

Improved Quantum Linear Key-Recovery Attacks with Walsh-Spectrum Puncturing

Qihang Ding¹, Jiongjiong Ren¹, Shaozheng He¹, and Shaozhen Chen¹

Information Engineering University, Zhengzhou, 450001, Henan, China
jiongjiong_fun@163.com

Abstract. Quantum linear key-recovery attacks using QFT-based correlation state preparation can achieve superquadratic speedups, but their complexity grows rapidly with the active subkey dimension, potentially erasing the advantage over Grover’s search. WSP (EUROCRYPT 2024) provides a natural mechanism for alleviating this bottleneck by replacing the key-recovery map with a spectrally punctured approximation of smaller active dimension.

In this work, we establish a puncture-aware framework for QFT-based quantum key recovery, extending quantum linear cryptanalysis to key-recovery problems whose active-subkey dimension can be reduced through spectral puncturing. We further propose Q-Puncture, an automated optimization method tailored to quantum key-recovery complexity. Applied to Serpent, RECTANGLE-128, and PIPO-128, Q-Puncture yields a new family of punctured quantum linear key-recovery attacks whose modeled quantum complexities improve upon generic Grover search. The resulting attacks improve upon comparable QFT-based results, extend them to new settings, and provide further examples in which structured quantum cryptanalysis surpasses the generic quadratic speedup of Grover’s search.

Keywords: Quantum Cryptanalysis, Linear Cryptanalysis, Quantum Fourier Transform, Walsh-Spectrum Puncturing

1 Introduction

As a fundamental quantum algorithm for unstructured search, Grover’s algorithm reduces exhaustive search over a K -bit key space from 2^K classical evaluations to roughly $2^{K/2}$ coherent evaluations, and therefore provides the standard baseline for quantum key-recovery attacks against symmetric primitives [22]. Many quantum cryptanalytic results obtain a comparable quadratic speedup by adapting established classical techniques, including differential and linear cryptanalysis [26], meet-in-the-middle attacks [35] and rebound attacks [2] to the quantum setting. Improvements beyond this baseline are much less common and typically rely on exploitable algebraic or periodic structure, as illustrated by Simon-type period-finding [37] and quantum slide attacks [10], hidden-shift

techniques [9], Q1 attacks derived from offline period- or shift-finding methods [11,29], and automated searches for Simon-type attacks [13]. Quantum key-recovery attacks achieving superquadratic speedups are therefore of particular importance in quantum cryptanalysis.

Quantum linear key recovery. Linear cryptanalysis exploits statistical correlations among plaintext, ciphertext, and key bits to distinguish the correct key hypothesis from incorrect ones [30,5]. Early quantum linear attacks used quantum amplitude estimation to accelerate correlation evaluation, thereby obtaining a quadratic speedup over their classical counterparts [26]. At CRYPTO 2023, Schrottenloher introduced a QFT-based correlation state preparation algorithm in which candidate-subkey amplitudes are proportional to experimental correlations, yielding gains of about 1.58 and 2.81 bits over Grover’s search on FLY and PIPO, respectively [34]. In Schrottenloher’s original framework, active-subkey candidates are extended to full master-key candidates for verification. As later emphasized by Bashiri et al., this completion step is most direct when the key schedule admits an efficient completion procedure. They generalized the key-identification step by separating correlation-state preparation from candidate identification and introducing coherent distinguishers to recognize the correct active subkey [3]. The resulting QFT-based key-recovery framework, however, remains sensitive to the active subkey dimension. Reducing this dimension while retaining sufficient correlation signal can lower the quantum key-recovery cost and thereby preserve an advantage over Grover’s search.

Advanced key-recovery attacks based on Walsh-spectrum puncturing. Classical linear cryptanalysis has developed increasingly advanced key-recovery techniques to reduce the cost of searching large subkey spaces. FFT- and FWT-based fast key-recovery attacks evaluate the scores of many subkey hypotheses simultaneously, substantially reducing the time complexity of the key-recovery phase [17,18,39]. Multiple and multidimensional linear key-recovery attacks further improve the separation between correct and wrong subkey candidates by combining information from several linear approximations [23,15,28]. More recent techniques reduce the cost of key recovery by exploiting structure in the key-recovery map: pruned Walsh methods eliminate computations that are unnecessary for evaluating the relevant subkey hypotheses [20], whereas Walsh-spectrum puncturing (WSP) reduces the active dimension of the key-recovery problem by approximating the original key-recovery map with a spectrally punctured one [21]. Hou et al. further advanced WSP-based key recovery by characterizing two important families of punctured maps and introducing MILP models to automate the construction of dimension-reduced key-recovery instances [24]. More recently, Beyne et al. revisited the statistical model used in earlier WSP-based key-recovery analyses, identified inaccuracies in some of its underlying assumptions, and derived more accurate parameter estimates, particularly near the full-codebook regime [4]. Overall, WSP trades spectral information for a reduction in the active subkey dimension, thereby potentially reducing the cost of classical key recovery. This mechanism is especially relevant to QFT-based quantum

key recovery, whose complexity is likewise sensitive to the active-subkey dimension. It is therefore natural to ask whether spectral puncturing can reduce this bottleneck while preserving sufficient right-key signal for quantum key recovery.

The main difficulty is that spectral puncturing changes both the key-recovery computation and the correlation statistics on which the quantum attack relies. The existing QFT-based frameworks [34,3] are designed for unpunctured linear key recovery. They do not account for puncturing in correlation state preparation, amplitude amplification, or coherent key discrimination. Moreover, puncturing has two opposing effects. It reduces the active subkey dimension, which can lower the state-preparation cost. At the same time, it discards spectral energy, which weakens the correlation of the correct key. If puncturing is too aggressive, the loss of right-key signal can outweigh the gain from dimension reduction, increasing the overall attack cost (see Appendix H for a quantitative discussion). This leads to the central problem addressed in this work: how to integrate WSP into QFT-based quantum key recovery and optimize the puncturing strategy with respect to the total quantum attack complexity.

Our contributions. In this work, we establish a puncture-aware QFT-based framework for linear and multiple linear quantum key recovery. First, for linear key recovery, we formulate puncture-aware correlation state preparation and derive its complexity under spectral energy loss. We then characterize the outer amplitude-amplification cost under puncturing, showing how active-subkey dimension reduction and right-key signal attenuation jointly determine the quantum key-recovery complexity (Theorem 3). Second, we construct PLinDist for coherent key discrimination over real-valued punctured key-recovery maps and establish its approximation error and gate complexity. Finally, for multiple linear key recovery, we characterize the punctured multiple linear statistic (Theorem 4), construct the corresponding multiple-correlation state, and derive both its preparation and outer key-recovery complexities (Theorem 5).

We further propose Q-Puncture, a MILP-based automated Walsh-spectrum puncturing framework that directly optimizes quantum key-recovery complexity. Q-Puncture evaluates puncturing strategies according to the total key-recovery cost in our quantum cost model,

$$\mathcal{G}_{\text{total}} = R_{\text{outer}}(\mathcal{G}_{\text{NCOR}} + \mathcal{G}_{\text{Dist}}),$$

where $\mathcal{G}_{\text{NCOR}}$ denotes the cost of correlation state preparation, $\mathcal{G}_{\text{Dist}}$ the cost of coherent key discrimination, and R_{outer} the number of outer amplitude amplification iterations. We instantiate three optimization settings: SL-Basic serves as the baseline with a linear approximation and BasicDist; SL-PLin replaces BasicDist with the lower-cost PLinDist, reducing the complexity of coherent key discrimination; and ML-Basic extends the framework to multiple linear approximations while retaining BasicDist. Each setting uses a MILP to characterize structurally feasible puncturing strategies, which are then evaluated according to the quantum cost model above. Q-Puncture can thereby identify a strategy minimizing the modeled total quantum key-recovery cost and determine whether the resulting attack achieves a superquadratic speedup.

As applications, we instantiate these three optimization settings, together with their corresponding key-recovery algorithms, in attacks on Serpent [6,1], RECTANGLE-128 [41], and PIPO-128 [27], as summarized in Table 1. Of particular significance, these results include the first quantum linear key-recovery attacks on Serpent-192 and RECTANGLE-128. In our quantum cost model, the 11-round Serpent-192 and Serpent-256 attacks have complexities 10.05 and 15.43 bits below the corresponding Grover-search baselines, respectively. For RECTANGLE-128, the modeled complexities of the 14-round PLinDist and 16-round BasicDist attacks are 4.87 and 1.13 bits below the Grover-search baseline, respectively. For PIPO-128, we improve the previously analyzed 8-round setting and further extend the analysis to 9 rounds under the typical-statistic model of Theorem 5.

Table 1. Comparison with representative classical and quantum key-recovery attacks.

Target	Attack	Rds.	Time	Data	QRAM	Source
Serpent (192-bit)	Linear	11	$2^{121.23}$ En	$2^{121.23}$	–	[17,31]
	Linear	12	$2^{189.7}$ En	$2^{127.5}$	–	[21,4]
	Linear	12	$2^{184.8}$ En	$2^{127.8}$	–	[24]
	Grover’s search [†]	–	$2^{97.65}Q$	2	–	–
	Q1+BasicDist+WSP	11	$2^{87.6}Q$	2^{128}	2^{128} (QRACM [§])	Section 6
Serpent (256-bit)	Linear	12	$2^{210.4}$ En	$2^{126.3}$	–	[21,4]
	Linear	12	$2^{200.4}$ En	$2^{127.7}$	–	[24]
	Q2+boomerang	9	$2^{115.43}Q$	n.r.	$2^{55.68}$ (QRAQM [¶])	[42]
	Q2+boomerang	10	$2^{126.6}Q$	n.r.	None [#]	[42]
	Grover’s search [†]	–	$2^{129.65}Q$	2	–	–
	Q1+BasicDist+WSP	11	$2^{114.22}Q$	2^{128}	2^{128} (QRACM [§])	Section 6
RECTANGLE-128	Grover’s search [†]	–	$2^{65.65}Q$	2	–	–
	Q1+PLinDist+WSP	14	$2^{60.78}Q$	2^{64}	2^{64} (QRACM [§])	Section 7
	Q1+BasicDist+WSP	16	$2^{64.52}Q$	2^{64}	2^{64} (QRACM [§])	Section 7
PIPO-128	Q1+multiple linear	8	$2^{61.83}Q^{\ddagger}$	2^{64}	2^{71} (QRACM [§])	[34]
	Q1+differential-MITM	8	$2^{61.48}Q$	2^{57}	2^{57} (QRACM [§])	[38]
	Grover’s search [†]	–	$2^{65.65}Q$	2	–	–
	Q1+multiple linear+WSP	8	$2^{60.46}Q$	2^{64}	2^{64} (QRACM [§])	Section 8
	Q1+multiple linear+WSP	9	$2^{65.14}Q$	2^{64}	2^{64} (QRACM [§])	Section 8

[†] Grover’s baseline is round-independent for a fixed key size. [‡] Recomputed following [34] using the same $M = 24$ linear family as our attack. [§] QRACM stores classical data and allows coherent read access using a superposition of addresses. [¶] QRAQM stores quantum data and allows coherent access using quantum addresses. [#] No QRAM is used; the attack assumes Q2 access to the cipher. n.r.: not reported.

Paper layout. Section 2 recalls the necessary background on linear cryptanalysis and quantum algorithms. Sections 3 and 4 develop the punctured quantum linear key-recovery framework for linear and multiple linear approximations. Section 5 presents Q-Puncture, and Sections 6, 7 and 8 apply the framework to Serpent,

RECTANGLE-128 and PIPO-128. Further details of these attacks are provided in the Appendices.

2 Preliminaries

In this section, we first introduce the necessary background on quantum computing and quantum amplitude-amplification and amplitude-estimation techniques, and then provide an overview of classical and quantum linear cryptanalysis.

We briefly recall the Walsh-transform notation used throughout the paper. For a real-valued function $f : \mathbb{F}_2^n \rightarrow \mathbb{R}$, its Walsh transform and inverse transform are defined as

$$\widehat{f}(u) = \sum_{x \in \mathbb{F}_2^n} (-1)^{\langle u, x \rangle} f(x), \quad f(x) = 2^{-n} \sum_{u \in \mathbb{F}_2^n} (-1)^{\langle u, x \rangle} \widehat{f}(u),$$

where $\langle u, x \rangle$ denotes the binary inner product. We use

$$\langle f, g \rangle = 2^{-n} \sum_x f(x)g(x), \quad \|f\|_2^2 = \langle f, f \rangle.$$

Under this convention, Parseval and Plancherel give

$$\|f\|_2^2 = 2^{-2n} \sum_u \widehat{f}(u)^2, \quad \langle f, g \rangle = 2^{-2n} \sum_u \widehat{f}(u)\widehat{g}(u).$$

For $f, g : \mathbb{F}_2^n \rightarrow \mathbb{R}$, define their convolution by

$$(f * g)(z) = \sum_{x \in \mathbb{F}_2^n} f(x)g(x \oplus z).$$

The convolution theorem gives

$$\widehat{f * g}(u) = \widehat{f}(u)\widehat{g}(u).$$

2.1 Quantum Computing Preliminaries

We recall only the quantum query models, memory assumptions, and the amplitude amplification and amplitude estimation procedures used below; standard quantum-computing notation follows [32].

Quantum query models and QRACM. In the Q1 model, the adversary makes only classical queries to the secret-key primitive, whereas in the Q2 model it has coherent access; for an n -bit block cipher E_K ,

$$U_{E_K} |p\rangle |y\rangle = |p\rangle |y \oplus E_K(p)\rangle.$$

We use quantum random-access classical memory (QRACM) for coherent access to classically stored data $D = (d_0, \dots, d_{N-1})$ through

$$U_D |i\rangle |y\rangle = |i\rangle |y \oplus d_i\rangle.$$

QRACM does not imply Q2 access to the cipher. We denote by Q the gate cost of one reversible cipher evaluation under the counting convention specified for each application and normalize the reported gate complexities by Q .

Quantum Amplitude Amplification. Quantum amplitude amplification (QAA) amplifies a good component of probability p using $O(p^{-1/2})$ applications of the preparation procedure and its inverse [12]. For our complexity estimates, we use the following standard variant when only a lower bound on p is known.

Lemma 1 ([36, Lemma 2]). *Let $p \geq p_{\min} > 0$ and $M = \lceil 1.21/\sqrt{p_{\min}} \rceil$. There exists an amplitude-amplification procedure that returns a good state with probability at least $1/2$ and uses, on average, at most $2M + 2$ applications of U or $U^\dagger$.*

Quantum Amplitude Estimation. Quantum amplitude estimation (QAE) estimates the success probability p of a quantum procedure quadratically faster than classical sampling. We use the following standard error bound [12].

Theorem 1 ([12, Theorem 12]). *For $p \in [0, 1]$ and a positive integer q , QAE outputs an estimate $\tilde{p}$ such that, with probability at least $8/\pi^2$,*

$$|\tilde{p} - p| \leq \frac{2\pi\sqrt{p(1-p)}}{q} + \frac{\pi^2}{q^2}. \quad (1)$$

The algorithm uses $O(q)$ applications of the underlying quantum procedure and its inverse.

2.2 Linear Cryptanalysis and Walsh-Spectrum Puncturing

Linear approximation. Let $E_K : \mathbb{F}_2^n \rightarrow \mathbb{F}_2^n$ be an n -bit block cipher. A linear approximation with input mask $\alpha \in \mathbb{F}_2^n$ and output mask $\beta \in \mathbb{F}_2^n$ has correlation

$$\text{cor}_K(\alpha, \beta) = 2^{-n} \sum_{p \in \mathbb{F}_2^n} (-1)^{\langle \alpha, p \rangle \oplus \langle \beta, E_K(p) \rangle}.$$

Its average strength over the key space is measured by the expected linear potential

$$\text{ELP}(\alpha, \beta) = \mathbb{E}_K [\text{cor}_K(\alpha, \beta)^2].$$

Linear key recovery. As illustrated in Fig. 1, the guessed subkey is divided into an external part k_{ext} and an internal part k_{int} . Let $x = x(p, c) \in \mathbb{F}_2^d$ denote the public plaintext–ciphertext variables involved in the approximation. The corresponding key-recovery map can be written in sign representation as

$$\nu(p, c, k_{\text{ext}}, k_{\text{int}}) = f_{k_{\text{int}}}(x(p, c) \oplus k_{\text{ext}}). \quad (2)$$

Given a data set $\mathcal{D}$ of N plaintext–ciphertext pairs, the experimental correlation is

$$\widehat{\text{cor}}_f(k_{\text{ext}}, k_{\text{int}}) = \frac{1}{N} \sum_{(p, c) \in \mathcal{D}} f_{k_{\text{int}}}(x(p, c) \oplus k_{\text{ext}}). \quad (3)$$




Fig. 1. Generic linear key-recovery attack.

Define the distillation table

$$\mathbf{a}[x] = \#\{(p, c) \in \mathcal{D} : x(p, c) = x\}.$$

Then

$$\widehat{\text{cor}}_f(k_{\text{ext}}, k_{\text{int}}) = \frac{1}{N} (\mathbf{a} * f_{k_{\text{int}}})(k_{\text{ext}}). \quad (4)$$

Hence, for each fixed k_{int} , the correlations for all k_{ext} can be evaluated simultaneously in the Walsh domain [17]. This convolution representation also underlies the QFT-based correlation state preparation used later.

Statistical model and Walsh-spectrum puncturing. Let B denote the sampling correction factor. For N distinct plaintexts, $B = \frac{2^n - N}{2^n - 1}$, and hence $B = 0$ in the full-codebook setting $N = 2^n$. Under the standard right- and wrong-key statistical models [19, 8],

$$\widehat{\text{cor}}_f^W \sim \mathcal{N}\left(0, \frac{B}{N} + 2^{-n}\right), \quad (5)$$

$$\widehat{\text{cor}}_f^R \sim \mathcal{N}\left(c, \frac{B}{N} + \text{ELP} - c^2\right), \quad (6)$$

where c is the mean right-key correlation. When dominant trails induce different right-key classes, the unconditional right-key distribution may instead be a mixture of normal components; in such cases, the above model is understood conditionally on a fixed right-key class.

Walsh-spectrum puncturing (WSP) replaces the Boolean key-recovery map $f : \mathbb{F}_2^d \rightarrow \{-1, +1\}$ by a real-valued approximation $g : \mathbb{F}_2^d \rightarrow \mathbb{R}$. For a puncturing set $\mathcal{P} \subseteq \mathbb{F}_2^d \setminus \{0\}$, define

$$g = f - 2^{-d} \sum_{u \in \mathcal{P}} \widehat{f}(u) (-1)^{\langle u, x \rangle}. \quad (7)$$

The puncturing loss and retained spectral energy are

$$\epsilon_{\mathcal{P}} = 2^{-2d} \sum_{u \in \mathcal{P}} \widehat{f}(u)^2, \quad \eta_{\mathcal{P}} = 1 - \epsilon_{\mathcal{P}}. \quad (8)$$

For $\|f\|_2 = 1$, Parseval gives

$$\|g\|_2^2 = \langle f, g \rangle = \eta_{\mathcal{P}}.$$

Under the corrected puncturing model of Beyne et al. [4], the wrong- and right-key correlations of g are approximated by

$$\widehat{\text{cor}}_g^W \sim \mathcal{N}\left(0, \eta_{\mathcal{P}}\left(2^{-n} + \frac{B}{N}\right)\right), \quad (9)$$

$$\widehat{\text{cor}}_g^R \sim \mathcal{N}\left(\eta_{\mathcal{P}}c, \eta_{\mathcal{P}}^2(\text{ELP} - c^2) + \eta_{\mathcal{P}}(1 - \eta_{\mathcal{P}})2^{-n} + \eta_{\mathcal{P}}\frac{B}{N}\right). \quad (10)$$

These distributions are used throughout the subsequent analysis of punctured correlation state preparation and key discrimination.

We consider two WSP strategies [21,24]. *Bit puncturing* removes Walsh coefficients involving selected input coordinates, thereby reducing the active external-key dimension. *LAT-subspace puncturing* restricts the retained spectrum to a low-dimensional affine subspace and can reduce the effective internal-key dimension by identifying key guesses that induce the same retained spectrum.

Multiple linear cryptanalysis. Multiple linear cryptanalysis combines several linear approximations to improve key discrimination [25,7]. For M approximations (α_i, β_i) , define their capacity by

$$C(K) = \sum_{i=1}^M \text{cor}_K(\alpha_i, \beta_i)^2, \quad C = \mathbb{E}_K[C(K)] = \sum_{i=1}^M \text{ELP}_i. \quad (11)$$

The corresponding key-recovery statistic is

$$\hat{q}(z) = \sum_{i=1}^M \widehat{\text{cor}}_i(z)^2. \quad (12)$$

2.3 Quantum Linear Cryptanalysis

QFT-based correlation state preparation was introduced for quantum linear key recovery by Schrottenloher [34], with a related quantum convolution construction independently studied in [14]. We use the quantum convolution algorithm (QCA) formulation of [33], restated below under our Walsh-transform and normalized ℓ_2 -norm conventions.

Let $\{f_z\}_{z \in \mathcal{Z}}$ and $\{g_z\}_{z \in \mathcal{Z}}$ be families of real-valued functions over $\mathbb{F}_2^d$, and define

$$\mathcal{N}_f = \sum_{z \in \mathcal{Z}} \|f_z\|_2^2, \quad \mathcal{N}_{f*g} = \sum_{z \in \mathcal{Z}} \|f_z * g_z\|_2^2.$$

The QCA prepares the convolution state

$$|\text{Conv}(f, g)\rangle = \frac{1}{\sqrt{2^d \mathcal{N}_{f*g}}} \sum_{z \in \mathcal{Z}} \sum_{x \in \mathbb{F}_2^d} (f_z * g_z)(x) |x\rangle |z\rangle. \quad (13)$$

The construction assumes an initialization operator

$$\text{Init} |0\rangle = \frac{1}{\sqrt{2^d \mathcal{N}_f}} \sum_{z \in \mathcal{Z}} \sum_{x \in \mathbb{F}_2^d} f_z(x) |x\rangle |z\rangle, \quad (14)$$

and a coherent implementation of the Walsh coefficients of g_z ,

$$g \text{ Fourier} : |u\rangle |z\rangle |0\rangle \mapsto |u\rangle |z\rangle \left(\frac{\widehat{g}_z(u)}{G} |0\rangle + |\perp_{u,z}\rangle \right), \quad (15)$$

where

$$G \geq \max_{z,u} |\widehat{g}_z(u)|$$

and the ancillary component of $|\perp_{u,z}\rangle$ is orthogonal to the all-zero state.

Theorem 2 (Quantum convolution algorithm [33]). *Given implementations of Init and $g \text{ Fourier}$, there exists a quantum algorithm that returns a state close to the convolution state in Eq. (13) with gate complexity*

$$\mathcal{G}(\text{QCA}) \leq \frac{\pi G \sqrt{\mathcal{N}_f}}{2 \sqrt{\mathcal{N}_{f*g}}} (\mathcal{G}(\text{Init}) + \mathcal{G}(g \text{ Fourier})), \quad (16)$$

up to polynomially bounded lower-order costs.

The algorithm applies Walsh transforms to encode the convolution in the amplitudes and uses amplitude amplification to obtain the normalized state. A detailed description of the QCA is given in Appendix A.

3 Puncture-Aware Quantum Linear Key Recovery

Building on the punctured key-recovery model introduced in Section 2.2, we develop a puncture-aware QFT-based quantum key-recovery framework for a linear approximation.

3.1 Correlation State Preparation under Puncturing

For a fixed internal-key guess z_{in} , the experimental correlation of the punctured map g can be written as

$$\widehat{\text{cor}}_g(z_{\text{in}}, z) = \frac{1}{N} (\mathbf{a} * g_{z_{\text{in}}})(z), \quad (17)$$

where $d = m = |k_{\text{ext}}|$ and $z \in \mathbb{F}_2^m$. The corresponding normalized correlation state is

$$|\text{Cor}_{z_{\text{in}}}^g\rangle = \frac{\sum_{z \in \mathbb{F}_2^m} \widehat{\text{cor}}_g(z_{\text{in}}, z) |z\rangle}{\sqrt{\sum_{z \in \mathbb{F}_2^m} \widehat{\text{cor}}_g(z_{\text{in}}, z)^2}}. \quad (18)$$

Algorithm 1 Subroutine of NCORCOMP

Input: $ z_{\text{in}}\rangle 0^d\rangle$	
1: Apply $\text{Init}_{\mathbf{a}}$	$\triangleright z_{\text{in}}\rangle \frac{1}{\ \mathbf{a}\ _{\ell_2}} \sum_{x \in \mathbb{F}_2^d} \mathbf{a}[x] x\rangle$
2: Apply $H^{\otimes d}$	$\triangleright z_{\text{in}}\rangle \frac{1}{\ \mathbf{a}\ _{\ell_2} 2^{d/2}} \sum_{u \in \mathbb{F}_2^d} \widehat{\mathbf{a}}[u] u\rangle$
3: Append $ 0^{r_a}\rangle$	$\triangleright z_{\text{in}}\rangle \frac{1}{\ \mathbf{a}\ _{\ell_2} 2^{d/2}} \sum_u \widehat{\mathbf{a}}[u] u\rangle 0^{r_a}\rangle$
4: Apply g Fourier	$\triangleright z_{\text{in}}\rangle \frac{1}{\ \mathbf{a}\ _{\ell_2} 2^{d/2}} \sum_u \widehat{\mathbf{a}}[u] u\rangle \left(\frac{\widehat{g}_{z_{\text{in}}}(u)}{G} 0^{r_a}\rangle + \perp_{z_{\text{in}},u}\rangle \right)$
5: Collect the all-zero ancilla component	$\triangleright z_{\text{in}}\rangle \frac{1}{G \ \mathbf{a}\ _{\ell_2} 2^{d/2}} \sum_u \widehat{\mathbf{a}}[u] \widehat{g}_{z_{\text{in}}}(u) u\rangle 0^{r_a}\rangle + \perp\rangle$
6: Apply $H^{\otimes d}$	$\triangleright z_{\text{in}}\rangle \frac{1}{G \ \mathbf{a}\ _{\ell_2}} \sum_{x \in \mathbb{F}_2^d} (\mathbf{a} * g_{z_{\text{in}}})(x) x\rangle 0^{r_a}\rangle + \perp\rangle$
7: Rewrite using the experimental correlation	$\triangleright z_{\text{in}}\rangle \frac{N}{G \ \mathbf{a}\ _{\ell_2}} \sum_{x \in \mathbb{F}_2^d} \widehat{\text{Cor}}_g(z_{\text{in}}, x) x\rangle 0^{r_a}\rangle + \perp\rangle$

To prepare this state, we assume coherent implementations of

$$\text{Init}_{\mathbf{a}} : |0^d\rangle \mapsto \frac{1}{\|\mathbf{a}\|_{\ell_2}} \sum_{x \in \mathbb{F}_2^d} \mathbf{a}[x] |x\rangle, \quad (19)$$

and

$$g \text{ Fourier} : |z_{\text{in}}\rangle |u\rangle |0^{r_a}\rangle \mapsto |z_{\text{in}}\rangle |u\rangle \left(\frac{\widehat{g}_{z_{\text{in}}}(u)}{G} |0^{r_a}\rangle + |\perp_{z_{\text{in}},u}\rangle \right), \quad (20)$$

where $G \geq \max_{z_{\text{in}},u} |\widehat{g}_{z_{\text{in}}}(u)|$ and the ancillary component of $|\perp_{z_{\text{in}},u}\rangle$ is orthogonal to $|0^{r_a}\rangle$. The register of width r_a is used only for amplitude encoding. The distillation table $\mathbf{a}$ is constructed classically from the available plaintext-ciphertext data and stored in QRACM. We make the explicit state-preparation assumption that the coherent operator $\text{Init}_{\mathbf{a}}$ in Eq. (19) is available. This assumption concerns amplitude-state preparation rather than a QRACM lookup alone.

We retain $G(\text{Init}_a)$ in the formal state-preparation complexity. In the concrete instantiations below, we use the working assumption $\mathcal{G}(\text{Init}_a) < \mathcal{G}(g \text{ Fourier})$ and therefore omit $\mathcal{G}(\text{Init}_a)$ from the leading-order numerical estimates.

For the complexity analysis, we assume $d = m$ and the full-codebook setting $N = 2^n$, so that $B = 0$. We use the typical bucket-count model for $\mathbf{a}$ and independent wrong-key correlations. On the correct internal-key branch, we further assume that the correlation norm is dominated by the wrong-key contribution.

Lemma 2. *There exists a quantum algorithm NCORCOMP such that, on the correct internal-key branch,*

$$\text{NCORCOMP} |k_{\text{int}}\rangle |0\rangle \simeq |k_{\text{int}}\rangle |\text{Cor}_{k_{\text{int}}}^g\rangle,$$

up to the negligible preparation error described below. For $N = 2^n$ and $d = m$, its estimated Toffoli-gate count is

$$\left(\frac{\pi}{2} \frac{G}{\sqrt{\eta_{\mathcal{P}}}} \sqrt{2^{-m} + 2^{n-2m}} + 3 \right) (\mathcal{G}(\text{Init}_{\mathbf{a}}) + \mathcal{G}(g \text{ Fourier})), \quad (21)$$

up to polynomially bounded lower-order costs.

Proof. Algorithm 1 defines one preparation attempt. On the correct internal-key branch, its all-zero ancilla component has probability

$$p_R = \frac{N^2}{G^2 \|\mathbf{a}\|_{\ell_2}^2} \sum_{z \in \mathbb{F}_2^m} \widehat{\text{cor}}_g(k_{\text{int}}, z)^2, \quad (22)$$

and, conditioned on this component, the output state is $|\text{Cor}_{k_{\text{int}}}^g\rangle$.

By the assumed bucket-count model,

$$\mathbb{E}[\|\mathbf{a}\|_{\ell_2}^2] = N + N(N-1)2^{-m},$$

and hence, for $N = 2^n$,

$$\|\mathbf{a}\|_{\ell_2}^2 \simeq 2^n + 2^{2n-m}.$$

For the correlation norm, the $2^m - 1$ wrong external-key guesses follow the punctured wrong-key model in Eq. (9). Under the independence assumption,

$$\frac{1}{\eta_{\mathcal{P}} 2^{-n}} \sum_{z \neq k_{\text{ext}}} \widehat{\text{cor}}_g(k_{\text{int}}, z)^2 \approx \chi_{2^m-1}^2.$$

Therefore, this sum is sharply concentrated around $\eta_{\mathcal{P}}(2^m - 1)2^{-n}$. Under the wrong-key-dominated norm assumption, the contribution of the right-key term is lower order, and thus

$$\sum_{z \in \mathbb{F}_2^m} \widehat{\text{cor}}_g(k_{\text{int}}, z)^2 \simeq \eta_{\mathcal{P}} 2^{m-n}.$$

Substituting these estimates into Eq. (22) gives the typical preparation probability

$$p_0 \simeq \frac{\eta_{\mathcal{P}}}{G^2(2^{-m} + 2^{n-2m})}.$$

The concentration of the correlation norm implies that the actual preparation probability differs from p_0 only by a negligible relative amount. Consequently, amplitude amplification calibrated to p_0 incurs only a negligible preparation error and uses at most $\frac{\pi}{2\sqrt{p_0}} + 3$ calls to the preparation circuit or its inverse. $\square$

Unlike the construction in [34], NCORCOMP allows $m = |k_{\text{ext}}|$ to exceed the block size n . Although this may increase the correlation state preparation cost, Walsh-spectrum puncturing can reduce the active key dimension and thereby lower the overall key-recovery complexity.

3.2 Puncture-Aware Outer Key Recovery

Following Bashiri et al. [3], we use a coherent distinguisher to mark the correct active subkey $(k_{\text{int}}, k_{\text{ext}})$ directly. This differs from the original QFT-based

attack of Schrottenloher [34] in which the active subkey is extended to a full master-key candidate before verification. Since the analysis of [3] relies on the unpunctured right- and wrong-key distributions, we re-evaluate the outer key-recovery complexity under the corrected puncturing model of [4], where

$$v_W = \eta_P \left(2^{-n} + \frac{B}{N} \right), \quad v_R = \eta_P^2 (\text{ELP} - c^2) + \eta_P (1 - \eta_P) 2^{-n} + \eta_P \frac{B}{N}.$$

For the correct key after puncturing, $\widehat{\text{cor}}_g(k_{\text{int}}, k_{\text{ext}})$ satisfies, with probability approximately 0.95,

$$\begin{aligned} |\widehat{\text{cor}}_g(k_{\text{int}}, k_{\text{ext}}) - \eta_P c| &\leq 2\sqrt{v_R} \\ \Rightarrow \quad \eta_P c - 2\sqrt{v_R} &\leq |\widehat{\text{cor}}_g(k_{\text{int}}, k_{\text{ext}})| \leq \eta_P c + 2\sqrt{v_R}. \end{aligned} \quad (23)$$

For the single-dominant-trail applications considered below, we specialize to a fixed right-key class and write $c > 0$ for the magnitude of its mean correlation. When the dominant trail accounts for the full linear potential, this gives $c^2 = \text{ELP}$. The general punctured right-key model above keeps c and ELP as separate parameters. In the concrete single-trail instantiations below, we use the single-dominant-trail approximation $\text{ELP} \simeq c^2$.

The state in Eq. (18) is always normalized with its exact denominator. Under the punctured wrong-key randomization hypothesis, each wrong-key correlation has variance $v_W = \eta_P(2^{-n} + B/N)$. Therefore, summing the squared correlations over the 2^m external-key guesses gives

$$\sum_{z \in \mathbb{F}_2^m} \widehat{\text{cor}}_g(z_{\text{in}}, z)^2 \simeq 2^m \left(\eta_P \left(2^{-n} + \frac{B}{N} \right) \right) = 2^m v_W. \quad (24)$$

This is only a typical-case estimate for the correlation norm; it is not used as an exact normalization when the correct key is present.

We first analyze the outer search with an ideal active-subkey marking unitary Dist_{eff} . We work in the full-codebook setting and apply NCORCOMP to a uniform superposition of internal-key guesses. We further assume that the right-key event in Eq. (23) holds and that the correlation norm is dominated by the wrong-key contribution.

Theorem 3. *Let $c = \sqrt{\text{ELP}}$ and assume $\sqrt{\eta_P} c > 2\sqrt{1 - \eta_P} 2^{-n/2}$. Let Dist_{eff} be the ideal marking unitary defined by*

$$\text{Dist}_{\text{eff}} |z_{\text{in}}, z\rangle |b\rangle = \begin{cases} |z_{\text{in}}, z\rangle |b \oplus 1\rangle, & (z_{\text{in}}, z) = (k_{\text{int}}, k_{\text{ext}}), \\ |z_{\text{in}}, z\rangle |b\rangle, & \text{otherwise.} \end{cases}$$

Then the outer QAA returns $(k_{\text{int}}, k_{\text{ext}})$ with probability at least 0.5, using an estimated average number of

$$2 \left\lceil 1.21 \frac{2^{|k_{\text{int}}|/2} 2^{(m-n)/2}}{\sqrt{\eta_P} c - 2\sqrt{1 - \eta_P} 2^{-n/2}} \right\rceil + 2 \quad (25)$$

calls to NCORCOMP and Dist_{eff} .

Proof. Prepare a uniform superposition over the internal-key register and apply NCORCOMP. By Lemma 2, on the correct internal-key branch the external-key register is, up to the negligible preparation error, in the state $|\text{Cor}_{k_{\text{int}}}^g\rangle$. Therefore, neglecting this preparation error, the amplitude of the unique basis state marked by Dist_{eff} is

$$a_R = 2^{-|k_{\text{int}}|/2} \frac{\widehat{\text{cor}}_g^R}{\sqrt{\sum_{z \in \mathbb{F}_2^m} \widehat{\text{cor}}_g(k_{\text{int}}, z)^2}}.$$

In the full-codebook setting, $v_W = \eta_P 2^{-n}$. Using the wrong-key-dominated normalization estimate in Eq. (24),

$$\sum_{z \in \mathbb{F}_2^m} \widehat{\text{cor}}_g(k_{\text{int}}, z)^2 \simeq 2^m \eta_P 2^{-n},$$

and hence

$$\frac{1}{|a_R|} \simeq \frac{2^{|k_{\text{int}}|/2} \sqrt{\eta_P} 2^{(m-n)/2}}{|\widehat{\text{cor}}_g^R|}.$$

On the right-key event in Eq. (23),

$$|\widehat{\text{cor}}_g^R| \geq \eta_P c - 2\sqrt{\eta_P(1 - \eta_P)} 2^{-n/2}.$$

Therefore,

$$\frac{1}{|a_R|} \lesssim \frac{2^{|k_{\text{int}}|/2} 2^{(m-n)/2}}{\sqrt{\eta_P} c - 2\sqrt{1 - \eta_P} 2^{-n/2}}.$$

Substituting this modeled inverse marked amplitude into the amplitude-amplification call-count expression gives Eq. (25). $\square$

The marking operation Dist_{eff} can be instantiated by different coherent distinguishers. We distinguish three representative strategies: the key-completion strategy of Schrottenloher [34], the Simon-based and linear distinguishers of Bashiri et al. [3]. In our terminology, these correspond to BasicDist, SimonDist, and LinDist, respectively. Appendix B summarizes BasicDist and LinDist and briefly recalls SimonDist for context. Under puncturing, we extend LinDist to PLinDist as described below. We focus on BasicDist and PLinDist in this work.

3.3 PLinDist for Punctured Key-Recovery Maps

LinDist introduced by Bashiri et al. [3] uses quantum amplitude estimation, coherent repetition, and majority voting to distinguish subkey candidates from their experimental correlations. However, it is designed for Boolean key-recovery maps. After puncturing, the map is generally real-valued, so the experimental correlation becomes a weighted average rather than a Boolean count. We therefore introduce PLinDist to coherently estimate this real-valued statistic.

As in [3], we consider either the Q2 model or the Q1 model with the full codebook stored in QRACM, and let Q denote the cost of one coherent block-cipher evaluation. Let g be the punctured key-recovery map defined in Section 2.2, with retained spectral energy η_P , and write $\bar{z} = (z_{\text{int}}, z_{\text{ext}})$.

For $x \in \mathbb{F}_2^n$, let $h_{\bar{z}}(x) = h(x, E_K(x), \bar{z})$ denote the input to g , and let $f_0(x, E_K(x)) \in \{-1, 1\}$ denote the subkey-independent part of the approximation. In the full-codebook setting, define

$$s_{\bar{z}}(x) = f_0(x, E_K(x))g(h_{\bar{z}}(x)), \quad C_g(\bar{z}) = 2^{-n} \sum_{x \in \mathbb{F}_2^n} s_{\bar{z}}(x). \quad (26)$$

To estimate $C_g(\bar{z})$ coherently, let B_g be a known upper bound on $|s_{\bar{z}}(x)|$ and set

$$a_{\bar{z}}(x) = \frac{1}{2} \left(1 + \frac{s_{\bar{z}}(x)}{B_g} \right) \in [0, 1].$$

Let $\mathcal{A}_g$ satisfy

$$\mathcal{A}_g : |\bar{z}\rangle|0^n\rangle|0\rangle \mapsto |\bar{z}\rangle \frac{1}{2^{n/2}} \sum_{x \in \mathbb{F}_2^n} |x\rangle (\sqrt{1 - a_{\bar{z}}(x)}|0\rangle + \sqrt{a_{\bar{z}}(x)}|1\rangle). \quad (27)$$

The probability of obtaining $|1\rangle$ is

$$p_g(\bar{z}) = 2^{-n} \sum_{x \in \mathbb{F}_2^n} a_{\bar{z}}(x) = \frac{1}{2} \left(1 + \frac{C_g(\bar{z})}{B_g} \right), \quad (28)$$

and therefore

$$C_g(\bar{z}) = B_g(2p_g(\bar{z}) - 1). \quad (29)$$

Thus, QAE can coherently estimate the punctured correlation through $p_g(\bar{z})$. Algorithm 2 gives the resulting threshold distinguisher.

Algorithm 2 PLinDist

- 1: **for** $i = 1, \dots, L$ **do**
 - 2: Coherently run QAE $_q$ using $\mathcal{A}_g$ to estimate $p_g(\bar{z})$; denote the output by $\tilde{p}_i$
 - 3: Compute $\tilde{C}_i \leftarrow B_g(2\tilde{p}_i - 1)$
 - 4: **end for**
 - 5: Compute $\text{cnt} \leftarrow |\{i : |\tilde{C}_i| \geq T_g\}|$
 - 6: **if** $\text{cnt} \geq (L + 1)/2$ **then**
 - 7: $b \leftarrow b \oplus 1$
 - 8: **end if**
 - 9: Uncompute cnt , all $\tilde{C}_i$, and all QAE registers
-

We write the gate cost of one implementation of $\mathcal{A}_g$ as $\mathcal{G}(\mathcal{A}_g) = Q + \mathcal{G}_{\text{enc}}$, where $\mathcal{G}_{\text{enc}}$ includes the reversible evaluation and uncomputation of the key-recovery map, the fixed-point arithmetic needed to obtain $a_{\bar{z}}(x)$, and the controlled rotation used for amplitude encoding. For the concrete instances below,

we assume $\mathcal{G}_{\text{enc}} \ll Q$ and therefore use $\mathcal{G}(\mathcal{A}_g) \simeq Q$ in the leading-order numerical estimates.

Lemma 3. *Let $\kappa_g = |k_{\text{int}}| + |k_{\text{ext}}|$. In the full-codebook setting, put $c = \sqrt{\text{ELP}} > 0$ and assume $\sqrt{\eta_P} c \geq 4\sqrt{2\kappa_g} 2^{-n/2}$. Let L be a positive odd integer and let $T_g = \frac{3}{8}\eta_P\sqrt{\text{ELP}}, q = \left\lceil 32\pi B_g \eta_P^{-1} \text{ELP}^{-\frac{1}{2}} \right\rceil$. Conditioned on the statistical separation event of Lemma 6 and on independent QAE repetitions, PLinDist approximates Dist_{eff} with error $\|\text{PLinDist} - \text{Dist}_{\text{eff}}\| \leq 2^{\kappa_g/2 - 0.045L + 2}$. Its gate complexity is approximately $8qLQ$.*

Proof. The proof is given in Appendix D. $\square$

If the outer key-recovery procedure invokes Dist_{eff} at most R_{outer} times, replacing each ideal invocation by PLinDist accumulates an error of at most $R_{\text{outer}} \|\text{PLinDist} - \text{Dist}_{\text{eff}}\|$. It therefore suffices to choose

$$L = \left\lceil \frac{1}{0.045} \left(\frac{\kappa_g}{2} + \log_2 R_{\text{outer}} + 2n + 2 \right) \right\rceil$$

and round it up to the next odd integer. With this choice, the accumulated error is at most 2^{-2n} . Up to additive constants, this has the same form as the original LinDist.

Compared with LinDist, puncturing changes the estimation factor from $\frac{1}{\sqrt{\text{ELP}}}$ to $\frac{B_g}{\eta_P \sqrt{\text{ELP}}}$, while reducing the active-subkey dimension κ_g in the repetition count L .

4 Puncture-Aware Quantum Key Recovery with Multiple Linear Approximations

Section 3 considers a linear approximation. We now extend the framework to multiple linear approximations. While the unpunctured multiple linear QFT setting was studied by Schrottenloher [34], Walsh-spectrum puncturing changes both the multiple linear statistic and the resulting quantum key-recovery complexity. We therefore characterize the punctured statistic, construct the corresponding multiple-correlation state, and derive its state-preparation and outer key-recovery complexities. Throughout this section, we assume full-codebook access, $N = 2^n$.

For each approximation i and internal-key guess z_{in} , let $g_{i,z_{\text{in}}}$ denote the corresponding punctured key-recovery map. Following the standard multiple-linear statistical model used in prior QFT-based analyses [34], we model the M linear approximations as statistically independent. They share a common internal-key register, external-key dimension m , and retained spectral energy η_P , and their capacity is $C = \sum_{i=1}^M \text{ELP}_i$. The independence assumption is statistical and is not implied by the structural construction of the approximation family.

4.1 Punctured Multiple Linear Key-Recovery Statistic

For a subkey guess (z_{in}, z) , define the multiple linear statistic

$$\hat{q}_g(z_{\text{in}}, z) = \sum_{i=1}^M \widehat{\text{cor}}_{g_i}(z_{\text{in}}, z)^2. \quad (30)$$

Within a fixed right-key class, let

$$v_i = \eta_{\mathcal{P}}^2(\text{ELP}_i - c_i^2) + \eta_{\mathcal{P}}(1 - \eta_{\mathcal{P}})2^{-n}$$

denote the variance of the i th right-key correlation in the full-codebook setting.

Theorem 4. *Under the independent Gaussian correlation model with $N = 2^n$, the wrong-key statistic satisfies*

$$\hat{q}_g^W \approx \eta_{\mathcal{P}} 2^{-n} \chi_M^2. \quad (31)$$

The right-key statistic has mean and variance

$$S_R := \mathbb{E}[\hat{q}_g^R] = \eta_{\mathcal{P}}^2 C + M \eta_{\mathcal{P}} (1 - \eta_{\mathcal{P}}) 2^{-n}, \quad (32)$$

$$\text{Var}(\hat{q}_g^R) = \sum_{i=1}^M (2v_i^2 + 4\eta_{\mathcal{P}}^2 c_i^2 v_i). \quad (33)$$

For sufficiently large M , if $c_i = 0$ and the v_i are approximately equal, then $\hat{q}_g^R \approx \mathcal{N}(S_R, 2S_R^2/M)$. For dominant trails, these expressions apply within each key class; the unconditional law is a mixture over classes.

Proof. The wrong-key result follows by squaring and summing M independent $\mathcal{N}(0, \eta_{\mathcal{P}} 2^{-n})$ variables. For the right key, write the i -th correlation as an independent normal variable with mean $\eta_{\mathcal{P}} c_i$ and variance v_i . The identities $\mathbb{E}[X^2] = v + \mu^2$ and $\text{Var}(X^2) = 2v^2 + 4\mu^2 v$ for $X \sim \mathcal{N}(\mu, v)$ give Eqs. (32) and (33). The last approximation is the equal-variance normal approximation to the resulting scaled χ^2 statistic. $\square$

Under the above normal approximation,

$$\hat{q}_g^R = S_R(1 + \delta), \quad |\delta| \leq 3\sqrt{\frac{2}{M}}$$

at the three-standard-deviation level.

4.2 Multiple Correlation State Preparation under Puncturing

For a fixed internal-key guess z_{in} , define

$$\begin{aligned} S_{z_{\text{in}}}^{(M)} &= \sum_{i=1}^M \sum_{z \in \mathbb{F}_2^n} \widehat{\text{cor}}_{g_i}(z_{\text{in}}, z)^2, \\ |\text{MCor}_{z_{\text{in}}}^g\rangle &= \frac{1}{\sqrt{S_{z_{\text{in}}}^{(M)}}} \sum_{i=1}^M \sum_{z \in \mathbb{F}_2^n} \widehat{\text{cor}}_{g_i}(z_{\text{in}}, z) |i\rangle |z\rangle. \end{aligned} \quad (34)$$

As in Section 3.1, assume a coherent Walsh-coefficient circuit

$$g \text{ FMULT} : |z_{\text{in}}, i\rangle |x\rangle |0^{r_a}\rangle \mapsto |z_{\text{in}}, i\rangle |x\rangle \left(\frac{\widehat{g}_{i, z_{\text{in}}}(x)}{G} |0^{r_a}\rangle + |\perp_{z_{\text{in}}, i, x}\rangle \right), \quad (35)$$

where $G \geq \max_{z_{\text{in}}, i, x} |\widehat{g}_{i, z_{\text{in}}}(x)|$.

Lemma 4. NCORMULT prepares $|\text{MCor}_{k_{\text{int}}}^g\rangle$ on the correct internal-key branch up to the same negligible preparation error as in Lemma 2. For $N = 2^n$ and $d = m$, its estimated Toffoli-gate count is

$$\left(\frac{\pi G}{2\sqrt{\eta_P}} \sqrt{2^{-m} + 2^{n-2m}} + 3 \right) (\mathcal{G}(\text{Init}_{\mathbf{a}}) + \mathcal{G}(g \text{ FMULT})).$$

Proof. Initialize the approximation register to $M^{-1/2} \sum_{i=1}^M |i\rangle$ and apply the same QCA construction as in Lemma 2, with $g \text{ FMULT}$ replacing $g \text{ Fourier}$. Under the same norm approximation, $\sum_{i,z} \widehat{\text{cor}}_{g_i}(k_{\text{int}}, z)^2 \simeq M\eta_P 2^{m-n}$. The factor M cancels the $1/M$ factor arising from the uniform approximation register, so the preparation probability and amplitude-amplification factor are the same as in Lemma 2, giving the stated complexity. $\square$

4.3 Outer Key Recovery from Multiple-Correlation States

We reuse the ideal marking operation Dist_{eff} from Section 3.2, acting trivially on the approximation-index register. As in the linear case, we apply NCORMULT to a uniform superposition of internal-key guesses and assume that the correlation norm is dominated by the wrong-key contribution.

Theorem 5. Under the above assumptions, using the typical right-key value $\widehat{q}_g^R \simeq S_R$ and the wrong-key-dominated normalization, the outer-QAA cost for returning $(k_{\text{int}}, k_{\text{ext}})$ is estimated by

$$2 \left\lceil 1.21 2^{(|k_{\text{int}}|+m)/2} \left(1 - \eta_P + \eta_P \frac{2^n C}{M} \right)^{-1/2} \right\rceil + 2 \quad (36)$$

calls to NCORMULT and Dist_{eff} .

Proof. Prepare a uniform superposition over the internal-key register and apply NCORMULT. Neglecting the preparation error of Lemma 4, the probability of the subspace marked by Dist_{eff} is

$$p = 2^{-|k_{\text{int}}|} \frac{\widehat{q}_g^R}{\sum_{z \in \mathbb{F}_2^m} \widehat{q}_g(k_{\text{int}}, z)}.$$

Under the wrong-key-dominated normalization assumption and the full-codebook model,

$$\sum_{z \in \mathbb{F}_2^m} \widehat{q}_g(k_{\text{int}}, z) \simeq 2^m M \eta_P 2^{-n}.$$

Using the typical right-key value $\hat{q}_g^R \simeq S_R$ from Eq. (32), where

$$S_R = \eta_{\mathcal{P}}^2 C + M \eta_{\mathcal{P}} (1 - \eta_{\mathcal{P}}) 2^{-n},$$

gives the typical marked probability

$$p_{\text{typ}} \simeq 2^{-(|k_{\text{int}}|+m)} \left(1 - \eta_{\mathcal{P}} + \eta_{\mathcal{P}} \frac{2^n C}{M} \right).$$

Substituting p_{typ} into the outer amplitude-amplification call-count expression gives Eq. (36). $\square$

Thus, puncturing affects the outer-QAA cost through $(1 - \eta_{\mathcal{P}} + \eta_{\mathcal{P}} 2^n C/M)^{-1/2}$ and the multiple-correlation state preparation cost through $G/\sqrt{\eta_{\mathcal{P}}}$, while simultaneously reducing the active-subkey dimensions. These quantities form the tradeoff optimized by Q-Puncture in the next section.

5 Q-Puncture: Automated Walsh-Spectrum Puncturing for Quantum Key Recovery

Sections 3 and 4 quantify how spectral puncturing affects quantum key recovery through the retained spectral energy, the active-subkey dimensions, and the spectral normalization factor. We now turn these quantities into an optimization criterion for selecting puncturing strategies.

Q-Puncture starts from the automated WSP model of Hou et al. [24], whose bit-puncturing and affine-subspace constraints characterize the admissible retained Walsh supports; these structural constraints are recalled in Appendix C. For each feasible support, we derive the retained spectral energy, the active-subkey dimensions, and the spectral normalization parameter, and map them to the quantum costs established in Sections 3 and 4. Thus, the MILP of Hou et al. serves as the structural search layer of Q-Puncture, whereas the feasibility conditions and optimization criterion are determined by the puncture-aware quantum key-recovery analysis established above. For a candidate puncturing strategy, we evaluate

$$\mathcal{G}_{\text{total}} = R_{\text{outer}} (\mathcal{G}_{\text{NCOR}} + \mathcal{G}_{\text{Dist}}), \quad (37)$$

where $\mathcal{G}_{\text{NCOR}}$ is the correlation state preparation cost, $\mathcal{G}_{\text{Dist}}$ is the coherent distinguishing cost, and R_{outer} is the outer amplitude-amplification factor.

5.1 Puncture-Aware Spectral Normalization

We establish the spectral normalization bound used to evaluate a punctured key-recovery map. For one side of the map, let $\eta_{\mathcal{P}} = 2^{-d_{\mathcal{P}}}$, $l = |k_{\text{ext}}|$, $r = |k_{\text{int}}|$, $s = l - r$, where $0 \leq r \leq l$. Here, $d_{\mathcal{P}}$ is the spectral-loss exponent and need not be an integer. Under the puncture-aware random-sign model, fixing an r -dimensional intermediate configuration leaves an s -dimensional local Fourier sum whose typical scale is reduced by $2^{-d_{\mathcal{P}}/2}$.

Lemma 5. *Let $G_{X,M}$ denote the spectral normalization parameter for one side X of the key-recovery map when $M \geq 1$ linear approximations are evaluated simultaneously. Suppose that every punctured local Fourier contribution $A_{i,u}(\xi)$ satisfies*

$$\Pr\left[|A_{i,u}(\xi)| \geq t 2^{(s-d_P)/2}\right] \leq 2e^{-t^2/6}.$$

Then, with probability at least 0.99,

$$G_{X,M} \leq 2^{(l+r-d_P)/2} H(l, M), \quad (38)$$

where

$$H(l, M) = \sqrt{6(\ln 100 + (l+1)\ln 2 + \ln M)}.$$

Proof. There are $M2^r2^s = M2^l$ triples (i, u, ξ) . By the assumed tail bound and a union bound, the choice

$$H(l, M) = \sqrt{6(\ln 100 + (l+1)\ln 2 + \ln M)}$$

ensures, with probability at least 0.99, that

$$|A_{i,u}(\xi)| \leq 2^{(s-d_P)/2} H(l, M)$$

simultaneously for all (i, u, ξ) . For every complete Fourier coefficient,

$$\widehat{g}_{i,z}(\omega) = \sum_{u \in \mathbb{F}_2^r} \chi_{i,z,\omega}(u) A_{i,u}(\xi_{u,\omega}), \quad |\chi_{i,z,\omega}(u)| = 1.$$

Hence,

$$|\widehat{g}_{i,z}(\omega)| \leq 2^r 2^{(s-d_P)/2} H(l, M) = 2^{(l+r-d_P)/2} H(l, M),$$

which proves the claim. $\square$

In the concrete automated search, Eq. (38) is used as a model-based high-probability spectral bound under the above random-sign hypothesis; whenever a cipher-specific value of G is available, we use it directly.

The automated model below applies puncturing independently to both sides of the key-recovery map and combines their side-specific spectral bounds multiplicatively. If one side contains only one recovery round, or if its spectral normalization parameter can otherwise be computed exactly, that side need not be punctured: its exact value is substituted directly. Thus, bilateral puncturing is the general case, while the one-sided model is an immediate specialization.

5.2 From Spectral Puncturing to Quantum Cost

We now translate the spectral bound of Lemma 5 into the correlation state preparation costs derived in Sections 3 and 4.

For each side $X \in \{L, R\}$, let d_X , r_X , and l_X denote the spectral-loss exponent and the internal- and external-key dimensions, respectively. We write

$d_{\mathcal{P}} = d_L + d_R$, $\eta_{\mathcal{P}} = 2^{-d_{\mathcal{P}}}$, and, when the two sides use disjoint active-subkey coordinates, $|k_{\text{int}}| = r_L + r_R$, $|k_{\text{ext}}| = l_L + l_R$.

Puncturing may be applied independently to the two sides. If one side contains only one recovery round, or if its spectral normalization parameter can otherwise be computed exactly, its exact value is used directly and we set $d_X = 0$ for that side. Thus, bilateral puncturing is the general case, while one-sided puncturing is obtained as an immediate specialization.

Let $h(l, M) = \log_2 H(l, M)$. For $X \in \{L, R\}$, define

$$\gamma_X(M) = \begin{cases} \frac{l_X + r_X}{2} + h(l_X, M), & \text{if Lemma 5 is used,} \\ g_X, & \text{if } G_{X,M} = 2^{g_X} \text{ is known exactly.} \end{cases}$$

Set

$$\Gamma_M = \gamma_L(M) + \gamma_R(M), \quad G_M = G_{L,M} G_{R,M}.$$

Applying Lemma 5 to both sides and using a union bound, the two spectral bounds hold simultaneously with probability at least 0.98, and hence

$$G_M \leq 2^{\Gamma_M - d_{\mathcal{P}}/2}. \quad (39)$$

Since $\eta_{\mathcal{P}} = 2^{-d_{\mathcal{P}}}$, Eq. (39) gives

$$\frac{G_M}{\sqrt{\eta_{\mathcal{P}}}} \leq 2^{\Gamma_M}. \quad (40)$$

Thus, the factor $2^{-d_{\mathcal{P}}/2}$ gained in the spectral normalization bound cancels the factor $1/\sqrt{\eta_{\mathcal{P}}}$ caused by spectral-energy loss in the state-preparation costs of Lemmas 2 and 4. The remaining effect of puncturing on state preparation is therefore captured by Γ_M together with the reduced active-subkey dimensions.

Let $\mathcal{G}(E) = 2^a$, and let 2^b denote the gate cost of the relevant coherent Walsh-coefficient circuit, namely $\mathcal{G}(g \text{ Fourier})$ in the linear setting and $\mathcal{G}(g \text{ FMULT})$ in the multiple linear setting. Ignoring ceilings and polynomially bounded lower-order terms, and omitting the cost of $\text{Init}_{\mathbf{a}}$ as in Section 3.1, the dominant correlation state preparation cost satisfies

$$\mathcal{G}_{\text{NCOR}}^{(M)} \leq \frac{\pi}{2} 2^{\Gamma_M} \sqrt{2^{-|k_{\text{ext}}|} + 2^{n-2|k_{\text{ext}}|}} 2^b. \quad (41)$$

For convenience, define

$$\delta_n(t) = \frac{1}{2} \log_2(1 + 2^{t-n}). \quad (42)$$

Then

$$\log_2 \sqrt{2^{-t} + 2^{n-2t}} = \frac{n}{2} - t + \delta_n(t).$$

This form will be used below to derive the feasibility constraints for the different Q-Puncture settings.

5.3 Q-Puncture Optimization Settings

Together with the WSP support constraints and cipher-specific propagation rules, we instantiate Q-Puncture in three settings: SL-Basic, SL-PLin, and ML-Basic. They share the same structural search model, but differ in the number of linear approximations, the coherent distinguisher, and the corresponding quantum key-recovery cost.

SL-Basic. This setting uses a linear approximation together with BasicDist. With $r(E) = 4$ as in Appendix B, the cost of BasicDist is $\mathcal{G}_{\text{Basic}} = \pi 2^{a+(K-|k_{\text{int}}|-|k_{\text{ext}}|)/2}$. Requiring the correlation state preparation cost to remain below the distinguishing cost, $\mathcal{G}_{\text{NCOR}}^{(1)} < \mathcal{G}_{\text{Basic}}$, and substituting Eq. (41) gives

$$\frac{|k_{\text{int}}| - |k_{\text{ext}}|}{2} + \Gamma_1 + \delta_n(|k_{\text{ext}}|) < a - b + \frac{K - n}{2} + 1. \quad (43)$$

SL-PLin. This setting uses a linear approximation together with the puncture-aware distinguisher PLinDist. Let $B_g \geq \max_{\bar{z}, x} |s_{\bar{z}}(x)|$. Using Lemma 3, we estimate $q \simeq \frac{32\pi B_g}{\eta_{\mathcal{P}} \sqrt{\text{ELP}}}$, $\mathcal{G}_{\text{PLin}} \simeq 8qL 2^a$. We impose $\mathcal{G}_{\text{NCOR}}^{(1)} < \mathcal{G}_{\text{PLin}} < \mathcal{G}_{\text{Basic}}$, which yields

$$\begin{cases} \frac{n}{2} - |k_{\text{ext}}| + \Gamma_1 - d_{\mathcal{P}} + \delta_n(|k_{\text{ext}}|) < a - b + 9 + \log_2 \left(\frac{LB_g}{\sqrt{\text{ELP}}} \right), \\ d_{\mathcal{P}} + 9 + \log_2 \left(\frac{LB_g}{\sqrt{\text{ELP}}} \right) < \frac{K - |k_{\text{int}}| - |k_{\text{ext}}|}{2} + 1. \end{cases} \quad (44)$$

ML-Basic. This setting uses M linear approximations together with BasicDist. For common side profiles (d_X, r_X, l_X) , the condition $\mathcal{G}_{\text{NCOR}}^{(M)} < \mathcal{G}_{\text{Basic}}$ gives

$$\frac{|k_{\text{int}}| - |k_{\text{ext}}|}{2} + \Gamma_M + \delta_n(|k_{\text{ext}}|) < a - b + \frac{K - n}{2} + 1. \quad (45)$$

The inequalities above are feasibility constraints rather than optimization objectives. Subject to these constraints, the WSP support constraints, and the cipher-specific propagation rules, Q-Puncture minimizes Eq. (37). Here, $\mathcal{G}_{\text{Dist}} = \mathcal{G}_{\text{Basic}}$ for SL-Basic and ML-Basic, and $\mathcal{G}_{\text{Dist}} = \mathcal{G}_{\text{PLin}}$ for SL-PLin. The corresponding R_{outer} follows from Sections 3 and 4.

This objective captures the tradeoff underlying puncturing: increasing $d = -\log_2 \eta_{\mathcal{P}}$ reduces the active Walsh dimension and hence the state-preparation cost, but also weakens the right-key signal and increases the outer amplification cost. Consequently, the total complexity is generally non-monotonic in the puncturing level and is typically minimized at an intermediate value of d (see Appendix H).

6 Quantum Attacks on Serpent

Serpent [6,1] is a 128-bit block cipher and one of the finalists of the AES competition. It adopts a 32-round SPN structure, supports 128/192/256-bit keys,

and uses eight 4-bit S-boxes $S_0, \dots, S_7$ together with a linear transformation. A detailed specification is given in Appendix E.1.

6.1 Attack Setting

We follow the linear key-recovery setting considered in [21,24]. As illustrated in Fig. 2, we use the first eight rounds of the 9-round characteristic reported in [16]. The resulting trail covers R3–R10, follows the S-box sequence $S_3, S_4, S_5, S_6, S_7, S_0, S_1, S_2$, and has correlation 2^{-46} . Adding one plaintext-side and two ciphertext-side key-recovery rounds gives an 11-round attack over R2–R12 with structure $1 + 8 + 2$.




Fig. 2. Key recovery attack against 11-round Serpent.

On the plaintext side, R2 contains 15 active S_2 boxes. Four local Walsh supports have affine dimension 2 and the remaining eleven have dimension 4, giving an active external-key dimension of 52 and $G_L = 2^{45}$.

On the ciphertext side, the trail output activates seven S_3 boxes in R11. Before puncturing, their local supports give $r_0 = 28$, while propagation through LT^{-T} activates 28 nibbles in R12 and gives $l_0 = 112$. Puncturing is applied only to these two ciphertext-side recovery rounds. Hence, for a ciphertext-side profile $(d_{\mathcal{P}}, r, l)$, $|k_{\text{int}}| = r$, $|k_{\text{ext}}| = 52 + l$, $G = G_L G_R$.

Since the attack uses a linear approximation together with BasicDist, we instantiate the SL-Basic setting of Section 5.3. We count Toffoli gates and normalize quantum time by $Q = \mathcal{G}(E_{11}) = 3008 \simeq 2^{11.55}$. Both attacks below are evaluated in the Q1 model. We assume that the full codebook has been collected classically and stored in QRACM, which enables coherent access to the plaintext–ciphertext data during the quantum key-recovery phase. The detailed puncturing constraints and automated search procedure are given in Appendix E.3.

6.2 Quantum Key Recovery for 11-Round Serpent-256

For $K = 256$, the automated search returns $(d_{\mathcal{P}}, r, l) = (2.09, 24, 104)$, $\eta_{\mathcal{P}} = \frac{15}{64}$. Thus, $|k_{\text{int}}| = 24$, $|k_{\text{ext}}| = 156$. For this profile, $\mathcal{G}(g \text{ Fourier}) = 6648 \simeq 2^{12.70}$. The SL-Basic feasibility condition in Eq. (43) is satisfied:

$$\underbrace{24 + 4.43 + 14.00}_{r+h(104,1)+\delta_{128}(156)} = 42.43 < \underbrace{11.55 - 12.70 + 64 + 1 - 45 + 26}_{44.86}.$$

By Lemma 5, the spectral normalization factor is bounded by

$$\begin{aligned} G &\leq G_L 2^{(l+r-d_{\mathcal{P}})/2} H(l, 1) \\ &= 2^{45} 2^{(104+24-\log_2(64/15))/2} \sqrt{6(\ln 100 + 105 \ln 2)} \\ &\simeq 2^{45} 2^{62.95} 2^{4.43} \lesssim 2^{112.38}. \end{aligned}$$

The first factor below is the outer amplification cost, while the two terms in the second factor correspond to NCORCOMP and BasicDist, respectively. Normalizing by $Q = \mathcal{G}(E_{11}) \simeq 2^{11.55}$, the total key-recovery cost is

$$\begin{aligned} \mathcal{G}_{\text{total}} &\simeq \left(\frac{2(1.21)2^{24/2}2^{(156-128)/2}}{\sqrt{15/64}2^{-46} - 2\sqrt{49/64}2^{-64}} + 2 \right) \\ &\quad \times \left(\frac{\pi 2^{112.38}}{2\sqrt{15/64}} \sqrt{2^{-156} + 2^{128-2\cdot 156}} 2^{12.70-11.55} + \pi 2^{(256-24-156)/2} \right) Q \\ &\simeq 2^{74.32} (2^{37.23} + 2^{39.65}) Q \simeq 2^{114.22} Q. \end{aligned}$$

The retained Walsh supports and their propagation to the active R12 S-boxes are given in Appendix E.4.

6.3 Quantum Key Recovery for 11-Round Serpent-192

For $K = 192$, the automated search returns $(d_{\mathcal{P}}, r, l) = (12, 4, 84)$ with $\eta_{\mathcal{P}} = 2^{-12}$, giving $|k_{\text{int}}| = 4$ and $|k_{\text{ext}}| = 136$. For this puncturing strategy, $\mathcal{G}(g \text{ Fourier}) = 5894 \simeq 2^{12.53}$, the SL-Basic feasibility condition in Eq. (43) is satisfied, and Lemma 5 gives $G \lesssim 2^{87.29}$. Substituting these parameters into the punctured state-preparation cost of Lemma 2, the outer key-recovery cost of Theorem 3, and the BasicDist cost gives a total key-recovery complexity of $\mathcal{G}_{\text{total}} \simeq 2^{87.60} Q$, which is 10.05 bits below the Grover-search baseline. The retained Walsh supports and the corresponding R12 activity pattern are given in Appendix E.4.

7 Quantum Attacks on RECTANGLE-128

RECTANGLE [41] is a 64-bit lightweight block cipher that supports 80- and 128-bit keys and has been extensively analyzed using classical cryptanalytic techniques. In this section, we focus on RECTANGLE-128 and apply the quantum linear key-recovery framework developed above. A detailed description of RECTANGLE is given in Appendix F.

7.1 Attack Setting

We consider two quantum linear key-recovery attacks on RECTANGLE-128. The first combines an 11-round linear trail with one plaintext-side and two ciphertext-side recovery rounds, yielding a $1 + 11 + 2$ construction. Since it uses one linear approximation together with PLinDist, it is instantiated in the **SL-PLin** setting. The second extends the central trail to 12 rounds and adds two recovery rounds on both sides, yielding a $2 + 12 + 2$ construction in the **SL-Basic** setting with BasicDist.

As shown in Fig. 3, the two trails share the same input mask. The 11-round trail terminates with correlation 2^{-25} , while extending it by one round gives the 12-round trail with correlation 2^{-28} . The corresponding parameters are summarized in Table 2, and the detailed active-subkey derivations are given in Appendix F.

Both attacks are evaluated in the Q1 model, with the full codebook collected classically and stored in QRACM.

7.2 14-Round Attack in the SL-PLin Setting

We first consider the $1 + 11 + 2$ construction based on the 11-round linear trail with correlation $c = 2^{-25}$. Before puncturing, the ciphertext-side recovery has $(r, l) = (6, 40)$, and the active-subkey dimensions are $|k_{\text{int}}| = 6, |k_{\text{ext}}| = 46$. The automated search punctures two of the three trail-adjacent ciphertext-side components and returns $d_P = 4, \eta_P = 2^{-4}, r = 2, l = 24$, so that $|k_{\text{int}}| = 2, |k_{\text{ext}}| = 30$.

For this profile, $\kappa_g = 32$, and the signal-separation condition required by PLinDist holds with equality: $\sqrt{\eta_P} c = 2^{-27} = 4\sqrt{2\kappa_g} 2^{-32}$. The retained local




Fig. 3. The 11- and 12-round linear trails used in the RECTANGLE-128 attacks.

Table 2. Parameters of the RECTANGLE-128 attacks.

	14-round SL-PLin	16-round SL-Basic
Attack structure	1 + 11 + 2	2 + 12 + 2
Correlation	2^{-25}	2^{-28}
Input mask	0000000000060002	0000000000060002
Output mask	0600004000002000	4000820000420000
G_L	2^9	—
$\mathcal{G}(E)$	$2^{10.81}$	2^{11}
$\mathcal{G}(g \text{ Fourier})$	$2^{11.07}$	$2^{11.95}$

spectra give $B_g = 2^{-2}$, and hence $q \simeq 2^{33.65}$, $L \simeq 2^{11.76}$. The selected supports and the calculation of B_g are given in Appendix F.5 and Appendix F.4, respectively.

By Lemma 5, the spectral normalization factor is bounded by

$$\begin{aligned}
G &\leq G_L 2^{(l+r-d_P)/2} H(l, 1) \\
&= 2^9 2^{(24+2-4)/2} \sqrt{6(\ln 100 + 25 \ln 2)} \\
&\simeq 2^9 2^{11} 2^{3.52} \lesssim 2^{23.52}.
\end{aligned}$$

For this profile,

$$\mathcal{G}_{\text{NCOR}} \simeq 2^{28.43} Q < \mathcal{G}_{\text{PLin}} \simeq 2^{48.41} Q < \mathcal{G}_{\text{Basic}} \simeq 2^{49.65} Q,$$

so the SL-PLin feasibility requirements in Eq. (44) are satisfied.

The first factor below is the outer amplification cost, while the two terms in the second factor correspond to NCORCOMP and PLinDist, respectively. Normalizing by $Q = \mathcal{G}(E_{14}) \simeq 2^{10.81}$, the total key-recovery cost is

$$\begin{aligned}
\mathcal{G}_{\text{total}} &\simeq \left(\frac{2(1.21)2^{2/2}2^{(30-64)/2}}{\sqrt{2^{-4}} 2^{-25} - 2\sqrt{1-2^{-4}} 2^{-32}} + 2 \right) \\
&\quad \times \left(\frac{\pi 2^{23.52}}{2\sqrt{2^{-4}}} \sqrt{2^{-30} + 2^{64-2\cdot 30}} 2^{11.07-10.81} + 8 \cdot 2^{33.65} \cdot 2^{11.76} \right) Q \\
&\simeq 2^{12.37} (2^{28.43} + 2^{48.41}) Q \simeq 2^{60.78} Q.
\end{aligned}$$

The resulting total cost is $\mathcal{G}_{\text{total}} \simeq 2^{60.78} Q$, about 4.87 bits below the Grover-search baseline of $2^{65.65} Q$ for RECTANGLE-128.

7.3 16-Round Attack in the SL-Basic Setting

We next use the 12-round trail with correlation 2^{-28} and add two recovery rounds on each side, yielding a $2+12+2$ construction. Before puncturing, the two recovery sides have $(r_L, l_L) = (8, 36)$ and $(r_R, l_R) = (12, 48)$. The joint SL-Basic optimization punctures only the plaintext side and returns $(d_L, r_L, l_L) = (2, 4, 28)$

and $(d_R, r_R, l_R) = (0, 12, 48)$. Hence, $d_P = 2$, $\eta_P = 2^{-2}$, $|k_{\text{int}}| = 16$, and $|k_{\text{ext}}| = 76$. Lemma 5 gives $G \lesssim 2^{52.57}$, and both the SL-Basic feasibility condition and the positive-signal condition of Theorem 3 are satisfied. Substituting these parameters into the modeled SL-Basic cost gives $\mathcal{G}_{\text{total}} \simeq 2^{64.52}Q$, which is about 1.13 bits below Grover’s search. The detailed numerical evaluation and selected Walsh supports are given in Appendix F.

8 Quantum Attacks on PIPO-128

PIPO-128 is a 13-round, 64-bit block cipher with a 128-bit master key [27]. We analyze it in the ML-Basic setting. Following [34], we start from the 6000 boundary masks that activate at most two PIPO S-boxes on either side of the relevant linear layer. Instead of the previously used 112-approximation family, our search selects $M = 24$ five-round linear hulls, consisting of three rotational orbits (20, 23, and 24) with eight cyclic shifts each. Although structurally related, these approximations are evaluated under the independent-correlation model of Theorem 4, following prior multiple-linear QFT analyses [34]; the resulting complexities should therefore be interpreted as model-based estimates under this assumption. Each approximation has $\text{ELP} = 2^{-50}$, so $C = 24 \cdot 2^{-50}$ and $C/M = 2^{-50}$.

For this family, Q-Puncture minimizes the modeled ML-Basic cost in Eq. (37), with the outer amplification factor from Theorem 5; no key-schedule relations are exploited. For 8 rounds, puncturing is applied only to the plaintext-side recovery, yielding $\eta_P = 5/8$ and $(d_P, r, l) = (\log_2(8/5), 12, 64)$; for 9 rounds, both recovery sides are punctured jointly. Table 3 summarizes the selected profiles; further details are given in Appendix G.

For the 8-round profile, $\eta_P = 5/8$, $|k_{\text{int}}| = 12$, and $|k_{\text{ext}}| = 64$. Since $n = 64$ and $C/M = 2^{-50}$, Theorem 5 gives the typical outer amplification factor

$$R_{\text{outer}} \simeq 2 \left\lceil 1.21 \cdot 2^{38} \left(\frac{3}{8} + \frac{5}{8} 2^{14} \right)^{-1/2} \right\rceil + 2 \simeq 2^{32.61}.$$

Since $|k_{\text{int}}| + |k_{\text{ext}}| = 76$, BasicDist searches the remaining 52 master-key bits at cost $\pi 2^{26}Q \simeq 2^{27.65}Q$. Combining these costs according to Eq. (37) gives $\mathcal{G}_{\text{total}} \simeq 2^{60.46}Q$, about 5.19 bits below the Grover-search baseline. For nine rounds, joint puncturing gives $(d_P, r, l) = (9.83, 9, 80)$ and the corresponding typical-statistic estimate is $2^{65.14}Q$, only 0.51 bits below the Grover baseline; we therefore regard the nine-round result as a boundary-case estimate.

9 Conclusion

In this work, building on Walsh-spectrum puncturing, we develop a puncture-aware framework for QFT-based quantum linear key recovery. We establish a

Table 3. ML-Basic attacks on PIPO-128. Complexity exponents are reported as $\log_2(\mathcal{G}/Q)$, with the factor Q omitted.

Rounds	Puncturing profile	$\mathcal{G}_{\text{total}}$	Grover	Gain
8	$(d_{\mathcal{P}}, r, l) = (\log_2(8/5) \simeq 0.68, 12, 64)$	60.46	65.65	5.19
9	$(d_{\mathcal{P}}, r, l) = (\log_2(8192/9) \simeq 9.83, 9, 80)$	65.14	65.65	0.51

quantitative connection between spectral puncturing and quantum attack complexity by characterizing how the resulting dimension reduction and spectral-energy loss jointly affect correlation state preparation, right-key signal strength, and outer amplitude amplification. The resulting framework covers both linear and multiple linear key recovery, and we further develop PLinDist to support coherent key discrimination for real-valued punctured key-recovery maps. Based on these results, we propose Q-Puncture, an automated MILP-based optimization framework that searches over puncturing strategies and minimizes the modeled total quantum key-recovery cost by balancing active-subkey-dimension reduction against spectral information loss. A key advantage of our approach is that puncturing strategies are evaluated according to a modeled quantum key-recovery cost that jointly accounts for correlation-state preparation, coherent key discrimination, and outer amplitude amplification. The applications to Serpent, RECTANGLE-128 and PIPO-128 demonstrate that Walsh-spectrum puncturing can effectively improve quantum linear cryptanalysis when the structure of the key-recovery map is properly exploited. Several directions are of interest for future work:

- **Tighter spectral bounds.** Our automated model relies on an upper bound for the spectral normalization parameter G . Improving this bound, or deriving cipher-specific estimates with stronger guarantees, may enlarge the set of feasible puncturing strategies and lead to better quantum attacks.
- **Beyond correlation-limited attacks.** Current QFT-based quantum linear key-recovery attacks, including ours, are restricted by the correlation strength of the underlying linear approximations. Investigating whether additional quantum techniques can relax this correlation limitation and enable effective attacks with weaker correlations remains an important open problem.
- **Relationship with quantum distinguishers.** We mainly use BasicDist and PLinDist. Combining puncturing with other quantum distinguishers, such as SimonDist, and jointly optimizing the puncturing strategy and distinguisher remain interesting directions.

References

1. Anderson, R., Biham, E., Knudsen, L.: Serpent: A proposal for the advanced encryption standard. First Advanced Encryption Standard Candidate Conference (1998)

2. Bao, Z., Guo, J., Li, S., Pham, P.: Automatic quantum multi-collision distinguishers and rebound attacks with triangulation algorithm. In: Zhu, T., Li, Y. (eds.) ACISP 24: 29th Australasian Conference on Information Security and Privacy, Part II. Lecture Notes in Computer Science, vol. 14896, pp. 24–43. Springer, Singapore, Singapore, Sydney, NSW, Australia (Jul 15–17, 2024). https://doi.org/10.1007/978-981-97-5028-3_2
3. Bashiri, K., Bonnetain, X., Hosoyamada, A., Lang, N., Schrottenloher, A.: Improved quantum linear attacks and application to CAST. IACR Transactions on Symmetric Cryptology **2025**(2), 124–165 (2025). <https://doi.org/10.46586/tosc.v2025.i2.124-165>
4. Beyne, T., Flórez-Gutiérrez, A., Todo, Y.: Accurate parameter estimates for punctured key recovery linear attacks. IACR Communications in Cryptology **3**(2) (2026). <https://doi.org/10.62056/ava3z11w9p>
5. Biham, E.: On Matsui’s linear cryptanalysis. In: De Santis, A. (ed.) Advances in Cryptology – EUROCRYPT’94. Lecture Notes in Computer Science, vol. 950, pp. 341–355. Springer Berlin Heidelberg, Germany, Perugia, Italy (May 9–12, 1995). <https://doi.org/10.1007/BFb0053449>
6. Biham, E., Anderson, R.J., Knudsen, L.R.: Serpent: A new block cipher proposal. In: Vaudenay, S. (ed.) Fast Software Encryption – FSE’98. Lecture Notes in Computer Science, vol. 1372, pp. 222–238. Springer Berlin Heidelberg, Germany, Paris, France (Mar 23–25, 1998). https://doi.org/10.1007/3-540-69710-1_15
7. Biryukov, A., De Cannière, C., Quisquater, M.: On multiple linear approximations. In: Franklin, M. (ed.) Advances in Cryptology – CRYPTO 2004. Lecture Notes in Computer Science, vol. 3152, pp. 1–22. Springer Berlin Heidelberg, Germany, Santa Barbara, CA, USA (Aug 15–19, 2004). https://doi.org/10.1007/978-3-540-28628-8_1
8. Bogdanov, A., Rijmen, V.: Linear hulls with correlation zero and linear cryptanalysis of block ciphers. Designs, Codes and Cryptography **70**(3), 369–383 (2014). <https://doi.org/10.1007/s10623-012-9697-z>
9. Bonnetain, X., Naya-Plasencia, M.: Hidden shift quantum cryptanalysis and implications. In: Peyrin, T., Galbraith, S. (eds.) Advances in Cryptology – ASIACRYPT 2018, Part I. Lecture Notes in Computer Science, vol. 11272, pp. 560–592. Springer, Cham, Switzerland, Brisbane, Queensland, Australia (Dec 2–6, 2018). https://doi.org/10.1007/978-3-030-03326-2_19
10. Bonnetain, X., Naya-Plasencia, M., Schrottenloher, A.: On quantum slide attacks. In: Paterson, K.G., Stebila, D. (eds.) SAC 2019: 26th Annual International Workshop on Selected Areas in Cryptography. Lecture Notes in Computer Science, vol. 11959, pp. 492–519. Springer, Cham, Switzerland, Waterloo, ON, Canada (Aug 12–16, 2019). https://doi.org/10.1007/978-3-030-38471-5_20
11. 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
12. Brassard, G., Høyer, P., Mosca, M., Tapp, A.: Quantum amplitude amplification and estimation. Contemporary Mathematics **305**, 53–74 (2002). <https://doi.org/10.1090/conm/305/05215>
13. 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
14. Castellazo, G., Nguyen, Q.T., Palma, G.D., Englund, D., Lloyd, S., Kiani, B.T.: Quantum algorithms for group convolution, cross-correlation, and equivariant transformations. *Physical Review A* **106**(3), 032402 (2022). <https://doi.org/10.1103/PhysRevA.106.032402>
 15. Cho, J.Y., Hermelin, M., Nyberg, K.: A new technique for multidimensional linear cryptanalysis with applications on reduced round Serpent. In: Lee, P.J., Cheon, J.H. (eds.) ICISC 08: 11th International Conference on Information Security and Cryptology. Lecture Notes in Computer Science, vol. 5461, pp. 383–398. Springer Berlin Heidelberg, Germany, Seoul, Korea (Dec 3–5, 2009). https://doi.org/10.1007/978-3-642-00730-9_24
 16. Collard, B., Standaert, F.X., Quisquater, J.J.: Improved and multiple linear cryptanalysis of reduced round serpent. In: Pei, D., Yung, M., Lin, D., Wu, C. (eds.) Information Security and Cryptology – Inscrypt 2007. Lecture Notes in Computer Science, vol. 4990, pp. 51–65. Springer (2007). https://doi.org/10.1007/978-3-540-79499-8_6
 17. Collard, B., Standaert, F.X., Quisquater, J.J.: Improving the time complexity of Matsui’s linear cryptanalysis. In: Nam, K.H., Rhee, G. (eds.) ICISC 07: 10th International Conference on Information Security and Cryptology. Lecture Notes in Computer Science, vol. 4817, pp. 77–88. Springer Berlin Heidelberg, Germany, Seoul, Korea (Nov 29–30, 2007). https://doi.org/10.1007/978-3-540-76788-6_7
 18. Collard, B., Standaert, F.X., Quisquater, J.J.: Experiments on the multiple linear cryptanalysis of reduced round Serpent. In: Nyberg, K. (ed.) Fast Software Encryption – FSE 2008. Lecture Notes in Computer Science, vol. 5086, pp. 382–397. Springer Berlin Heidelberg, Germany, Lausanne, Switzerland (Feb 10–13, 2008). https://doi.org/10.1007/978-3-540-71039-4_24
 19. Daemen, J., Rijmen, V.: Probability distributions of correlation and differentials in block ciphers. *Journal of Mathematical Cryptology* **1**(3), 221–242 (2007). <https://doi.org/10.1515/JMC.2007.011>
 20. Flórez-Gutiérrez, A.: Optimising linear key recovery attacks with affine Walsh transform pruning. In: Agrawal, S., Lin, D. (eds.) Advances in Cryptology – ASIACRYPT 2022, Part IV. Lecture Notes in Computer Science, vol. 13794, pp. 447–476. Springer, Cham, Switzerland, Taipei, Taiwan (Dec 5–9, 2022). https://doi.org/10.1007/978-3-031-22972-5_16
 21. Flórez-Gutiérrez, A., Todo, Y.: Improving linear key recovery attacks using Walsh spectrum puncturing. In: Joye, M., Leander, G. (eds.) Advances in Cryptology – EUROCRYPT 2024, Part I. Lecture Notes in Computer Science, vol. 14651, pp. 187–216. Springer, Cham, Switzerland, Zurich, Switzerland (May 26–30, 2024). https://doi.org/10.1007/978-3-031-58716-0_7
 22. Grover, L.K.: A fast quantum mechanical algorithm for database search. In: 28th Annual ACM Symposium on Theory of Computing. pp. 212–219. ACM Press, Philadelphia, PA, USA (May 22–24, 1996). <https://doi.org/10.1145/237814.237866>
 23. Hermelin, M., Cho, J.Y., Nyberg, K.: Multidimensional linear cryptanalysis of reduced round Serpent. In: Mu, Y., Susilo, W., Seberry, J. (eds.) ACISP 08: 13th Australasian Conference on Information Security and Privacy. Lecture Notes in Computer Science, vol. 5107, pp. 203–215. Springer Berlin Heidelberg,

- Germany, Wollongong, Australia (Jul 7–9, 2008). https://doi.org/10.1007/978-3-540-70500-0_15
24. Hou, C., Wang, S., Liu, M.: Walsh spectrum puncturing revisited: Toward automated linear key recovery attacks. *IACR Transactions on Symmetric Cryptology* **2026**(1), 318–344 (2026). <https://doi.org/10.46586/tosc.v2026.i1.318-344>
 25. Kaliski, Jr., B.S., Robshaw, M.J.B.: Linear cryptanalysis using multiple approximations. In: Desmedt, Y. (ed.) *Advances in Cryptology – CRYPTO’94*. Lecture Notes in Computer Science, vol. 839, pp. 26–39. Springer Berlin Heidelberg, Germany, Santa Barbara, CA, USA (Aug 21–25, 1994). https://doi.org/10.1007/3-540-48658-5_4
 26. 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>
 27. Kim, H., Jeon, Y., Kim, G., Kim, J., Sim, B.Y., Han, D.G., Seo, H., Kim, S., Hong, S., Sung, J., Hong, D.: PIPO: A lightweight block cipher with efficient higher-order masking software implementations. In: Hong, D. (ed.) *ICISC 20: 23rd International Conference on Information Security and Cryptology*. Lecture Notes in Computer Science, vol. 12593, pp. 99–122. Springer, Cham, Switzerland, Seoul, Korea (Dec 2–4, 2020). https://doi.org/10.1007/978-3-030-68890-5_6
 28. Lee, J.K., Kim, W.H.: Multiple linear cryptanalysis using linear statistics. *IACR Transactions on Symmetric Cryptology* **2019**(4), 369–406 (2019). <https://doi.org/10.13154/tosc.v2019.i4.369-406>
 29. Liu, F., Zhu, X., Xu, R., Shi, D., Wang, P.: The offline quantum attack against modular addition variant of Even-Mansour cipher. In: Zhu, T., Li, Y. (eds.) *ACISP 24: 29th Australasian Conference on Information Security and Privacy, Part I*. Lecture Notes in Computer Science, vol. 14895, pp. 3–19. Springer, Singapore, Singapore, Sydney, NSW, Australia (Jul 15–17, 2024). https://doi.org/10.1007/978-981-97-5025-2_1
 30. Matsui, M.: Linear cryptanalysis method for DES cipher. In: Hellese, T. (ed.) *Advances in Cryptology – EUROCRYPT’93*. Lecture Notes in Computer Science, vol. 765, pp. 386–397. Springer Berlin Heidelberg, Germany, Lofthus, Norway (May 23–27, 1994). https://doi.org/10.1007/3-540-48285-7_33
 31. McLaughlin, J., Clark, J.A.: Filtered nonlinear cryptanalysis of reduced-round Serpent, and the wrong-key randomization hypothesis. In: Stam, M. (ed.) *14th IMA International Conference on Cryptography and Coding*. Lecture Notes in Computer Science, vol. 8308, pp. 120–140. Springer Berlin Heidelberg, Germany, Oxford, UK (Dec 17–19, 2013). https://doi.org/10.1007/978-3-642-45239-0_8
 32. Nielsen, M.A., Chuang, I.L.: *Quantum Computation and Quantum Information*. Cambridge University Press (2010)
 33. Pichollet-Mugnier, A., Schrottenloher, A.: Quantum truncated differential attacks using convolutions. *IACR Transactions on Symmetric Cryptology* **2026**(1), 148–188 (2026). <https://doi.org/10.46586/tosc.v2026.i1.148-188>
 34. Schrottenloher, A.: Quantum linear key-recovery attacks using the QFT. In: Handschuh, H., Lysyanskaya, A. (eds.) *Advances in Cryptology – CRYPTO 2023, Part V*. Lecture Notes in Computer Science, vol. 14085, pp. 258–291. Springer, Cham, Switzerland, Santa Barbara, CA, USA (Aug 20–24, 2023). https://doi.org/10.1007/978-3-031-38554-4_9
 35. Schrottenloher, A., Stevens, M.: Simplified MITM modeling for permutations: New (quantum) attacks. In: Dodis, Y., Shrimpton, T. (eds.) *Advances in Cryptology –*

- CRYPTO 2022, Part III. Lecture Notes in Computer Science, vol. 13509, pp. 717–747. Springer, Cham, Switzerland, Santa Barbara, CA, USA (Aug 15–18, 2022). https://doi.org/10.1007/978-3-031-15982-4_24
36. Schrottenloher, A., Stevens, M.: Quantum procedures for nested search problems: with applications in cryptanalysis. *IACR Communications in Cryptology (CiC)* **1**(3), 9 (2024). <https://doi.org/10.62056/aee0fhhbm0>
 37. Simon, D.R.: On the power of quantum computation. In: 35th Annual Symposium on Foundations of Computer Science. pp. 116–123. IEEE Computer Society Press, Santa Fe, NM, USA (Nov 20–22, 1994). <https://doi.org/10.1109/SFCS.1994.365701>
 38. Xu, Y.S., Cai, B.B., Yuan, Z., Qin, S.J., Gao, F., Wen, Q.Y.: Quantum differential meet-in-the-middle attack and some applications to lightweight ciphers. *Advanced Quantum Technologies* **7**(10), 2400157 (2024). <https://doi.org/10.1002/qute.202400157>
 39. Xu, Z., Chen, S., Wang, M., Wei, P.: Linear cryptanalysis and its variants with fast fourier transformation technique on MPC/FHE/ZK-friendly $\mathbb{F}_p$ -based ciphers. In: Simpson, L., Baee, M.A.R. (eds.) *ACISP 23: 28th Australasian Conference on Information Security and Privacy*. Lecture Notes in Computer Science, vol. 13915, pp. 25–52. Springer, Cham, Switzerland, Brisbane, QLD, Australia (Jul 5–7, 2023). https://doi.org/10.1007/978-3-031-35486-1_2
 40. Zajac, P., Jókay, M.: Multiplicative complexity of bijective 4x4 s-boxes. *Cryptography and Communications* **6**(3), 255–277 (2014). <https://doi.org/10.1007/s12095-014-0100-y>
 41. Zhang, W., Bao, Z., Lin, D., Rijmen, V., Yang, B., Verbauwhede, I.: RECTANGLE: A bit-slice lightweight block cipher suitable for multiple platforms. *Science China Information Sciences* **58** (2015). <https://doi.org/10.1007/s11432-015-5459-7>
 42. Zou, H., Zou, J., Luo, Y.: New results on quantum boomerang attacks. *Quantum Information Processing* **22**(4), 171 (2023). <https://doi.org/10.1007/s11128-023-03921-6>

A Quantum Convolution Algorithm

For completeness, we give the detailed realization of the quantum convolution algorithm recalled in Section 2.3.

B Coherent Key Distinguishers

The operator Dist_{eff} used in the outer amplitude-amplification procedure is the ideal specification of a coherent key distinguisher. For an active subkey candidate $\bar{z} = (z_{\text{int}}, z_{\text{ext}})$, it implements

$$\text{Dist}_{\text{eff}} |\bar{z}\rangle |b\rangle = |\bar{z}\rangle |b \oplus 1[\bar{z} = \bar{k}]\rangle,$$

where $\bar{k} = (k_{\text{int}}, k_{\text{ext}})$ denotes the correct active subkey. This equation is a specification, not an oracle that can compare a candidate with the unknown key for free. A concrete attack must supply an implementable realization whose correctness is proved on a statistical event. Different quantum key-recovery strategies

Algorithm 3 Quantum convolution algorithm [33]

1: Starting from the all-zero state, apply Init:

$$\frac{1}{\sqrt{2^d \mathcal{N}_f}} \sum_{z \in \mathcal{Z}} \sum_{x \in \mathbb{F}_2^d} f_z(x) |x\rangle |z\rangle.$$

2: Apply $H^{\otimes d}$ to the $|x\rangle$ register:

$$\frac{1}{2^d \sqrt{\mathcal{N}_f}} \sum_{z \in \mathcal{Z}} \sum_{u \in \mathbb{F}_2^d} \hat{f}_z(u) |u\rangle |z\rangle.$$

3: Append an ancillary register and apply g Fourier:

$$\begin{aligned} & \frac{1}{2^d \sqrt{\mathcal{N}_f}} \sum_{z \in \mathcal{Z}} \sum_{u \in \mathbb{F}_2^d} \hat{f}_z(u) |u\rangle |z\rangle \left(\frac{\hat{g}_z(u)}{G} |0\rangle + |\perp_{u,z}\rangle \right) \\ &= \frac{1}{2^d G \sqrt{\mathcal{N}_f}} \left(\sum_{z \in \mathcal{Z}} \sum_{u \in \mathbb{F}_2^d} \hat{f}_z(u) \hat{g}_z(u) |u\rangle |z\rangle \right) |0\rangle + |\perp\rangle, \end{aligned}$$

where the ancillary component of $|\perp\rangle$ is orthogonal to the all-zero state.

4: Apply $H^{\otimes d}$ again to the $|u\rangle$ register:

$$\begin{aligned} & \frac{1}{2^{d/2} G \sqrt{\mathcal{N}_f}} \sum_{z \in \mathcal{Z}} \sum_{x \in \mathbb{F}_2^d} (f_z * g_z)(x) |x\rangle |z\rangle |0\rangle + |\perp\rangle \\ &= \sqrt{\frac{\mathcal{N}_{f*g}}{G^2 \mathcal{N}_f}} |\text{Conv}(f, g)\rangle |0\rangle + |\perp\rangle. \end{aligned}$$

correspond to different realizations of this marking operation. We distinguish three types: BasicDist, LinDist, and SimonDist. SimonDist is another coherent distinguisher considered in [3]. Since it is not used in our attacks, we omit its details here.

BasicDist. BasicDist does not attempt to distinguish the active subkey directly from a statistical property. Instead, for a candidate $(z_{\text{int}}, z_{\text{ext}})$, the remaining master-key bits are searched over using Grover's algorithm, and each resulting full-key candidate is verified by reversible encryption.

For the concrete attacks below, we evaluate BasicDist under the following key-completion model. Conditioned on a correct active-subkey candidate of dimension $\kappa = |k_{\text{int}}| + |k_{\text{ext}}|$, the compatible master-key candidates are assumed to admit an efficiently reversible parametrization by $K - \kappa$ binary variables. We do not exploit additional cipher-specific key-schedule relations in this completion model. The exponential cost of key completion is therefore accounted for through a Grover search over $2^{K-\kappa}$ compatible master-key candidates. In the

concrete gate counts, the reversible circuit cost of the cipher-specific key schedule is omitted consistently with the counting convention used in the application sections.

Under this model, Grover search requires approximately $\frac{\pi}{4} 2^{(K-\kappa)/2}$ iterations. Let $r(E)$ denote the number of equivalent reversible block-cipher evaluations required by one Grover iteration for full-key verification. Following the counting convention used in our concrete estimates, we take $r(E) = 4$. Hence, neglecting lower-order operations, the gate complexity of BasicDist is

$$\mathcal{G}_{\text{Basic}} = \frac{\pi}{4} 2^{(K-\kappa)/2} r(E) Q = \pi 2^{(K-\kappa)/2} Q.$$

LinDist. LinDist marks an active-subkey candidate directly based on its linear statistic. Conceptually, it is the quantum analogue of Matsui’s Algorithm 2: for every candidate $\bar{z}$, quantum amplitude estimation is used to estimate its experimental correlation, and the candidate is accepted when the estimated absolute correlation exceeds a prescribed threshold. To obtain a sufficiently reliable coherent marking oracle, the estimation is repeated and a majority vote is performed.

After Walsh-spectrum puncturing, the key-recovery map is generally real-valued rather than Boolean. We therefore use the puncture-aware variant introduced in Section 3.3.

$$\mathcal{G}_{\text{Lin}} \simeq 8qLQ.$$

Here L is chosen so that the majority-vote error remains negligible over all invocations of the distinguisher. Thus, unlike BasicDist, LinDist avoids completing the remaining master-key bits, but its cost depends mainly on the strength of the linear approximation.

C MILP Modeling of Walsh-Spectrum Puncturing

This appendix summarizes the Walsh-spectrum puncturing constraints used in Q-Puncture. We follow the MILP formulation of Hou et al. [24].

Let $f : \mathbb{F}_2^t \rightarrow \mathbb{R}$ be a key-recovery map with Walsh spectrum $\hat{f}$. For each $u \in \mathbb{F}_2^t$, introduce a binary variable $S_u \in \{0, 1\}$ and define

$$\hat{g}(u) = S_u \hat{f}(u).$$

Thus, the retained Walsh support is

$$R = \{u \in \mathbb{F}_2^t : S_u = 1\}.$$

For a prescribed support restriction, retaining the original coefficients on R maximizes the correlation between f and g [24]. The retained spectral-energy ratio is therefore

$$\eta_{\mathcal{P}} = \frac{\sum_{u \in \mathbb{F}_2^t} S_u \hat{f}(u)^2}{\sum_{u \in \mathbb{F}_2^t} \hat{f}(u)^2}.$$

C.1 Bit Puncturing

For each input coordinate i , let $v_i \in \{0, 1\}$ indicate whether the corresponding variable remains active. The support constraints are

$$v_i \leq \sum_{\substack{u \in \mathbb{F}_2^t \\ u_i = 1}} S_u \leq 2^{t-1} v_i, \quad 0 \leq i < t.$$

Hence,

$$V = \sum_{i=0}^{t-1} v_i$$

is the number of active variables after puncturing.

C.2 Affine-Subspace Puncturing

For affine-subspace puncturing, the retained Walsh support is constrained to an affine subspace

$$R \subseteq A = a + V \subseteq \mathbb{F}_2^t, \quad \dim V = r.$$

Equivalently, every retained mask satisfies

$$u = a \oplus \bigoplus_{j=0}^{r-1} c_{u,j} b_j, \quad c_{u,j} \in \mathbb{F}_2,$$

where $b_0, \dots, b_{r-1}$ form a basis of V . Q-Puncture uses the MILP encoding of these affine-membership constraints from [24], with r representing the resulting active Walsh dimension.

D Proof of Lemma 3

In this subsection, let

$$c = \sqrt{\text{ELP}}, \quad \eta_{\mathcal{P}} = 1 - \epsilon_{\mathcal{P}}, \quad T_g = \frac{3}{8} \eta_{\mathcal{P}} c.$$

We work in the full-codebook setting. In the favorable case $c = \sqrt{\text{ELP}}$, the corrected puncturing model gives

$$\Pr \left[\eta_{\mathcal{P}} c - 2\sqrt{\eta_{\mathcal{P}}(1 - \eta_{\mathcal{P}})} 2^{-n/2} \leq |C_g(\bar{k})| \leq \eta_{\mathcal{P}} c + 2\sqrt{\eta_{\mathcal{P}}(1 - \eta_{\mathcal{P}})} 2^{-n/2} \right] \geq 0.95. \quad (46)$$

For every wrong active subkey $\bar{z} \neq \bar{k}$, we have

$$C_g(\bar{z}) \sim \mathcal{N}(0, \eta_{\mathcal{P}} 2^{-n}). \quad (47)$$

Lemma 6. *Let $\kappa_g = |k_{\text{int}}| + |k_{\text{ext}}|$. Assume that $\sqrt{\eta_{\mathcal{P}}} c \geq 4\sqrt{2\kappa_g} 2^{-n/2}$. We assume $\kappa_g \geq 1$. Then, with probability at least $0.95 - (2/e)^{\kappa_g} / \sqrt{2\kappa_g}$, the following inequalities hold simultaneously: $|C_g(\bar{k})| \geq \frac{4}{3} T_g$, $|C_g(\bar{z})| \leq \frac{2}{3} T_g$, where $\bar{z} \neq \bar{k}$.*

Proof. We first consider the correct active subkey. By the assumption of the lemma, we have

$$\sqrt{\eta_{\mathcal{P}}} c \geq 4\sqrt{2\kappa_g} 2^{-n/2} \geq 4\sqrt{1-\eta_{\mathcal{P}}} 2^{-n/2}.$$

By Eq. (46), with probability at least 0.95, we have

$$|C_g(\bar{k})| \geq \eta_{\mathcal{P}} c - 2\sqrt{\eta_{\mathcal{P}}(1-\eta_{\mathcal{P}})} 2^{-n/2} \geq \frac{\eta_{\mathcal{P}} c}{2} = \frac{4}{3} T_g.$$

We now consider the wrong active subkeys. Fix $\bar{z} \neq \bar{k}$, and define

$$Z_{\bar{z}} = C_g(\bar{z}) \eta_{\mathcal{P}}^{-1/2} 2^{n/2}.$$

By Eq. (47), $Z_{\bar{z}} \sim \mathcal{N}(0, 1)$.

Moreover,

$$\frac{\eta_{\mathcal{P}} c / 4}{\sqrt{\eta_{\mathcal{P}}} 2^{-n/2}} = \frac{\sqrt{\eta_{\mathcal{P}}} c 2^{n/2}}{4} \geq \sqrt{2\kappa_g}.$$

Therefore,

$$\Pr \left[|C_g(\bar{z})| \geq \frac{\eta_{\mathcal{P}} c}{4} \right] \leq \Pr [|Z_{\bar{z}}| \geq \sqrt{2\kappa_g}] \leq \frac{1}{\sqrt{\pi\kappa_g}} e^{-\kappa_g}.$$

There are at most $2^{\kappa_g} - 1 < 2^{\kappa_g}$ wrong active subkeys. Applying the union bound, we obtain

$$\Pr \left[\exists \bar{z} \neq \bar{k} : |C_g(\bar{z})| \geq \frac{\eta_{\mathcal{P}} c}{4} = \frac{2}{3} T_g \right] \leq 2^{\kappa_g} \frac{e^{-\kappa_g}}{\sqrt{\pi\kappa_g}} = \frac{(2/e)^{\kappa_g}}{\sqrt{\pi\kappa_g}} \leq \frac{(2/e)^{\kappa_g}}{\sqrt{2\kappa_g}}.$$

Therefore, except with probability at most $(2/e)^{\kappa_g} / \sqrt{2\kappa_g}$, all wrong active subkeys simultaneously satisfy

$$|C_g(\bar{z})| < \frac{2}{3} T_g.$$

Combining the right-key event and the event for all wrong keys by a union bound, the two inequalities hold simultaneously with probability at least

$$0.95 - \frac{(2/e)^{\kappa_g}}{\sqrt{2\kappa_g}}.$$

□

Lemma 7. Let $B_g \geq \max_{\bar{z}, x} |s_{\bar{z}}(x)|$ and assume that $0 < \eta_{\mathcal{P}} c \leq B_g$. Let $q = \lceil 32\pi B_g / \eta_{\mathcal{P}} c \rceil = \lceil 32\pi B_g / \eta_{\mathcal{P}} \sqrt{\text{ELP}} \rceil$. For any fixed active-subkey candidate $\bar{z}$, one execution of QAE_q produces an estimate $\tilde{C}_g(\bar{z})$ satisfying $|\tilde{C}_g(\bar{z}) - C_g(\bar{z})| < \frac{T_g}{3}$ with probability at least $8/\pi^2$. Consequently, with the same probability, when $\bar{z} = \bar{k}$, we have $|\tilde{C}_g(\bar{z})| \geq T_g$, whereas when $\bar{z} \neq \bar{k}$, we have $|\tilde{C}_g(\bar{z})| < T_g$.

Proof. For a fixed candidate $\bar{z}$, the amplitude-encoding procedure defines

$$p_g(\bar{z}) = \frac{1}{2} \left(1 + \frac{C_g(\bar{z})}{B_g} \right).$$

For simplicity, write $p_g = p_g(\bar{z})$.

Quantum amplitude estimation with parameter q returns an estimate $\tilde{p}_g$ such that

$$|\tilde{p}_g - p_g| \leq \frac{2\pi\sqrt{p_g(1-p_g)}}{q} + \frac{\pi^2}{q^2}$$

with probability at least $8/\pi^2$.

Since $p_g(1-p_g) \leq \frac{1}{4}$, we obtain

$$|\tilde{p}_g - p_g| \leq \frac{\pi}{q} + \frac{\pi^2}{q^2}.$$

The estimated punctured statistic is obtained from $\tilde{p}_g$ through the affine transformation

$$\tilde{C}_g(\bar{z}) = B_g(2\tilde{p}_g - 1).$$

It follows that

$$|\tilde{C}_g(\bar{z}) - C_g(\bar{z})| = 2B_g|\tilde{p}_g - p_g| \leq \frac{2\pi B_g}{q} + \frac{2\pi^2 B_g}{q^2}.$$

By the definition of q ,

$$q \geq \frac{32\pi B_g}{\eta_{\mathcal{P}} c}.$$

Therefore,

$$\frac{2\pi B_g}{q} \leq \frac{\eta_{\mathcal{P}} c}{16},$$

and

$$\frac{2\pi^2 B_g}{q^2} \leq \frac{2\pi^2 B_g (\eta_{\mathcal{P}} c)^2}{(32\pi B_g)^2} = \frac{\eta_{\mathcal{P}}^2 c^2}{512 B_g} \leq \frac{\eta_{\mathcal{P}} c}{512}.$$

Combining the above, we have

$$|\tilde{C}_g(\bar{z}) - C_g(\bar{z})| \leq \left(\frac{1}{16} + \frac{1}{512} \right) \eta_{\mathcal{P}} c = \frac{33}{512} \eta_{\mathcal{P}} c < \frac{\eta_{\mathcal{P}} c}{8} = \frac{T_g}{3}.$$

For the correct active subkey, Lemma 6 gives

$$|\tilde{C}_g(\bar{k})| \geq |C_g(\bar{k})| - |\tilde{C}_g(\bar{k}) - C_g(\bar{k})| > \frac{4}{3}T_g - \frac{1}{3}T_g = T_g.$$

For every wrong active subkey,

$$|\tilde{C}_g(\bar{z})| \leq |C_g(\bar{z})| + |\tilde{C}_g(\bar{z}) - C_g(\bar{z})| < \frac{2}{3}T_g + \frac{1}{3}T_g = T_g.$$

Thus, a single QAE execution makes the correct threshold decision with probability at least $8/\pi^2$. $\square$

Proof (Lemma 3). By Lemmas 6 and 7, each coherent threshold test returns the correct decision with probability at least

$$p \geq 8/\pi^2 > 0.8.$$

Here, this probability is the squared norm of the correct-decision subspace of one QAE output; fresh QAE work registers make the L outputs a tensor product, so the usual binomial tail applies to the majority projector. Let U_1 denote the unitary operator that executes L QAE instances, and let U_2 denote the flag flip controlled by the majority vote. Then

$$\text{PLinDist}_g = U_1^\dagger U_2 U_1.$$

For a fixed candidate $\bar{z}$, let cnt be the number of correct threshold decisions among the L executions. Then $\text{cnt} \sim \text{Bin}(L, p)$, with $p \geq 0.8$. By the Chernoff bound,

$$\Pr \left[\text{cnt} < \frac{L}{2} \right] \leq 2^{-0.09L}.$$

Therefore, the norm of the subspace on which the majority vote is incorrect is at most $2^{-0.045L}$.

On the complementary subspace, PLinDist_g acts identically to Dist_{eff} . Since both operators are unitary, for every computational basis state $|\bar{z}, b\rangle$,

$$\|\text{PLinDist}_g |\bar{z}, b\rangle - \text{Dist}_{\text{eff}} |\bar{z}, b\rangle\| \leq 2^{1-0.045L}.$$

Now let

$$|\psi\rangle = \sum_{\bar{z}, b} \alpha_{\bar{z}, b} |\bar{z}, b\rangle$$

be an arbitrary normalized state. Since the effective-candidate register contains κ_g qubits, the number of basis states $|\bar{z}, b\rangle$ is 2^{κ_g+1} . By the triangle inequality and the Cauchy-Schwarz inequality,

$$\begin{aligned} & \|(\text{PLinDist}_g - \text{Dist}_{\text{eff}}) |\psi\rangle\| \\ & \leq 2^{1-0.045L} \sum_{\bar{z}, b} |\alpha_{\bar{z}, b}| \\ & \leq 2^{1-0.045L} \sqrt{2^{\kappa_g+1}} \left(\sum_{\bar{z}, b} |\alpha_{\bar{z}, b}|^2 \right)^{1/2} \\ & \leq 2^{\kappa_g/2-0.045L+2}. \end{aligned}$$

Therefore,

$$\|\text{PLinDist}_g - \text{Dist}_{\text{eff}}\| \leq 2^{\kappa_g/2-0.045L+2}.$$

Finally, one QAE execution uses $O(q)$ applications of the punctured amplitude-encoding operator and its inverse. Including the reverse execution required for

uncomputation, and using the same conservative counting as for the original LinDist, the cost of each repetition is at most $8qQ$. Hence,

$$\mathcal{G}(\text{PLinDist}_g) \approx 8qLQ.$$

Substituting $q = \left\lceil \frac{32\pi B_g}{\eta_P c} \right\rceil$ and $c = \sqrt{\text{ELP}}$ yields

$$\mathcal{G}(\text{PLinDist}_g) = O\left(L \frac{B_g}{\eta_P \sqrt{\text{ELP}}} Q\right).$$

This completes the proof. $\square$

E Details of the Serpent Attack

E.1 Specification of Serpent

We use the final AES-candidate version of Serpent [1] in its bitslice representation. Serpent is a 128-bit block cipher with 32 rounds and supports 128-, 192-, and 256-bit master keys. The 128-bit state is represented by four 32-bit words

$$X = (X_0, X_1, X_2, X_3),$$

where bit 0 is the least significant bit of each word. We use the bitslice representation throughout and therefore omit the conventional initial and final permutations. The key schedule generates 33 128-bit round keys, denoted by $K_0, \dots, K_{32}$.

Let S_i denote the parallel bitslice application of $S_{i \bmod 8}$. Encryption is written as

$$B_0 = P, \quad B_{i+1} = \begin{cases} L(S_i(B_i \oplus K_i)), & 0 \leq i \leq 30, \\ \mathcal{S}_{31}(B_{31} \oplus K_{31}) \oplus K_{32}, & i = 31. \end{cases}$$

For each bit position $j \in \{0, \dots, 31\}$, the four bits $X_0[j], X_1[j], X_2[j], X_3[j]$ form the input to one bitsliced S-box, with

$$x_j = X_0[j] + 2X_1[j] + 4X_2[j] + 8X_3[j].$$

The eight S-boxes are listed in Table 4.

The linear transformation L is evaluated sequentially as

$$\begin{aligned} X_0 &\leftarrow X_0 \lll 13, & X_2 &\leftarrow X_2 \lll 3, \\ X_1 &\leftarrow X_1 \oplus X_0 \oplus X_2, & X_3 &\leftarrow X_3 \oplus X_2 \oplus (X_0 \ll 3), \\ X_1 &\leftarrow X_1 \lll 1, & X_3 &\leftarrow X_3 \lll 7, \\ X_0 &\leftarrow X_0 \oplus X_1 \oplus X_3, & X_2 &\leftarrow X_2 \oplus X_3 \oplus (X_1 \ll 7), \\ X_0 &\leftarrow X_0 \lll 5, & X_2 &\leftarrow X_2 \lll 22. \end{aligned}$$

Here $\lll$ denotes left rotation and $\ll$ denotes logical left shift with truncation to 32 bits.

Table 4. The eight S-boxes of Serpent.

x	0	1	2	3	4	5	6	7	8	9	A	B	C	D	E	F
$S_0(x)$	3	8	F	1	A	6	5	B	E	D	4	2	7	0	9	C
$S_1(x)$	F	C	2	7	9	0	5	A	1	B	E	8	6	D	3	4
$S_2(x)$	8	6	7	9	3	C	A	F	D	1	E	4	0	B	5	2
$S_3(x)$	0	F	B	8	C	9	6	3	D	1	2	4	A	7	5	E
$S_4(x)$	1	F	8	3	C	0	B	6	2	5	4	A	9	E	7	D
$S_5(x)$	F	5	2	B	4	A	9	C	0	3	E	8	D	6	7	1
$S_6(x)$	7	2	C	5	8	4	6	B	E	9	1	F	D	3	A	0
$S_7(x)$	1	D	F	0	E	8	2	B	7	4	C	A	9	3	5	6

E.2 Serpent-Specific Cost Parameters

The plaintext-side map contains 15 active S_2 components. For every component mask occurring in the trail, the largest absolute LAT entry is 8. Since the key-dependent terms only change the signs of the Walsh coefficients,

$$G_L = 8^{15} = 2^{45}.$$

We count Toffoli gates and omit the key-schedule overhead. Key additions and the linear layer are Clifford operations. Using $\text{MC}(S_3) = \text{MC}(S_7) = 5$ and $\text{MC}(S_i) = 4$ for the remaining Serpent S-boxes [40], clean computation and uncomputation over the 11 attacked rounds give

$$\mathcal{G}(E_{11}) = 2 \cdot 32(3 \cdot 5 + 8 \cdot 4) = 3008 \simeq 2^{11.55}.$$

For g Fourier, let $C_{\text{LAT}} = 64$ denote the cost of a reversible local LAT lookup, and use a clean controlled-increment cost

$$C_{\text{inc}}(w) = 2w - 1.$$

The 15 plaintext-side components give

$$\mathcal{G}(G_1) = 2 \cdot 15 \cdot 64 + 14 + 14C_{\text{inc}}(4) = 2032.$$

If the ciphertext side contains t local components in total and $w = \lceil \log_2(t+1) \rceil$, its cost is

$$\mathcal{G}(G_2) = 2t \cdot 64 + (t-1) + (t-1)C_{\text{inc}}(w).$$

For the selected Serpent-256 profile, $t = 7 + 26 = 33$ and $w = 6$, whereas for Serpent-192, $t = 7 + 21 = 28$ and $w = 5$. Including the final flag/counter combination gives

$$\mathcal{G}(g \text{ Fourier})_{256} = 2032 + (2 \cdot 33 \cdot 64 + 32 + 32 \cdot 11) + 8 = 6648 \simeq 2^{12.70},$$

$$\mathcal{G}(g \text{ Fourier})_{192} = 2032 + (2 \cdot 28 \cdot 64 + 27 + 27 \cdot 9) + 8 = 5894 \simeq 2^{12.53}.$$

E.3 Serpent-Specific Q-Puncture Model

For the 11-round attack of Section 6.1, puncturing is applied to the two ciphertext-side recovery rounds adjacent to the 8-round trail of correlation 2^{-46} . The seven active S_3 components in R_{11} are

$$\mathcal{A}_R = \{2, 9, 12, 14, 19, 24, 29\}, \quad (\alpha_j)_{j \in \mathcal{A}_R} = (1, E, 2, B, 3, B, B).$$

For each $j \in \mathcal{A}_R$, define

$$f_j(y) = (-1)^{\langle \alpha_j, S_3^{-1}(y) \rangle}, \quad \hat{f}_j(u) = L_{S_3}(\alpha_j, u).$$

For every admissible nonempty retained support

$$U_{j,p} \subseteq \text{Supp}(\hat{f}_j),$$

let

$$r_{j,p} = \dim \text{Aff}(U_{j,p}), \quad d_{j,p} = \log_2 \frac{256}{\sum_{u \in U_{j,p}} \hat{f}_j(u)^2}.$$

With binary selection variables $x_{j,p}$,

$$\sum_p x_{j,p} = 1, \quad r = \sum_{j,p} r_{j,p} x_{j,p}, \quad d_{\mathcal{P}} = \sum_{j,p} d_{j,p} x_{j,p}.$$

The retained masks are propagated from R_{11} through LT^{-T} . Let $y_q \in \{0, 1\}$ indicate whether nibble q of R_{12} is activated. Then

$$l = 4 \sum_{q=0}^{31} y_q, \quad |k_{\text{int}}| = r, \quad |k_{\text{ext}}| = 52 + l.$$

E.4 Selected Puncturing Supports for Serpent

Serpent-256. The selected profile is

$$(d_{\mathcal{P}}, r, l) = \left(\log_2 \frac{64}{15}, 24, 104 \right), \quad \eta_{\mathcal{P}} = \frac{15}{64}.$$

Table 5 lists the retained Walsh coefficients, where $u : c$ denotes $\hat{f}_j(u) = c$.

Propagation through LT^{-T} activates

$$\mathcal{C}_{256} = \{0, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 17, 19, 20, 21, 22, 24, 26, 27, 29, 31\},$$

giving $|\mathcal{C}_{256}| = 26, l = 104$.

Serpent-192. The selected profile is

$$(d_{\mathcal{P}}, r, l) = (12, 4, 84), \quad \eta_{\mathcal{P}} = 2^{-12}.$$

The retained Walsh coefficients are listed in Table 6.

Propagation through LT^{-T} activates

$$\mathcal{C}_{192} = \{0, 2, 3, 4, 5, 7, 8, 10, 12, 13, 14, 15, 17, 19, 21, 22, 24, 26, 27, 29, 31\},$$

giving $|\mathcal{C}_{192}| = 21, l = 84$.

Table 5. Selected R_{11} puncturing strategy for 11-round Serpent-256. Indices having the same retained Walsh coefficient are grouped together.

Position	α_j	Retained Walsh spectrum	d_j	r_j
2	1	+8 : {E}, -8 : {6} +4 : {1, 5, 7, 9, D, F} -4 : {3, B}	0	4
9	E	+8 : {8}, -8 : {2} +4 : {1, 3, 5, 9, B, F} -4 : {7, D}	0	4
12	2	+8 : {E} +4 : {2, 6, C}, -4 : {8}	1	3
14	B	+8 : {8}, -8 : {9} +4 : {2, 3, 4, 5, C, D} -4 : {A, B}	0	4
19	3	+8 : {1, B} +4 : {2, 3, 7, C} -4 : {6, 8, 9, D}	0	4
24	B	+8 : {8}, -8 : {9} +4 : {4, 5, C, D}	$\log_2(4/3)$	3
29	B	+8 : {8}, -8 : {9} +4 : {C, D}	$\log_2(8/5)$	2
Total			$\simeq 2.09$	24

Table 6. Selected R_{11} puncturing strategy for 11-round Serpent-192. Indices having the same retained Walsh coefficient are grouped together.

Position	α_j	Retained Walsh spectrum	d_j	r_j
2	1	+8 : {E}	2	0
9	E	-8 : {2}	2	0
12	2	+8 : {E}	2	0
14	B	+4 : {2, 3, 4, 5}	2	2
19	3	+8 : {1}	2	0
24	B	+8 : {8}, -8 : {9}	1	1
29	B	+8 : {8}, -8 : {9}	1	1
Total			12	4

F Details of the RECTANGLE Attacks

F.1 Specification of RECTANGLE

RECTANGLE [41] is a lightweight SPN block cipher with a 64-bit block size and 80- or 128-bit master keys. We consider RECTANGLE-128, which consists

of 25 rounds followed by a final AddRoundKey operation. Rows and columns are indexed from zero.

The 64-bit state is represented as a 4×16 binary array

$$A = (a_{i,j})_{0 \leq i < 4, 0 \leq j < 16},$$

where column j represents the nibble

$$\text{Col}(j) = a_{3,j} \| a_{2,j} \| a_{1,j} \| a_{0,j}.$$

Each round consists of AddRoundKey, SubColumn, and ShiftRow. SubColumn applies the 4-bit S-box

$$\begin{array}{c|cccccccccccccccc} x & 0 & 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 & 9 & A & B & C & D & E & F \\ \hline S(x) & 6 & 5 & C & A & 1 & E & 7 & 9 & B & 0 & 3 & D & 8 & F & 4 & 2 \end{array}$$

independently to the 16 columns. ShiftRow rotates rows 0, 1, 2, 3 to the left by 0, 1, 12, 13 positions, respectively.

Let $K_0, \dots, K_{25}$ denote the 64-bit round keys generated from the 128-bit master key. If X_i denotes the state before round i , then

$$X_{i+1} = \text{ShiftRow}(\text{SubColumn}(X_i \oplus K_i)), \quad 0 \leq i \leq 24,$$

followed by

$$C = X_{25} \oplus K_{25}.$$

F.2 Active-subkey Dimensions

14-round attack. The 11-round trail used in the $1 + 11 + 2$ construction is

$$\alpha_{11} = 00000000000060002, \quad \beta_{11} = 0600004000002000, \quad c = 2^{-25}.$$

The plaintext-side recovery contains the components $(3, 2), (8, 4), (15, 2)$, whose local Walsh supports each have affine dimension two. Hence $|k_{\text{int}}^L| = 0, |k_{\text{ext}}^L| = 6$.

The ciphertext-side components are $(3, 2), (9, 4), (14, 6)$, whose local Walsh supports have total affine dimension $|k_{\text{int}}^R| = 6$. Their propagation activates

$$\mathcal{C}_{11} = \{0, 3, 4, 5, 6, 9, 10, 11, 14, 15\},$$

and therefore $|k_{\text{ext}}^R| = 4|\mathcal{C}_{11}| = 40$. The resulting unpunctured effective-key dimensions are $|k_{\text{int}}^R| = 6$ and $|k_{\text{ext}}^R| = 46$.

16-round attack. The 12-round trail used in the $2 + 12 + 2$ construction is

$$\alpha_{12} = 00000000000060002, \quad \beta_{12} = 4000820000420000, \quad c = 2^{-28}.$$

The plaintext-side components are $(3, 2), (8, 4), (15, 2)$. Their full local Walsh supports have affine dimensions 2, 4, 2 and propagate to nine outer columns, giving $(r_L, l_L) = (8, 36)$.

The ciphertext-side components are $(4, 2), (5, 4), (10, 2), (11, 8), (15, 4)$, with affine dimensions 2, 2, 2, 4, 2. They propagate to twelve outer columns, giving $(r_R, l_R) = (12, 48)$. Thus, before puncturing, $(|k_{\text{int}}|, |k_{\text{ext}}|) = (r_L + r_R, l_L + l_R) = (20, 84)$.

F.3 Circuit and Spectral Parameters

For the 14-round PLinDist attack, the three plaintext-side component masks have maximum absolute LAT entry 8. Hence

$$G_L = 8^3 = 2^9.$$

We count Toffoli gates and omit the key-schedule overhead. Since AddRound-Key and ShiftRow are linear and $\text{MC}(S) = 4$, clean computation and uncomputation over R rounds give

$$\mathcal{G}(E_R) = 2 \cdot 16 \cdot R \cdot 4.$$

Consequently,

$$\mathcal{G}(E_{14}) = 1792 \simeq 2^{10.81}, \quad \mathcal{G}(E_{16}) = 2048 = 2^{11}.$$

For g Fourier, let

$$C_{\text{LAT}} = 64, \quad C_{\text{inc}}(t) = 2t - 1.$$

For the 14-round attack,

$$\mathcal{G}(g \text{ Fourier}^{14}) = 392 + 1760 + 4 = 2156 \simeq 2^{11.07}.$$

For the bilateral 16-round attack, the left side contains $3 + 9 = 12$ local LAT factors and the right side contains $5 + 12 = 17$. Hence

$$\begin{aligned} C_L &= 2 \cdot 12 \cdot 64 + 11 + 11C_{\text{inc}}(4) = 1624, \\ C_R &= 2 \cdot 17 \cdot 64 + 16 + 16C_{\text{inc}}(5) = 2336. \end{aligned}$$

Combining the two sides gives

$$\mathcal{G}(g \text{ Fourier}^{16}) = 1624 + 2336 + 1 + 7 = 3968 \simeq 2^{11.95}.$$

F.4 Computation of B_g

For a ciphertext-side component with mask m and retained Walsh support P_m , define

$$f_m(x) = (-1)^{\langle m, S^{-1}(x) \rangle},$$

and

$$g_{m, P_m}(x) = \frac{1}{16} \sum_{u \in P_m} L_{S^{-1}}(u, m) (-1)^{\langle u, x \rangle}.$$

Let

$$B_{m, P_m} = \max_{x \in \mathbb{F}_2^4} |g_{m, P_m}(x)|.$$

For the 14-round attack, the three ciphertext-side components give

Component	P_m	$\{u : L_{S^{-1}}(u, m)\}_{u \in P_m}$	B_{m, P_m}
(3, 2)	{6}	{6 : 8}	1/2
(9, 4)	{6}	{6 : -8}	1/2
(14, 6)	{8, 9, C, D}	{8 : 8, 9 : 8, C : -8, D : 8}	1

where {8, 9, C, D} is the complete nonzero Walsh support for the component with mask 6.

Since the three local S-box inputs are disjoint,

$$B_g = \prod_j B_{m_j, P_j} = \frac{1}{2} \cdot \frac{1}{2} \cdot 1 = \frac{1}{4} = 2^{-2}.$$

F.5 Selected Puncturing Supports for RECTANGLE

Columns are numbered from zero, with column zero corresponding to the right-most hexadecimal nibble.

14-round attack. For the ciphertext side of the 1 + 11 + 2 construction, the selected supports are

Component	Retained support	d_j	r_j
(3, 2)	{6}	2	0
(9, 4)	{6}	2	0
(14, 6)	{8, 9, C, D}	0	2

and hence $d_{\mathcal{P}} = 4, r = 2$. Propagation activates

$$\mathcal{C}_{14} = \{4, 5, 10, 11, 14, 15\},$$

giving $l = 24, |k_{\text{int}}| = 2, |k_{\text{ext}}| = 30$.

16-round attack. For the plaintext side of the 2+12+2 construction, the selected supports are

Component	Retained support	d_j	r_j
(3, 2)	{5, 7}	1	1
(8, 4)	{3, 7}	1	1
(15, 2)	{5, 7, D, F}	0	2

with $d_L = 2, r_L = 4$. The propagated outer columns are

$$\mathcal{C}_L = \{2, 3, 7, 8, 12, 14, 15\},$$

giving $l_L = 28$.

On the ciphertext side, the selected supports are

Component	Retained support	d_j	r_j
(4, 2)	$\{6, 7, A, B\}$	0	2
(5, 4)	$\{6, 7, E, F\}$	0	2
(10, 2)	$\{6, 7, A, B\}$	0	2
(11, 8)	$\{3, 4, 5, 6, 7, 9, C, D, E, F\}$	0	4
(15, 4)	$\{6, 7, E, F\}$	0	2

with $d_R = 0, r_R = 12$. Propagation activates

$$\mathcal{C}_R = \{0, 1, 2, 4, 5, 6, 7, 8, 10, 11, 12, 15\},$$

giving $l_R = 48$.

The aggregate puncturing profile is therefore

$$(d_{\mathcal{P}}, r, l) = (d_L + d_R, r_L + r_R, l_L + l_R) = (2, 16, 76), \quad \eta_{\mathcal{P}} = 2^{-2}.$$

G Details of the PIPO-128 Attacks

G.1 Description of PIPO-128

PIPO is a lightweight bitsliced block cipher proposed by Kim et al. [27]. We consider the PIPO-64/128 variant, which has a 64-bit block size, a 128-bit master key, and 13 rounds.

The internal state is represented as an 8×8 binary array

$$X = (x_{j,c})_{0 \leq j, c < 8},$$

or equivalently as eight 8-bit rows $X[0], \dots, X[7]$. Column c is interpreted as the input to one 8-bit S-box,

$$x_c = \sum_{j=0}^7 2^j x_{j,c},$$

so that row 0 supplies the least significant bit. The nonlinear layer $\mathcal{S}$ applies the fixed PIPO 8-bit S-box S_8 independently to the eight columns.

For PIPO-128, write the master key as

$$K = K_1 \| K_0, \quad K_0, K_1 \in \mathbb{F}_2^{64}.$$

The whitening and round keys are

$$RK_i = K_{i \bmod 2}, \quad 0 \leq i \leq 13.$$

Starting from

$$X^{(0)} = P \oplus RK_0,$$

round i , for $1 \leq i \leq 13$, is

$$X^{(i)} = \mathcal{R}\left(\mathcal{S}(X^{(i-1)})\right) \oplus RK_i \oplus c_i,$$

where $c_i = i$ is the round constant. The linear layer $\mathcal{R}$ rotates the eight state rows by

$$(0, 7, 4, 3, 6, 5, 1, 2)$$

positions, respectively.

For linear analysis, a boundary mask is written as

$$\gamma = (\gamma_0, \dots, \gamma_7), \quad \gamma_c \in \mathbb{F}_2^8,$$

where γ_c is the 8-bit mask applied to column c . We use the notation (c, m) for a nonzero mask m on column c . The LAT of S_8 is used throughout the transition and puncturing computations below.

G.2 5-Round Approximation Family

We consider the 6000 boundary masks having one or two nonzero byte components whose inverse propagation through the PIPO linear layer also activates at most two columns. Let $P \in \mathbb{R}^{6000 \times 6000}$ be the sparse one-round linear-probability transition matrix derived from the PIPO S-box LAT. For a valid one-round transition, let (a_c, b_c) denote the input/output masks of the S-box in column c . Its transition weight is

$$P_{u,v} = \prod_{c=0}^7 \left(\frac{L_{S_8}(a_c, b_c)}{2^8} \right)^2,$$

with $P_{u,v} = 0$ when the masks are incompatible with the R-layer propagation. The five-round expected linear potential between boundary masks i and j is $\text{ELP}_{i,j} = (P^5)_{i,j}$.

For the $M = 24$ approximation family used in Section 8, the three rotational representatives are listed in Table 7. These are also illustrated in Fig. 4. In the notation (c, m) , $c \in \{0, \dots, 7\}$ denotes a zero-based PIPO column and m an 8-bit hexadecimal mask. Each representative generates eight cyclic shifts under

$$(c, m) \mapsto ((c + t) \bmod 8, m), \quad 0 \leq t < 8.$$

Table 7. Representatives of the three selected rotational orbits.

Orbit	Left boundary	Right boundary	ELP per shift
20	(5, 28), (7, 40)	(3, 10), (4, 80)	2^{-50}
23	(5, 28), (7, 40)	(3, 50), (4, 80)	2^{-50}
24	(5, 28), (7, 40)	(1, 80), (4, A0)	2^{-50}




Fig. 4. Representative 5-round linear trails for the three selected rotational orbits used in the PIPO-128 linear-hull analysis.

G.3 PIPO-Specific Puncturing Model

For an active 8-bit S-box component, let $W(u)$ denote the corresponding Walsh row or column. For each enumerated affine flat $A \subseteq \mathbb{F}_2^8$, define

$$\mathcal{P}(A) = A \cap \text{supp}(W).$$

For every nonempty retained support, we record

$$E(A) = \sum_{u \in \mathcal{P}(A)} W(u)^2, \quad d(A) = -\log_2 \frac{E(A)}{2^{16}},$$

together with

$$r(A) = \dim A, \quad q(A) = \bigvee_{u \in \mathcal{P}(A)} u.$$

Different affine flats may induce the same retained set; $r(A)$ denotes the dimension of the selected flat.

The mask $q(A)$ is propagated through the PIPO bit permutation. If the resulting union of active columns is $\mathcal{C}$, the external dimension is $l = 8|\mathcal{C}|$. For multiple local components, retained energies multiply, while the puncturing losses and internal dimensions add. Dominated candidates with no better retained energy, dimension, or propagated activity are discarded.

The resulting profiles are evaluated using the ML-Basic objective in Section 8. The same aggregate profile is required for the three rotational representatives in Table 7; their cyclic shifts inherit the corresponding profile.

Table 8. Selected puncturing profile for the 8-round PIPO-128 attack.

Component	Retained set	d_i	r_i	q_i
(5, 28)	all nonzero coefficients	0	8	FF
(7, 40)	affine-flat intersection above $\log_2(8/5)$	4	4	3F
Total		$\log_2(8/5)$	12	$l = 64$

G.4 Selected 8-Round Puncturing Profile

For the 8-round WSP attack reported in Section 8, puncturing is applied to the plaintext-side recovery. All three selected rotational orbits have the same two active boundary components.

For component (5, 28), the complete nonzero Walsh spectrum is retained. For component (7, 40), the retained support is the intersection with

$$0C + \langle 01, 02, 10, 20 \rangle.$$

Its retained nonzero Walsh coefficients are

$$\{0C:64, 0D:64, 0E:64, 0F:-64, 1C:-64, 1D:-64, 1E:-64, 1F:64, \\ 2C:-32, 2D:-32, 2E:-32, 2F:32, 3C:-32, 3D:-32, 3E:-32, 3F:32\}.$$

Thus, the aggregate puncturing profile is

$$(d_P, r, l) = \left(\log_2 \frac{8}{5}, 12, 64 \right), \quad \eta_P = \frac{5}{8}.$$

G.5 Selected 9-Round Puncturing Profile

For the 9-round attack, puncturing is applied to both recovery sides.

On the plaintext side, the retained coefficients are

$$(5, 28) : \{48:64\}, \quad (7, 40) : \{0C:64\}.$$

Each selected affine candidate has dimension zero and puncturing loss four. Their propagation activates columns

$$\mathcal{C}_L = \{2, 3, 4\},$$

giving

$$(d_L, r_L, l_L) = (8, 0, 24).$$

Table 9. Selected ciphertext-side witnesses for the 9-round PIPO-128 attack.

Orbit	First component	Second component	Propagated columns
20	$(3, 10) : \mathcal{S}_{10}, r = 3, q = \text{D7}$	$(4, 80) : \mathcal{S}_{80}, r = 6, q = \text{FB}$	$\{1, 2, 3, 4, 5, 6, 7\}$
23	$(3, 50) : \mathcal{S}_{50}, r = 4, q = \text{DF}$	$(4, 80) : \mathcal{S}_{80}, r = 5, q = \text{FB}$	$\{1, 2, 3, 4, 5, 6, 7\}$
24	$(1, 80) : \mathcal{S}_{80}, r = 4, q = \text{FB}$	$(4, \text{A0}) : \mathcal{S}_{\text{A0}}, r = 5, q = \text{BF}$	$\{0, 1, 2, 3, 4, 6, 7\}$

On the ciphertext side, define

$$\begin{aligned}
\mathcal{S}_{10} &= \{16:-32, 52:32, 83:64, 87:64, \text{C3}:64, \text{C7}:-64\}, \\
\mathcal{S}_{50} &= \{0\text{A}:32, 12:32, 44:64, 45:64, 5\text{C}:-32, 5\text{D}:32, 82:32, 83:32, \\
&\quad 9\text{A}:32, 9\text{B}:-32, \text{CC}:-32, \text{D4}:32\}, \\
\mathcal{S}_{80} &= \{92:64, 93:-64, 9\text{A}:-64, 9\text{B}:64, \text{B2}:64, \text{B3}:-64, \text{BA}:64, \text{BB}:64, \\
&\quad \text{D2}:-64, \text{D3}:-64, \text{DA}:64, \text{DB}:64, \text{F2}:-64, \text{F3}:64, \text{FA}:-64, \text{FB}:64\}, \\
\mathcal{S}_{\text{A0}} &= \{0\text{B}:32, 13:-32, 17:-32, 1\text{F}:-32, 27:-64, 2\text{B}:32, 33:-32, 37:32, \\
&\quad 3\text{F}:-32, 88:-32, 94:16, 98:-32, 9\text{C}:16, \text{A4}:32, \text{A8}:-32, \text{B4}:-16, \text{B8}:32, \text{BC}:16\}.
\end{aligned}$$

The retained energy of each of $\mathcal{S}_{10}$, $\mathcal{S}_{50}$, and $\mathcal{S}_{\text{A0}}$ is 18432, giving

$$d_{\mathcal{P}} = -\log_2 \frac{18432}{65536} = \log_2 \frac{32}{9}.$$

The support $\mathcal{S}_{80}$ retains full energy and therefore has zero puncturing loss.

Thus, each rotational representative realizes

$$(d_R, r_R, l_R) = \left(\log_2 \frac{32}{9}, 9, 56 \right).$$

Combining the two recovery sides gives

$$d_{\mathcal{P}} = 8 + \log_2 \frac{32}{9} = \log_2 \frac{8192}{9} \simeq 9.83, \quad r = 9, \quad l = 80,$$

and hence

$$\eta_{\mathcal{P}} = \frac{9}{8192}.$$

H Puncturing Loss versus Quantum Cost

We illustrate the effect of puncturing on the quantum attack cost using 11-round Serpent-256 and 16-round RECTANGLE-128. For each attainable puncturing loss $d = -\log_2 \eta_{\mathcal{P}}$, we select the profile with minimum total cost

$$\mathcal{G}_{\text{total}} = R_{\text{outer}}(\mathcal{G}_{\text{NCOR}} + \mathcal{G}_{\text{Basic}}),$$

and evaluate $\mathcal{G}_{\text{NCOR}}$ for the same profile. The attack parameters are those used in Sections 6 and 7.




Fig. 5. Puncturing loss versus quantum cost for Serpent-256 and RECTANGLE-128. At each attainable d , both curves correspond to the profile minimizing the total cost. The dashed line denotes the Grover reference, and the upper axis reports $\eta_P = 2^{-d}$.

As shown in Fig. 5, $\mathcal{G}_{\text{NCOR}}$ generally decreases with increasing puncturing, whereas the total cost is non-monotonic. The minimum occurs at $d \simeq 2.09$ for Serpent-256 and $d = 2$ for RECTANGLE-128, illustrating that puncturing should be optimized with respect to the modeled total quantum key-recovery cost rather than the state-preparation cost alone.