., Maitra, S., Roy, D.: On (Noisy) Simon's (Quantum) Algorithm for Multi-shift Boolean Functions. In: Mesnager, S., Stănică, P., Debnath, S.K. (eds.) Security and Privacy. pp. 62–76. Springer Nature Switzerland, Cham (2025). https://doi.org/10.1007/978-3-031-90587-2_5
30. Even, S., Mansour, Y.: A construction of a cipher from a single pseudorandom permutation. Journal of Cryptology **10**(3), 151–162 (Jun 1997). https://doi.org/10.1007/s001459900025
31. Frixons, P., Naya-Plasencia, M., Schrottenloher, A.: Quantum boomerang attacks and some applications. In: AlTawy, R., Hülsing, A. (eds.) SAC 2021: 28th Annual International Workshop on Selected Areas in Cryptography. Lecture Notes in Computer Science, vol. 13203, pp. 332–352. Springer, Cham, Switzerland, Virtual Event (Sep 29 – Oct 1, 2022). https://doi.org/10.1007/978-3-030-99277-4_16
32. Fulton, W., Harris, J.: Representation Theory. Graduate Texts in Mathematics, Springer, New York, NY (2004). https://doi.org/10.1007/978-1-4612-0979-9, http://link.springer.com/10.1007/978-1-4612-0979-9
33. Grover, L.K.: A fast quantum mechanical algorithm for database search (Nov 1996). https://doi.org/10.48550/arXiv.quant-ph/9605043, http://arxiv.org/abs/quant-ph/9605043, arXiv:quant-ph/9605043
34. Helstrom, C.W.: Quantum detection and estimation theory. Journal of Statistical Physics **1**(2), 231–252 (Jun 1969). https://doi.org/10.1007/BF01007479, https://doi.org/10.1007/BF01007479
35. Hodzic, S., Ramkilde, L.K., Kidmose, A.B.: On quantum distinguishers for type-3 generalized Feistel network based on separability. In: Ding, J., Tillich, J.P. (eds.) Post-Quantum Cryptography - 11th International Conference, PQCrypto 2020. pp. 461–480. Springer, Cham, Switzerland, Paris, France (Apr 15–17, 2020). https://doi.org/10.1007/978-3-030-44223-1_25
36. Holevo, A.S.: Statistical decision theory for quantum systems. Journal of Multivariate Analysis **3**(4), 337–394 (Dec 1973). https://doi.org/10.1016/0047-259X(73)90028-6, https://www.sciencedirect.com/science/article/pii/0047259X73900286
37. Hosoyamada, A.: Quantum speed-up for multidimensional (zero correlation) linear distinguishers. In: Guo, J., Steinfeld, R. (eds.) Advances in Cryptology – ASIACRYPT 2023, Part III. Lecture Notes in Computer Science, vol. 14440, pp. 311–345. Springer, Singapore, Singapore, Guangzhou, China (Dec 4–8, 2023). https://doi.org/10.1007/978-981-99-8727-6_11
38. Hosoyamada, A.: Quantum algorithms for fast correlation attacks on LFSR-based stream ciphers. In: Chung, K.M., Sasaki, Y. (eds.) Advances in Cryptology – ASIACRYPT 2024, Part VIII. Lecture Notes in Computer Science, vol. 15491, pp. 396–430. Springer, Singapore, Singapore, Kolkata, India (Dec 9–13, 2024). https://doi.org/10.1007/978-981-96-0944-4_13

39. Hosoyamada, A., Sasaki, Y.: Quantum Demiric-Selçuk meet-in-the-middle attacks: Applications to 6-round generic Feistel constructions. In: Catalano, D., De Prisco, R. (eds.) SCN 18: 11th International Conference on Security in Communication Networks. Lecture Notes in Computer Science, vol. 11035, pp. 386–403. Springer, Cham, Switzerland, Amalfi, Italy (Sep 5–7, 2018). https://doi.org/10.1007/978-3-319-98113-0_21
40. Hu, K., Zhang, C., Chang, C., Zhang, J., Wang, M., Peyrin, T.: Unlocking mix-basis potential: Geometric approach for combined attacks. In: Kalai, Y.T., Kamara, S.F. (eds.) Advances in Cryptology – CRYPTO 2025, Part V. Lecture Notes in Computer Science, vol. 16004, pp. 293–334. Springer, Cham, Switzerland, Santa Barbara, CA, USA (Aug 17–21, 2025). https://doi.org/10.1007/978-3-032-01901-1_10
41. Ito, G., Hosoyamada, A., Matsumoto, R., Sasaki, Y., Iwata, T.: Quantum chosen-ciphertext attacks against Feistel ciphers. In: Matsui, M. (ed.) Topics in Cryptology – CT-RSA 2019. Lecture Notes in Computer Science, vol. 11405, pp. 391–411. Springer, Cham, Switzerland, San Francisco, CA, USA (Mar 4–8, 2019). https://doi.org/10.1007/978-3-030-12612-4_20
42. Kaplan, M., Leurent, G., Leverrier, A., Naya-Plasencia, M.: Breaking symmetric cryptosystems using quantum period finding. In: Robshaw, M., Katz, J. (eds.) Advances in Cryptology – CRYPTO 2016, Part II. Lecture Notes in Computer Science, vol. 9815, pp. 207–237. Springer Berlin Heidelberg, Germany, Santa Barbara, CA, USA (Aug 14–18, 2016). https://doi.org/10.1007/978-3-662-53008-5_8
43. Kaplan, M., Leurent, G., Leverrier, A., Naya-Plasencia, M.: Quantum differential and linear cryptanalysis. IACR Transactions on Symmetric Cryptology **2016**(1), 71–94 (2016). https://doi.org/10.13154/tosc.v2016.i1.71-94, https://tosc.iacr.org/index.php/ToSC/article/view/536
44. Kilian, J., Rogaway, P.: How to protect DES against exhaustive key search (an analysis of DESX). Journal of Cryptology **14**(1), 17–35 (Jan 2001). https://doi.org/10.1007/s001450010015
45. Kuwakado, H., Morii, M.: Quantum distinguisher between the 3-round Feistel cipher and the random permutation. In: 2010 IEEE International Symposium on Information Theory. pp. 2682–2685 (Jun 2010). https://doi.org/10.1109/ISIT.2010.5513654, https://ieeexplore.ieee.org/document/5513654
46. Kuwakado, H., Morii, M.: Security on the quantum-type Even-Mansour cipher. In: 2012 International Symposium on Information Theory and its Applications. pp. 312–316 (Oct 2012), https://ieeexplore.ieee.org/document/6400943
47. Leander, G., May, A.: Grover meets Simon - quantumly attacking the FX-construction. In: Takagi, T., Peyrin, T. (eds.) Advances in Cryptology – ASIACRYPT 2017, Part II. Lecture Notes in Computer Science, vol. 10625, pp. 161–178. Springer, Cham, Switzerland, Hong Kong, China (Dec 3–7, 2017). https://doi.org/10.1007/978-3-319-70697-9_6
48. Liu, Q., Wang, H., Wang, J., Li, B., Wang, M.: Quantum periodic distinguisher construction: Symbolization method and automated tool. In: Hanaoka, G., Yang, B.Y. (eds.) Advances in Cryptology – ASIACRYPT 2025, Part VIII. Lecture Notes in Computer Science, vol. 16252, pp. 224–255. Springer, Singapore, Singapore, Melbourne, VIC, Australia (Dec 8–12, 2025). https://doi.org/10.1007/978-981-95-5125-5_8
49. Magniez, F., Nayak, A., Roland, J., Santha, M.: Search via quantum walk. In: Proceedings of the thirty-ninth annual ACM symposium on Theory of computing. pp. 575–584. ACM, San Diego California USA (Jun 2007). https://doi.org/10.1145/1250790.1250874, https://dl.acm.org/doi/10.1145/1250790.1250874

50. May, A., Schlieper, L., Schwinger, J.: Noisy simon period finding. In: Paterson, K.G. (ed.) Topics in Cryptology – CT-RSA 2021. Lecture Notes in Computer Science, vol. 12704, pp. 75–99. Springer, Cham, Switzerland, Virtual Event (May 17–20, 2021). https://doi.org/10.1007/978-3-030-75539-3_4
51. Nielsen, M.A., Chuang, I.L.: Quantum computation and quantum information. Cambridge university press, Cambridge, 10th anniversary edition edn. (2010). https://doi.org/10.1017/CBO9780511976667
52. Ozols, M., Roetteler, M., Roland, J.: Quantum rejection sampling. In: Proceedings of the 3rd Innovations in Theoretical Computer Science Conference. pp. 290–308 (Jan 2012). https://doi.org/10.1145/2090236.2090261, http://arxiv.org/abs/1103.2774, arXiv:1103.2774 [quant-ph]
53. SAPV, T., Bera, D., Maitra, A., Maitra, S.: Simon's Algorithm and Autocorrelation Spectrum. In: SAPV, T., Bera, D., Maitra, A., Maitra, S. (eds.) Quantum Algorithms for Cryptographically Significant Boolean Functions: An IBMQ Experience, pp. 89–111. Springer, Singapore (2021). https://doi.org/10.1007/978-981-16-3061-3_4, https://doi.org/10.1007/978-981-16-3061-3_4
54. Selçuk, A.A.: On probability of success in linear and differential cryptanalysis. Journal of Cryptology **21**(1), 131–147 (Jan 2008). https://doi.org/10.1007/s00145-007-9013-7
55. Shi, R., Xie, H., Feng, H., Yuan, F., Liu, B.: Quantum zero correlation linear cryptanalysis. Quantum Information Processing **21**(8), 293 (Aug 2022). https://doi.org/10.1007/s11128-022-03642-2, https://doi.org/10.1007/s11128-022-03642-2
56. Shor, P.: Algorithms for quantum computation: discrete logarithms and factoring. In: Proceedings 35th Annual Symposium on Foundations of Computer Science. pp. 124–134 (Nov 1994). https://doi.org/10.1109/SFCS.1994.365700, https://ieeexplore.ieee.org/document/365700
57. Simon, D.: On the power of quantum computation. In: Proceedings 35th Annual Symposium on Foundations of Computer Science. pp. 116–123 (Nov 1994). https://doi.org/10.1109/SFCS.1994.365701, https://ieeexplore.ieee.org/document/365701
58. Steinberg, B.: Representation Theory of Finite Monoids. Universitext, Springer International Publishing : Imprint: Springer, Cham, 1st ed. 2016 edn. (2016). https://doi.org/10.1007/978-3-319-43932-7
59. Sun, B., Liu, Z., Rijmen, V., Li, R., Cheng, L., Wang, Q., AlKhzaimi, H., Li, C.: Links among impossible differential, integral and zero correlation linear cryptanalysis. In: Gennaro, R., Robshaw, M.J.B. (eds.) Advances in Cryptology – CRYPTO 2015, Part I. Lecture Notes in Computer Science, vol. 9215, pp. 95–115. Springer Berlin Heidelberg, Germany, Santa Barbara, CA, USA (Aug 16–20, 2015). https://doi.org/10.1007/978-3-662-47989-6_5
60. Szegedy, M.: Quantum speed-up of Markov chain based algorithms. In: 45th Annual IEEE Symposium on Foundations of Computer Science. pp. 32–41 (Oct 2004). https://doi.org/10.1109/FOCS.2004.53, https://ieeexplore.ieee.org/abstract/document/1366222
61. Wang, L., Song, L., Wu, B., Rahman, M., Isobe, T.: Revisiting the Boomerang Attack From a Perspective of 3-Differential. IEEE Transactions on Information Theory **70**(7), 5343–5357 (Jul 2024). https://doi.org/10.1109/TIT.2023.3324738, https://ieeexplore.ieee.org/document/10292670
62. Xiang, Z., Wang, X., Yu, B., Sun, B., Zhang, S., Zeng, X., Shen, X., Li, N.: Links between quantum distinguishers based on simon's algorithm and truncated differentials. IACR Transactions on Symmetric Cryptology **2024**(2), 296–321 (2024). https://doi.org/10.46586/tosc.v2024.i2.296-321

63. Xie, H., Yang, L.: Using bernstein-vazirani algorithm to attack block ciphers. Designs, Codes and Cryptography **87**(5), 1161–1182 (2019). https://doi.org/10.1007/s10623-018-0510-5
64. Zhandry, M.: How to model unitary oracles. In: Kalai, Y.T., Kamara, S.F. (eds.) Advances in Cryptology – CRYPTO 2025, Part II. Lecture Notes in Computer Science, vol. 16001, pp. 237–268. Springer, Cham, Switzerland, Santa Barbara, CA, USA (Aug 17–21, 2025). https://doi.org/10.1007/978-3-032-01878-6_8

A On QSD for Distinguishers

This section present a brief discussion about QSD compared to CEA-based EPDHT distinguishers (Theorem 9) by giving a trivial example (Example 6), to demonstrate the advantage from the permit of querying oracle O_F. In other words, it may be suboptimal to apply QSD directly as a subroutine of quantum cryptanalytic distinguishers.

Based on quantum cryptanalysis practice, we only consider pure states. In addition, we distinguish between two states $|\psi_1\rangle, |\psi_2\rangle$ with uniform a prior probability. For a single copy of the state to be distinguished, the minimal error probability p_{err}^{min} is characterized by the Helstrom bound [34,36] as

$$p_{err}^{\min} = \frac{1}{2}\left(1 - \sqrt{1 - |\langle\psi_1 \mid \psi_2\rangle|^2}\right). \tag{46}$$

For multiple-copy discrimination (needs multiple qubits accordingly), $p_{err}^{\min}$ follows the quantum Chernoff bound [2]. Suppose we have n copies, then

$$p_{err}^{\min} \sim |\langle\psi_1 \mid \psi_2\rangle|^{2n}. \tag{47}$$

A trivial example demonstrates that the speed-up from QSD Chernoff bound is less than the speed-up from EPDHT. Note that QSD requires full-state knowledge.

Example 6. Let

$$\begin{aligned} |\psi_1\rangle &:= \cos\theta_1 |0\rangle + \sin\theta_1 |1\rangle, \\ |\psi_2\rangle &:= \cos\theta_2 |0\rangle + \sin\theta_2 |1\rangle, \end{aligned} \tag{48}$$

where $\theta_1, \theta_2 \in [0, \pi/4]$, $\theta_1 < \theta_2$ and $\Delta\theta := |\theta_1 - \theta_2| \ll 1$. Note that we assume $\theta_2 \leq \pi/4$ without loss of generality because of the sine-cosine symmetry. The example corresponds to Theorem 9 by $p_1 := \sin^2\theta_1, p_2 := \sin^2\theta_2$.

QSD and EPDHT are applied separately to distinguish between $|\Psi_1\rangle$ and $|\Psi_2\rangle$, provided an oracle to prepare the unknown state. Observe that

$$\langle\psi_1 \mid \psi_2\rangle = \cos\Delta\theta \approx 1 - \frac{1}{2}\Delta\theta^2. \tag{49}$$

To achieve a constant-order success probability, p_{err} should be also of constant order. Hence

$$n_{\mathrm{QSD}} \sim -\frac{1}{\log|\langle\psi_1 \mid \psi_2\rangle|} = O(\frac{1}{\Delta\theta^2}). \tag{50}$$

The EPDHT query complexity is

$$n_{\mathrm{EPDHT}} = \frac{\pi}{\sin\theta_2 - \sin\theta_1} \approx \frac{\pi}{\cos\theta_1 \cdot \Delta\theta} = O\left(\frac{1}{\Delta\theta}\right). \tag{51}$$

Therefore, EPDHT achieves a quadratic improvement in query complexity over QSD in this regime.

This demonstrates that, when oracle queries are permitted, QAE-based distinguishers can be significantly more efficient than optimal state discrimination, especially for small probability deviations.

B Subset State Preparation via Quantum Rejection Sampling

We first recite [52, Theorem 1] and then apply it to subset states.

Definition 12 (Quantum resampling [52]). *Let $\pi, \sigma \in \mathbb{R}^n$ be fixed vectors with non-negative components, $\|\pi\|_2 = \|\sigma\|_2 = 1$ and $\pi_k > 0$ for all $k \in \{1, 2, \ldots, n\}$. For an unknown list of normalized quantum states $\xi = (|\xi_1\rangle, \ldots, |\xi_n\rangle)$ with $|\xi_k\rangle \in \mathbb{C}^d$, the oracle O_ξ satisfies*

$$O_\xi |0\rangle = |\pi^\xi\rangle := \sum_{k=1}^{n} \pi_k |\xi_k\rangle |k\rangle . \tag{52}$$

Given oracle access to O_ξ and its inverse, the task $\mathrm{QSAMPLING}_{\pi\to\sigma}$ *is to prepare the target state*

$$|\sigma^\xi\rangle := \sum_{k=1}^{n} \sigma_k |\xi_k\rangle |k\rangle \tag{53}$$

with success probability p; i.e. to output a state ρ such that $\langle\sigma^\xi|\rho|\sigma^\xi\rangle \geq p$.

Definition 13 (Water-filling vector [52]). *For π, σ define*

$$\begin{aligned} p_{\min} &:= (\sigma^\top \pi)^2, \quad \gamma_{\min} := \min_{k:\pi_k>0} (\pi_k/\sigma_k), \\ p_{\max} &:= \sum_{k:\pi_k>0} \sigma_k^2, \quad \gamma_{\max} := \max_k (\pi_k/\sigma_k). \end{aligned} \tag{54}$$

For $\gamma \in [\gamma_{\min}, \gamma_{\max}]$ set

$$\epsilon_k(\gamma) := \min\{\pi_k, \gamma\sigma_k\}, \qquad p(\gamma) := \Big(\frac{\sigma^\top \epsilon(\gamma)}{\|\epsilon(\gamma)\|_2}\Big)^2. \tag{55}$$

For any $p \in [p_{\min}, p_{\max}]$ take $\bar{\gamma}$ with $p(\bar{\gamma}) = p$ and put the water-filling vector by

$$\epsilon^p_{\pi\to\sigma} := \epsilon(\bar{\gamma}). \tag{56}$$

Theorem 17 (Quantum rejection sampling [52]). *For $p \in [p_{\min}, p_{\max}]$ the bounded-error quantum query complexity of* $\mathrm{QSAMPLING}_{\pi\to\sigma}$ *is*

$$Q_{1-p}\big(\mathrm{QSAMPLING}_{\pi\to\sigma}\big) = \Theta\bigg(\frac{1}{\|\epsilon^{p}_{\pi\to\sigma}\|_2}\bigg). \tag{57}$$

Each call to O_ξ or $O_\xi^\dagger$ counts as one query. If $p < p_{\min}$ the complexity is 1; if $p > p_{\max}$ the problem is infeasible.

Let $T \subseteq \mathbb{F}_2^n$ be a non-empty subset and suppose we are given an oracle O_T that prepares the uniform superposition

$$O_T\,|0\rangle = |T\rangle := |T|^{-1/2} \sum_{x\in T} |x\rangle\,. \tag{58}$$

Our goal is to prepare the target subset state $|S\rangle$ with $S \subseteq T$ and

$$|S\rangle := |S|^{-1/2} \sum_{x\in S} |x\rangle\,, \tag{59}$$

and we require success probability $p = 1$ (exact preparation).

Theorem 18. *$|S\rangle$ can be exactly prepared from $|T\rangle$ using $\Theta(\sqrt{|T|/|S|})$ queries to O_T and $O_T^\dagger$.*

Proof. We cast this problem as an instance of $\mathrm{QSAMPLING}_{\pi\to\sigma}$ (Definition 12). Define vectors $\pi, \sigma \in \mathbb{R}^{2^n}$ by

$$\pi_x = \begin{cases} |T|^{-1/2}, & x \in T, \\ 0, & x \notin T, \end{cases} \qquad \sigma_x = \begin{cases} |S|^{-1/2}, & x \in S, \\ 0, & x \notin S. \end{cases} \tag{60}$$

Obviously $\|\pi\|_2 = \|\sigma\|_2 = 1$ and $\pi_x, \sigma_x \geq 0$. Set the unknown states $|\xi_x\rangle = |0\rangle$ (the computational basis). Because the state $|0\rangle$ is constant and unentangled, it can be omitted from the notation without ambiguity. Thus we identify $|\pi^\xi\rangle = |T\rangle$ and the unknown "states" $|\xi_x\rangle$ are trivial. From Definition 13 we compute the relevant parameters:

$$p_{\min} = (\sigma^\top \pi)^2 = \Big(\sum_{x\in S} |T|^{-1/2}|S|^{-1/2}\Big)^2 = \frac{|S|}{|T|}, \tag{61}$$

$$p_{\max} = \sum_{x:\pi_x>0} \sigma_x^2 = \sum_{x\in S} \frac{1}{|S|} = 1. \tag{62}$$

We aim for success probability $p = 1$. Recall the definitions

$$\gamma_{\min} = \min_{k:\pi_k>0} (\pi_k/\sigma_k), \qquad \gamma_{\max} = \max_k (\pi_k/\sigma_k). \tag{63}$$

For $x \in S$ we have $\pi_x/\sigma_x = \sqrt{|S|/|T|}$; for $x \in T \setminus S$ the ratio is undefined (division by zero) and such indices are excluded from the min in $\gamma_{\min}$, while in $\gamma_{\max}$ they contribute an infinite value. Hence

$$\gamma_{\min} = \sqrt{\frac{|S|}{|T|}}, \qquad \gamma_{\max} = +\infty. \tag{64}$$

For a "water-level" $\gamma \in [\gamma_{\min}, \gamma_{\max}]$ we set

$$\epsilon_k(\gamma) = \min\{\pi_k, \gamma\sigma_k\}, \qquad p(\gamma) = \Big(\frac{\sigma^\top \epsilon(\gamma)}{\|\epsilon(\gamma)\|_2}\Big)^2. \tag{65}$$

Choosing $\bar{\gamma} = \gamma_{\min} = \sqrt{|S|/|T|}$ we obtain

$$\epsilon_x(\bar{\gamma}) = \begin{cases} \pi_x = |T|^{-1/2}, & x \in S, \\ 0, & x \notin S, \end{cases} \tag{66}$$

hence

$$\|\epsilon(\bar{\gamma})\|_2 = \sqrt{|S|} \cdot |T|^{-1/2} = \sqrt{\frac{|S|}{|T|}}, \qquad \sigma^\top \epsilon(\bar{\gamma}) = \sum_{x \in S} \frac{1}{\sqrt{|S|}} \cdot \frac{1}{\sqrt{|T|}} = \sqrt{\frac{|S|}{|T|}}, \tag{67}$$

and therefore $p(\bar{\gamma}) = 1$. Thus $\epsilon^1_{\pi\to\sigma} = \epsilon(\bar{\gamma})$ and $\|\epsilon^1_{\pi\to\sigma}\|_2 = \sqrt{|S|/|T|}$.

Thus, by applying Theorem 17 we obtain the exact quantum query complexity for preparing $|S\rangle$ with certainty:

$$Q_0\big(\mathrm{QSAMPLING}_{\pi\to\sigma}\big) = \Theta\Big(\frac{1}{\|\epsilon^1_{\pi\to\sigma}\|_2}\Big) = \Theta\left(\sqrt{\frac{|T|}{|S|}}\right). \tag{68}$$

C Proof of Theorem 9

We prove the theorem in four steps.

I. Classical data complexity. For distinguishing between two distributions D_1 and D_2 based on i.i.d. samples, the optimal asymptotic error rate is governed by the Kullback-Leibler (KL) divergence, implemented via the log-likelihood ratio (LLR) test. For binary hypothesis testing, the required number of classical samples N_C satisfies $N_C = \Omega\left(1/D_{\mathrm{KL}}(D_1\|D_2)\right)$, where

$$D_{\mathrm{KL}}(D_1\|D_2) := \sum_{x \in S} D_1(x) \log \frac{D_1(x)}{D_2(x)}. \tag{69}$$

For EPDHT, consider the Bernoulli random variables $B_1 \sim \mathrm{Bern}(p_1)$ and $B_2 \sim \mathrm{Bern}(p_2)$ obtained by checking whether a sample falls in G. Since we have no

knowledge other than p_1 and p_2, it is reasonable to consider the worst-case distinguishing scenario, where the two distributions are identical outside G. Therefore

$$D_{\mathrm{KL}}(D_1\|D_2) = D_{\mathrm{KL}}(B_1\|B_2) = p_1 \log \frac{p_1}{p_2} + (1-p_1)\log\frac{1-p_1}{1-p_2}. \tag{70}$$

Hence, the classical sample complexity N_C is at least $\Omega(1/D_{\mathrm{KL}}(B_1\|B_2))$.

II. Quantum data complexity. Meanwhile, on quantum computers, we can apply QAE. We estimate p_1 (i.e., to tell hypothesis 1 from hypothesis 2 and obtain a constant-order advantage) with N queries to O and $O^\dagger$ each, after which we obtain an estimation $\tilde{p_1}$ whose distance to p_1 is bounded with high probability by Theorem 5:

$$|p_1 - \tilde{p_1}| \le \frac{2\pi\sqrt{p_1(1-p_1)}}{N} + \frac{\pi^2}{N^2}. \tag{71}$$

Now we try to bound $\tilde{p_1} \le p_2$. It suffices if

$$\frac{2\pi\sqrt{p_1(1-p_1)}}{N} + \frac{\pi^2}{N^2} \le \frac{2\pi\sqrt{p_1}}{N} + \frac{\pi^2}{N^2} \le p_2 - p_1. \tag{72}$$

Solve the rightmost inequality in Equation (72) for N:

$$N \ge \frac{\pi}{\sqrt{p_2}-\sqrt{p_1}}. \tag{73}$$

Therefore, the quantum query complexity is $N_Q = 2N = 2\pi/(\sqrt{p_2}-\sqrt{p_1})$.

III. Regime 1. Since $p_1 \ll p_2 \ll 1$, we have $N_Q = 2\pi/(\sqrt{p_2}-\sqrt{p_1}) \approx 2\pi/\sqrt{p_2}$. For the classical counterpart, we approximate the KL divergence using $\log(1+x) \approx$ and $x\log x \approx 0$ for small x:

$$D_{\mathrm{KL}}(B_1\|B_2) \approx -\log(1-p_2) \approx p_2. \tag{74}$$

Hence, the classical data complexity N_C is at least $\Theta(1/p_2)$. Comparing N_C and N_Q, we achieve a quantum quadratic speed-up.

IV. Regime 2. In this case, we express complexities by the probability deviation $\varepsilon := |p_1 - p_2|$, because it is a common symbol in cryptanalysis, e.g. linear attacks. The quantum data complexity is

$$N_Q = \frac{2\pi}{\sqrt{p_2}-\sqrt{p_1}} = \frac{2\pi(\sqrt{p_1}+\sqrt{p_2})}{\varepsilon} = \Theta(1/\varepsilon). \tag{75}$$

Expanding the KL divergence for small ε using $\log(1+x) = x - x^2/2 + o(x^3)$:

$$\begin{aligned} D_{\mathrm{KL}}(B_1\|B_2) =& -p_1\log\left(1+\frac{\varepsilon}{p_1}\right) - (1-p_1)\log\left(1-\frac{\varepsilon}{1-p_1}\right) \\ \approx& -p_1\left(\frac{\varepsilon}{p_1} - \frac{\varepsilon^2}{p^2}\right) - (1-p_1)\left(-\frac{\varepsilon}{1-p_1} - \frac{\varepsilon^2}{(1-p_1)^2}\right) \\ =& \frac{\varepsilon^2}{p_1(1-p_1)} = \Theta(\varepsilon^2). \end{aligned} \tag{76}$$

The classical complexity is hence $N_C = \Theta(1/\varepsilon^2)$. Again, we achieve a quantum quadratic speed-up.

D Hosoyamada's Advanced Quantum Multidimensional Linear Distinguishers

This section recalls the advanced quantum multidimensional linear distinguishers in [37], i.e. distinguishers $\mathcal{A}_2, \mathcal{A}_3$, and shows its connection with the geometric approach, especially countermeasures to the *dilution* problem. The discussions can naturally transfer to its multidimensional zero-correlation distinguishers $\mathcal{B}_2, \mathcal{B}_3$, and we will not repeat them.

Let $F : \mathbb{F}_2^n \to \mathbb{F}_2^m$ be a function, and let $V \subset \mathbb{F}_2^m \times \mathbb{F}_2^n$ be a vector space of input-output masks that will be used for the multidimensional linear attack.

Definition 14 (Linearly independent masks). *The linear masks are said to be linearly independent if V can be decomposed as a direct product $V = V_1 \times V_2$ for some subspaces $V_1 \subset \mathbb{F}_2^n, V_2 \subset \mathbb{F}_2^m$.*

For linearly independent masks, the basis of V can be chosen as the union of bases of V_1 and V_2.

Definition 15. *The linear masks are linearly completely dependent if there exists a basis $\{(\alpha_i, \beta_i)\}_{1 \le i \le \dim(V)}$ of V such that both $\{\alpha_i\}$ and $\{\beta_i\}$ are linearly independent in $\mathbb{F}_2^n$.*

In this case, there is an isomorphism between the input mask space and the output mask space induced by this basis. These two structures allow nontrivial quantum speed-ups beyond the generic quadratic advantage of Grover's search.

Both distinguishers are instances of the generic QAA-based procedure $\mathcal{QD}$ of [37, Proposition 2] with carefully chosen unitary U and Boolean predicate that flips the masked states (corresponding to O_G in Theorem 10).

For simplicity, we consider the case where the linear capacity is known and fixed. Let $E : \mathbb{F}_2^n \to \mathbb{F}_2^m$ be the target block cipher, and let V be the vector space of masks for which the sum of squared correlations satisfies

$$\sum_{(\alpha,\beta) \in V \setminus \{0\}} \mathrm{LAT}_E(\alpha, \beta)^2 = \mathrm{Cap}. \tag{77}$$

Now we examine distinguishers $\mathcal{A}_2, \mathcal{A}_3$.

- $\mathcal{A}_2$*: linearly independent masks.* Assume $V = V_1 \times V_2$ with $\dim(V_1) = u$, $\dim(V_2) = w$. Without loss of generality we apply an affine transformation on the output of the cipher so that $V_2 = \{\beta \| 0^{n-w} \mid \beta \in \mathbb{F}_2^w\}$ and $\beta_i = e_i := 0^{i-1}10^{w-i}$. We truncate the output of E to the first w bits and apply QAA. The algorithm succeeds with high probability and query complexity $\Theta(2^{w/2}/\mathrm{Cap}^{1/2})$.

– $\mathcal{A}_3$*: linearly completely dependent masks.* Now V admits a basis $\{(\alpha_i, \beta_i)\}_{i=1}^{\dim(V)}$ with both $\{\alpha_i\}$ and $\{\beta_i\}$ linearly independent. We can again apply an affine transformation on the output of the cipher to make $\beta_i = e_i$. Similarly, the output of E is truncated into its first $w = \dim(V)$ bits. Hence, the algorithm succeeds with high probability and query complexity $\Theta(2^{w/2}/\mathrm{Cap}^{1/2})$.

As hinted in [37], the three distinguishers $\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3$ can be unified by restricting inputs and outputs appropriately. We illustrate this unification for general space V, with linearly independent and linearly completely dependent masks being special cases. Formally, for an arbitrary mask space $V \subseteq \mathbb{F}_2^n \times \mathbb{F}_2^m$, define its projection onto the output part:

$$V_2 := \{\, \beta \in \mathbb{F}_2^m \mid \exists \alpha \in \mathbb{F}_2^n : (\alpha, \beta) \in V \,\}. \tag{78}$$

Let $w = \dim(V_2)$. Choose a linear isomorphism $L : \mathbb{F}_2^m \to \mathbb{F}_2^m$ such that $L(V_2) = \mathbb{F}_2^w \times \{0^{m-w}\}$. By applying L to the ciphertext (i.e. replacing the cipher with $L \circ E$), we may assume without loss of generality that

$$V_2 = \{\, \beta \| 0^{m-w} \mid \beta \in \mathbb{F}_2^w \,\}. \tag{79}$$

Then every mask $(\alpha, \beta) \in V$ satisfies that the last $n - w$ bits of β are zero. Then we truncate the function output. It is easy to verify that output truncation retains the linear correlations and thus the linear capacity Cap. By [37] (which agrees with Theorem 10), the algorithm requires $\Theta(2^{w/2}/\mathrm{Cap}^{1/2})$ superposition queries.

Dilution and its countermeasure. As discussed in Section 3, the linear correlation state $|\Psi_F\rangle$ suffers from a severe dilution problem: the amplitude of the distinguishing signal is spread over 2^m rows (the output mask space), making a naive application of the CEA (as in $\mathcal{A}_1$) inefficient. This necessitates countermeasures that concentrate the discriminating information onto a smaller set of basis states.

The distinguishers $\mathcal{A}_2$ and $\mathcal{A}_3$ introduced in [37] address this issue by truncating the output bits of the function. From the geometric viewpoint, truncation corresponds to restricting the output mask space to a subspace $V_2 \subseteq \mathbb{F}_2^m$ (see Equation (78)) and applying the Fourier change-of-basis only on the remaining $w = \dim(V_2)$ qubits. In the original implementation, this truncation requires two queries to O_F: one for computing the ciphertext and one for uncomputing after the desired bit slice has been copied via CNOT gates, as depicted in Figure 5.

Within our geometric framework, the same truncation can be performed with a single query by employing the phase-oracle technique developed in Subsection 4.1, recovering the technique of [39]. Concretely, choosing the initial state as $|0\rangle\,|S_O\rangle$ with $S_O = V_2$ in Theorem 13 prepares a reduced correlation state $|\Psi_F'\rangle$ that retains only those rows of C^F indexed by $\lambda_O \in V_2$. The entire procedure uses exactly one call to O_F (and its inverse for uncomputation). Hence the query complexity of $\mathcal{A}_2$ and $\mathcal{A}_3$ remains $\Theta(2^{w/2}/\sqrt{\mathrm{Cap}})$, while the circuit depth (and

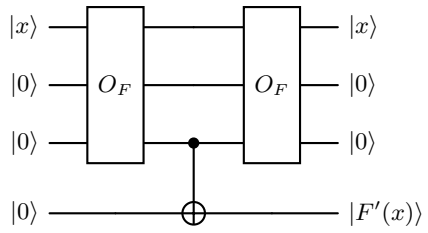


Fig. 5. Truncation by CNOT

the non-asymptotic complexity) is halved compared to the two-query CNOT-based approach. The circuit is depicted in Figure 6, where the Hadamard gate on the third register are cancelled in practice, thus recovering the truncation technique in [39].

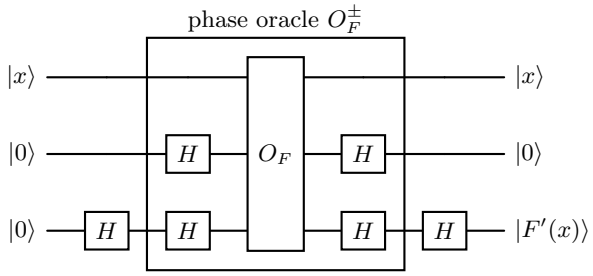


Fig. 6. Truncation by phase oracles

More generally, the phase-oracle formalism allows us to consider *arbitrary* subsets $S_O \subseteq \mathbb{F}_2^m$, not merely subspaces, provided that the subset state $|S_O\rangle$ can be prepared efficiently (cf. Subsection 3.2). This flexibility provides a more general countermeasure to dilution, as discussed in Subsection 4.1: any set of rows in the correlation matrix can be suppressed by an appropriate choice of S_O, typically using the techniques of quantum rejection sampling [52]. Thus the geometric CEA together with phase-oracle state preparation subsumes and generalizes the truncation technique of [37] while offering a unified algebraic interpretation of the achieved speed-up.

E Proof of Theorem 14

We compute the action of O'_F on a basis state:

$$\begin{aligned} O'_F \left|x\right\rangle \left|\lambda_O\right\rangle &= \left(I^{\otimes n} \otimes H_2\right) O_F \left(I^{\otimes n} \otimes H_1\right) \left|x\right\rangle \left|\lambda_O\right\rangle \\ &= \left(I^{\otimes n} \otimes H_2\right) O_F \left|x\right\rangle \left(H_1 \left|\lambda_O\right\rangle\right) \\ &= \left(I^{\otimes n} \otimes H_2\right) \sum_{y \in \mathbb{F}_2^m} \left\langle y|H_1|\lambda_O\right\rangle \left|x\right\rangle \left|y \oplus F(x)\right\rangle \\ &= \left|x\right\rangle \otimes \sum_{y, \lambda'_O \in \mathbb{F}_2^m} \left\langle \lambda'_O|H_2|y \oplus F(x)\right\rangle \left\langle y|H_1|\lambda_O\right\rangle \left|\lambda'_O\right\rangle . \end{aligned} \tag{80}$$

By assumption, the output must be proportional to $\left|x\right\rangle \left|\lambda_O\right\rangle$; hence for all $\lambda'_O \neq \lambda_O$ the coefficient of $\left|\lambda'_O\right\rangle$ must vanish:

$$\sum_{y \in \mathbb{F}_2^m} \left\langle \lambda'_O|H_2|y \oplus F(x)\right\rangle \left\langle y|H_1|\lambda_O\right\rangle = 0, \qquad \text{if } \lambda'_O \neq \lambda_O. \tag{81}$$

Because $F(x)$ can be chosen arbitrarily, we may set $y' := F(x)$ to be any element of $\mathbb{F}_2^m$. Introducing the translation operator $T^{y'} := \sum_{y \in \mathbb{F}_2^m} \left|y \oplus y'\right\rangle \left\langle y\right|$, the condition becomes

$$\left\langle \lambda'_O|H_2\, T^{y'}\, H_1|\lambda_O\right\rangle = 0, \qquad \text{if } \lambda'_O \neq \lambda_O. \tag{82}$$

Thus $H_2 T^{y'} H_1$ is a diagonal unitary for every y'. Denote this diagonal matrix by $D^{y'}$; then $T^{y'} = H_2^\dagger D^{y'} H_1^\dagger$, which is a singular value decomposition (SVD) of $T^{y'}$. On the other hand, $T^{y'}$ admits the spectral decomposition

$$T^{y'} = H^{\otimes m} \cdot \left(\sum_{\lambda \in \mathbb{F}_2^m} (-1)^{\lambda \cdot y'} \left|\lambda\right\rangle \left\langle \lambda\right| \right) \cdot H^{\otimes m}. \tag{83}$$

SVD is unique up to phase factors on the singular vectors and permutations within degenerate subspaces. Consider H_1, for any $\lambda' \in \mathbb{F}_2^m$, $H_1 \left|\lambda'\right\rangle$ lies in the subspace of single value $\lambda' \cdot y'$, i.e., $\mathrm{Span}\{H^{\otimes m} \left|\lambda\right\rangle : \lambda \cdot y' = \lambda' \cdot y'\}$. Taking intersections over all y' gives

$$H_1 \left|\lambda'\right\rangle \in \mathrm{Span}\left\{ H^{\otimes m} \left|\lambda\right\rangle : \lambda \cdot y' = \lambda' \cdot y', \forall y' \in \mathbb{F}_2^m \right\}. \tag{84}$$

Hence $\lambda = \lambda'$, resulting in $H_1 \left|\lambda'\right\rangle = d_{\lambda'} H^{\otimes m} \left|\lambda'\right\rangle$ for some phase factor $d_{\lambda'}$. Therefore, $H_1 = H^{\otimes m} D_1$ for some diagonal unitary D_1. A symmetric argument on $H_2^\dagger$ yields $H_2 = D_2 H^{\otimes m}$ for some diagonal unitary D_2.

F Proof of Theorem 15

We follow a scheme similar to [42, Theorem 1]. Its proof cannot automatically transfer to Theorem 15 due to the differences in its and our preparation-measurement template, so we present this proof for self-containment.

Notice that for a random variable $z \in \mathbb{F}_2$, its probability to be zero is $\Pr[z = 0] = \mathbb{E}[\langle z|0\rangle]$. Meanwhile, $\langle z|0\rangle = (1 + (-1)^z)/2$. Taking $z = \lambda_I \cdot \Delta_I \oplus \lambda_O \cdot \Delta_O$ for Equation (38), it therefore remains to prove

$$\mathbb{E}\left[(-1)^{\lambda_I \cdot \Delta_I \oplus \lambda_O \cdot \Delta_O}\right] = \Pr_x[F(x) \oplus F(x \oplus \Delta_I) = \Delta_O]. \tag{85}$$

The probability distribution over (λ_I, λ_O) is given by the squared amplitudes of the correlation state:

$$\Pr[\lambda_I \| \lambda_O] = \left|\langle \lambda_I, \lambda_O \mid \Psi_F \rangle\right|^2. \tag{86}$$

Substituting the expression for $\Pr[\lambda_I \| \lambda_O]$ (cf. Equation (12)) and simplifying using the orthogonality of characters yields

$$\begin{aligned}
&\sum_{\lambda_I \in \mathbb{F}_2^n} \sum_{\lambda_O \in \mathbb{F}_2^m} \Pr[\lambda_I || \lambda_O] \cdot (-1)^{\lambda_I \cdot \Delta_I \oplus \lambda_O \cdot \Delta_O} \\
&= \sum_{\lambda_I \in \mathbb{F}_2^n} \sum_{\lambda_O \in \mathbb{F}_2^m} \sum_{x, x' \in \mathbb{F}_2^n} 2^{-(2n+m)} \cdot (-1)^{\lambda_I \cdot \left(x \oplus x' \oplus \Delta_I\right) \oplus \lambda_O \cdot \left(F(x) \oplus F(x') \oplus \Delta_O\right)} \\
&= \sum_{x, x' \in \mathbb{F}_2^n} 2^{-n} \langle 0 \mid x \oplus x' \oplus \Delta_I \rangle \langle 0 \mid F(x) \oplus F(x') \oplus \Delta_O \rangle \\
&= \Pr_x[F(x) \oplus F(x \oplus \Delta_I) = \Delta_O],
\end{aligned} \tag{87}$$

which establishes (85) and completes the proof.