Single-Query Quantum Hidden Shift Attacks

Xavier Bonnetain[1] and André Schrottenloher[2]

[1] Université de Lorraine, CNRS, Inria, LORIA, Nancy, France
[2] Univ Rennes, Inria, CNRS, IRISA, Rennes, France
firstname.lastname@inria.fr

Abstract. Quantum attacks using superposition queries are known to break many classically secure modes of operation. While these attacks do not necessarily threaten the security of the modes themselves, since they rely on a strong adversary model, they help us to draw limits on their provable security.

Typically these attacks use the structure of the mode (stream cipher, MAC or authenticated encryption scheme) to embed a period-finding problem, which can be solved with a dedicated quantum algorithm. The hidden period can be recovered with a few superposition queries (e.g., $O(n)$ for Simon's algorithm), leading to state or key-recovery attacks. However, this strategy breaks down if the period changes *at each query*, e.g., if it depends on a nonce.

In this paper, we focus on this case and give dedicated state-recovery attacks on the authenticated encryption schemes Rocca, Rocca-S, Tiaoxin-346 and AEGIS-128L. These attacks rely on a procedure to find a Boolean hidden shift with a *single* superposition query, which overcomes the change of nonce at each query. This approach has the drawback of a lower success probability, meaning multiple independent (and parallelizable) runs are needed.

We stress that these attacks do not break any security claim of the authors, and do not threaten the schemes if the adversary only makes classical queries.

Keywords: Quantum cryptanalysis · Quantum Fourier Transform · Authenticated encryption · Boolean hidden shift · Rocca · Tiaoxin · AEGIS

1 Introduction

Since Shor's algorithm [Sho94], the enhanced computational power of quantum devices has been known to impact the security of public-key cryptosystems. Nowadays, *post-quantum* (public-key) cryptography is structured around several computational problems (e.g., lattice sieving, decoding random codes...) which are believed to remain intractable.

The situation is more favorable in symmetric (secret-key) cryptography, since most of it is expected to remain secure. Generic attacks on primitives are now well understood, for example Grover's quantum search [Gro96] that accelerates the recovery of a secret key from a time $\mathcal{O}(2^{\kappa})$ to $\mathcal{O}(2^{\kappa/2})$, or the BHT algorithm [BHT98] which accelerates n-bit collision search from $\mathcal{O}(2^{n/2})$ to $\mathcal{O}(2^{n/3})$. Many dedicated quantum attacks have also been introduced, whether on block ciphers [BNS19, KLLN16b] or hash functions [HS20]. Most of the time, these attacks reach at most a quadratic speedup (like Grover's search). In this paper, we focus on *superposition attacks* on modes of operation, which are known to allow super-quadratic speedups or sometimes total breaks of classically-secure schemes.

Superposition Queries. The literature separates quantum attacks on symmetric schemes in two categories. In the *Q1 setting*, the adversary has only classical access to the attacked function, typically an encryption scheme or MAC which contains secret information (the

key or internal states). Such attacks follow the main threat model of post-quantum cryptography, where the adversary records computations to decrypt them later in time. In the *Q2 setting*, also named *superposition query model*, the adversary can query the function as a *quantum oracle*, i.e., from within a quantum computation. Obviously, this cannot model a "store now, decrypt later" scenario anymore. Despite this lack of practical applications, Q2 attacks are still a relevant source of information on the quantum security of these schemes, as they are known to break many classically secure modes of operation [KM10, KM12, KLLN16a, LM17]. On the one hand, they can be used as a starting point or motivation for improved Q1 attacks [BHN+19, BSS22]. On the other hand, they can be seen as impossibility results, showing that any security proof must consider an adversary making classical queries to the scheme [ABKM22].

Principle of Q2 Breaks. Consider a symmetric scheme $E_K : \{0,1\}^n \to \{0,1\}^m$ with a secret key K, to which we have quantum access.

Typically, Q2 attacks will combine some *pre-processing* function f and *post-processing* function g so that the function $g \circ E_K \circ f$ has some property that can be exploited. For example, the Even-Mansour cipher: $E_{k_1,k_2} : x \mapsto k_1 \oplus \Pi(k_2 \oplus x)$, where Π is a public permutation, can be attacked by noticing that $E_{k_1,k_2} \oplus \Pi$ is a periodic function on $\mathbb{F}_2^n$, of period k_2 [KM12]. Simon's algorithm [Sim97] can recover this period in $\mathcal{O}(n)$ quantum queries. Other attacks (for example in [BLNS21], using a non-trivial f) may target an internal state instead. In MACs and authenticated encryption (AE) schemes, this can lead to forgeries.

A typical limitation of Q2 attacks is when the construction E_K admits a *nonce* N, like many MACs and AE schemes. It is indeed common [ATTU16] to assume that nonces remain classical values, and that they are not repeated from one Q2 query to another. While many attacks can also bypass the use of nonces [Bon17, BLNS21], they cannot apply in a situation where we would query: $E_{K,N}(x) = f(x \oplus s(K,N))$ where the secret internal state s depends on *K and N*.

New Strategy. In this paper, we use a hidden shift algorithm with a *single query* from [ORR13]. It follows a well-known strategy in quantum computing which was previously applied in [vDHI06, Röt10] and requires, in our case, a combination with a state preparation technique [SLSB19].

We consider several AE schemes, where the recovery of the internal state leads to forgery or key-recovery attacks. Our strategy is to perform a superposition query with several message blocks which, with proper post-processing, can be turned into an oracle:

$$|x\rangle |0\rangle \mapsto |x\rangle g(x \oplus s') |s\rangle ,$$

where g is a function to $\{-1,1\}$, and s and s' are values which, together, allow to determine a whole internal state. We measure s immediately, but we cannot use Simon's algorithm to obtain s' since it depends on the nonce, and will change at the next query.

Instead, we use the hidden shift algorithm from [ORR13]. This algorithm performs a Hadamard transform:

$$\frac{1}{2^{n/2}} \sum_x g(x \oplus s') |x\rangle \xmapsto{H} \frac{1}{2^n} \sum_y (-1)^{s' \cdot y} \widehat{g}(y) |y\rangle ,$$

with $\widehat{g}$ the Walsh-Hadamard transform of g. It then computes a multiplication by $1/\widehat{g}(y)$ in the amplitudes of this state. Such a multiplication cannot succeed with probability 1. In fact, the attack will require many trials, using each time a new random nonce, and even possibly a new secret key. When the multiplication succeeds, we obtain $\sum (-1)^{s' \cdot y} |y\rangle$ which, after another Hadamard transform, gives us s'. With s and s', we solve a system of equations which gives us the full internal state of the scheme.

Table 1: New quantum attacks and comparison with generic attacks ("Grover"). "Toffoli" is an approximate count of the total number of Toffoli gates applied during the attack, derived from the Toffoli count of AES. Approximately 10^3 to 10^4 qubits are required for all attacks, since the internal state of the schemes is of order 10^3 bits.

Target	Type	Setting	Queries per trial	Independent trials	Toffoli	Classical time	Method
Rocca	key	Q1	1 (encr.)	1	2^{145}	negl.	Grover
	forgery	Q2	2^{64} (decr.)	1	2^{80}	negl.	Grover
			1 (encr.)	2^{59}	2^{81}	negl.	Subsection 4.1
			1 (encr.)	$2^{46.4}$	$2^{68.4}$	$2^{125.4}$	Subsection 4.1
Rocca-S	key	Q1	1 (encr.)	1	2^{145}	negl.	Grover
	forgery	Q2	1 (encr.)	2^{30}	2^{101}	negl.	Subsection 4.3
			1 (encr.)	2^{65}	2^{87}	negl.	Subsection 4.3
Tiaoxin	key	Q1	1 (encr.)	1	2^{81}	negl.	Grover
		Q2	1 (encr.)	2^{34}	2^{56}	negl.	Subsection 4.4
AEGIS-128L	key	Q1	1 (encr.)	1	2^{81}	negl.	Grover
	forgery	Q2	2^{33} (decr.) $+1$ (encr.)	2^{27} $+2^{56}$	2^{78}	negl.	Subsection 4.5

The resulting attacks are summarized in Table 1. While we compare them with the gate and query counts of Grover search, one of their features is that the independent trials can be perfectly parallelized. It is well-known that reducing the depth of a Grover search by a factor S increases the computational cost by the same factor S. Therefore, under a limitation in depth, the advantage of our attacks becomes more significant.

Outline. We detail the targeted authenticated encryption schemes in Section 2. In Section 3, we give and analyze the quantum building blocks of linear post-processing, amplitude transduction and single-query hidden shift. In Section 4 we present our attacks. The SageMath [The24] and Python scripts that we used to write down formulas and compute the complexities in our applications are available at: gitlab.inria.fr/capsule/single-query-hidden-shift .

2   Description of the Schemes

In this section, we recall the Authenticated Encryption with Associated Data (AEAD) schemes AEGIS-128L [WP13b], Tiaoxin-346 [Nik16], Rocca [SLN+21] and Rocca-S [NFI]. Some details which are not relevant to our analysis will be omitted. In particular, we omit the processing of Associated Data and the padding of input messages.

The levels of security against key-recovery and forgery are set according to the generic attacks:

- **Key-recovery:** using a single classical known-plaintext query, an adversary can always find the κ-bit key in $\mathcal{O}(2^\kappa)$ computations ($\mathcal{O}(2^{\kappa/2})$ in the quantum setting using Grover's algorithm [Gro96]);

- **Forgery:** with a t-bit tag, an adversary that can make decryption queries can create a forgery in $\mathcal{O}(2^t)$ queries classically. This attack can be accelerated quantumly if one has access to a quantum decryption oracle. This would cost $\mathcal{O}(2^{t/2})$ quantum

Table 2: Summary of parameters for studied AE schemes and their bits of security (nonce-respecting) in the classical setting.

Cipher	Key size	Nonce size	Tag size	Forgeries	Key-recovery
AEGIS-128L	128	128	128	128	128
Tiaoxin-346	128	128	128	128	128
Rocca	256	128	128	128	256
Rocca-S	256	128	256	256	256

queries using Grover's algorithm. Classically or quantumly, this attack is relevant only if the key is larger than the tag.

The designs we study are known to be insecure in the nonce-misuse scenario, i.e., if the adversary is allowed to perform multiple chosen-plaintext queries with the same nonce.

2.1 AEGIS-128L

AEGIS was originally published at SAC [WP13a], and later submitted to the CAESAR competition [WP16]. We will focus here on the variant AEGIS-128L, which can be found in [WP13b]. In the CAESAR competition, AEGIS-128 appeared in the final portfolio for use case 2 (high-performance applications), and AEGIS-128L was a finalist for this use case.

All variants of AEGIS use a large internal state, made of several 128-bit *registers*, and a simple round function which updates this state and mixes it with additional registers of input (e.g., the message blocks). This round function is based on the block cipher standard AES [Nat01].

The AES Round. We denote the AES round function as: $A = \mathsf{MC} \circ \mathsf{SR} \circ \mathsf{SB}$. It applies on a state of 128 bits, represented as a 4×4 matrix of bytes, where the bytes are numbered from 0 to 15, top to bottom and left to right. SB (SubBytes) applies the AES S-Box (denoted SBox) in parallel to all bytes. SR (ShiftRows) shifts row number i in the matrix by i positions left. MC (MixColumns) multiplies each column by the AES MDS matrix.

AEGIS-128L Algorithm. AEGIS-128L accepts a key and a nonce of 128 bits each. The internal state is made of eight 128-bit registers denoted $S[i], 0 \leq i \leq 7$. The round function R takes two additional 128-bit inputs X_0, X_1 and outputs $S' = R(S, X_0, X_1)$ as:

$$
\begin{aligned}
S'[0] &= X_0 \oplus S[0] \oplus A(S[7]) & S'[4] &= X_1 \oplus S[4] \oplus A(S[3]) \\
S'[1] &= S[1] \oplus A(S[0]) & S'[5] &= S[5] \oplus A(S[4]) \\
S'[2] &= S[2] \oplus A(S[1]) & S'[6] &= S[6] \oplus A(S[5]) \\
S'[3] &= S[3] \oplus A(S[2]) & S'[7] &= S[7] \oplus A(S[6])
\end{aligned}
$$

Without AD, the algorithm has the following phases:

- **Initialization:** after loading the key K and nonce N into the state, we run 10 round updates $R(S, N, K)$

- **Encryption:** each round of encryption takes two plaintext blocks M_i, M_i' and returns two ciphertext blocks C_i, C_i'. For all $i = 0$ to $m - 1$:

$$
\begin{cases}
C_i = M_i \oplus S[1] \oplus S[6] \oplus \mathrm{AND}(S[2], S[3]) \\
C_i' = M_i' \oplus S[2] \oplus S[5] \oplus \mathrm{AND}(S[6], S[7]) \\
S \leftarrow R(S, M_i, M_i')
\end{cases}
\tag{1}
$$

Figure 1: Round function of Tiaoxin-346.

where AND denotes the bit-wise Boolean AND.

- **Finalization:** the state update function is called 6 times with X_0, X_1 depending on the AD length and message length. The authentication tag is obtained by XORing the 7 first registers.

Security. Third-party cryptanalysis has shown that AEGIS is insecure under nonce misuse [KEM17] and that it exhibits linear keystream biases [ENP19]. However, these attacks did not contradict its security claims. To the best of our knowledge, there has been no quantum security analysis of AEGIS.

2.2   Tiaoxin-346

Tiaoxin-346 was submitted to the CAESAR competition [Nik16] where it reached the third round. It accepts 128-bit keys and 128-bit tags. The internal state T is made of thirteen 128-bit registers separated into substates T_3, T_4, T_6 with 3, 4 and 6 registers respectively denoted as $T_j[i]$. The round function $R(T, X_0, X_1, X_2)$ takes a 3-register input X_0, X_1, X_2 and updates the state as shown on Figure 1. In particular, it can be noted that the round function processes independently the substates T_j.

In the initialization phase, the key and nonce are loaded in T, then, 15 rounds of the round function $R(T, Z_0, Z_1, Z_0)$ are applied where Z_0 and Z_1 are constants. In the encryption phase, message blocks are also encrypted by pairs M_i, M_i'. For all $i = 0$ to $m - 1$:

$$\begin{cases} T \leftarrow R(T, M_i, M_i', M_i \oplus M_i') \\ C_i = T_3[0] \oplus T_3[2] \oplus T_4[1] \oplus \text{AND}(T_6[3], T_4[3]) \\ C_i' = T_6[0] \oplus T_4[2] \oplus T_3[1] \oplus \text{AND}(T_6[5], T_3[2]) \end{cases} \tag{2}$$

It can be noted that the state update is performed *before* outputting the ciphertexts, and not after like the other designs in this section. Finally, the finalization performs 20 unkeyed rounds $R(T, Z_1, Z_0, Z_1)$ and outputs the tag as the XOR of all registers $T_j[i]$.

Security. An important difference between Tiaoxin and AEGIS is that the round function of Tiaoxin is invertible, as well as the initialization phase. Thus, recovering the internal state at any point of the ciphering process leads to a key-recovery. Furthermore it is enough to recover a single substate T_j.

Figure 2: Round function of Rocca.

A few third-party works have studied the security: a key-recovery attack in a nonce-misuse scenario has been proposed [KEM17], and Tiaoxin reduced to 8 rounds of initialization has been shown to have weak keys [LIMS21]. To the best of our knowledge, there has been no quantum security analysis of Tiaoxin.

2.3   Rocca

Rocca is an AEAD for beyond-5G applications. As such, it also aims at quantum security and uses keys of 256 bits. The internal state S is made of eight 128-bit registers denoted $S[i], 0 \leq i \leq 7$. The round function R (Figure 2) takes two additional 128-bit inputs X_0, X_1 and outputs $S' = R(S, X_0, X_1)$ defined as:

$$S'[0] = S[7] \oplus X_0 \qquad\qquad S'[4] = S[3] \oplus X_1$$
$$S'[1] = A(S[0]) \oplus S[7] \qquad\qquad S'[5] = A(S[4]) \oplus S[3]$$
$$S'[2] = S[1] \oplus S[6] \qquad\qquad S'[6] = A(S[5]) \oplus S[4]$$
$$S'[3] = A(S[2]) \oplus S[1] \qquad\qquad S[7] = S[0] \oplus S[6]$$

Algorithm.   The specification that we give here is from the latest version (2023-03-16) of the ePrint report [SLN+22]. After the publication of the conference version [SLN+21] and subsequent third-party cryptanalysis [HII+22], the authors added a key feedforward in the initialization phase to make it non-invertible, which was not present in the conference version.

The key is divided into two 128-bit key blocks K_0, K_1. The scheme also uses a pair of constants Z_0, Z_1. Rocca (without AD) runs as follows:

- **Initialization phase:** the state S is initialized using the nonce and key. Then, 20 rounds $R(S, Z_0, Z_1)$ are applied. Then, K_0, K_1 are XORed to $S[0], S[4]$ respectively.

- **Encryption:** message blocks are encrypted by pairs (M_i, M_i') into pairs of ciphertexts (C_i, C_i'). For all i from 0 to $m - 1$:

$$\begin{cases} C_i = A(S[1]) \oplus S[5] \oplus M_i \\ C_i' = A(S[0] \oplus S[4]) \oplus S[2] \oplus M_i' \\ S \leftarrow R(S, M_i, M_i') \end{cases}$$

- **Finalization:** the state is updated 20 times using $R(S, |AD|, |M|)$, where $|AD|$ and $|M|$ are the respective lengths of the AD and message, and the 128-bit tag is computed as the XOR of all state registers.

Figure 3: Round function of Rocca-S.

Classical Security. The authors of Rocca claimed 128-bit security against forgery attacks and 256-bit security against key-recovery attacks. Importantly, they did not make any claims in the nonce-misuse setting.

In [HII+22], Hosoyamada et al. introduced a nonce-misuse attack that could recover the internal state using only *one* nonce-repeated pair. It follows a strategy of introducing a difference in certain message blocks, in order to observe some output differences, and solving the obtained equations to recover state values. Since the finalization function is key-less, recovering the state allows to create forgeries.

They then observed that one could turn this attack into a nonce-respecting one, by making decryption queries (which are authorized to repeat the nonces). After making a first nonce-respecting query to the encryption oracle, the adversary introduces a difference in the obtained ciphertext and tries to decrypt by trying all possible tags. If the number of decryption queries is not limited, this will eventually succeed after 2^{128} such queries, leading to a recovery of the state. In the first version of Rocca, where the initialization phase was invertible, the state recovery led to a key-recovery attack, breaking the claims. However, with the modified initialization, a recovery of the state does not lead to a recovery of the key.

Quantum Security. The authors of Rocca made no claim against Q2 attacks. Anand and Isobe studied specifically the quantum security of Rocca [AI23] and found a forgery attack that requires 2^{75} superposition queries. This attack is nonce-respecting and makes Q2 decryption queries.

2.4 Rocca-S

Rocca-S is a new version of Rocca which was proposed for standardization by the IETF [NFI]. We refer to the version of the draft standard which is the latest one at the time of writing (published March 2^{nd}, 2023).

Round Function. The internal state of Rocca-S is made of 7 registers of 128 bits. The round function $S' = R(S, X_0, X_1)$ (Figure 3) updates this state as follows:

$$S'[0] = S[6] \oplus S[1] \qquad\qquad S'[4] = A(S[3]) \oplus X_1$$
$$S'[1] = A(S[0]) \oplus X_0 \qquad\qquad S'[5] = A(S[4]) \oplus S[3]$$
$$S'[2] = A(S[1]) \oplus S[0] \qquad\qquad S'[6] = A(S[5]) \oplus S[4]$$
$$S'[3] = A(S[2]) \oplus S[6]$$

Algorithm. The algorithm (without AD) runs as follows:

- **Initialization:** after loading the key K_0, K_1 and nonce, 16 rounds of $R(S, Z_0, Z_1)$ are applied, followed by a key addition in all state registers.

- **Encryption:** for all $i = 0$ to $m - 1$:

$$\begin{cases} C_i = A(S[3] \oplus S[5]) \oplus S[0] \oplus M_i \\ C_i' = A(S[4] \oplus S[6]) \oplus S[2] \oplus M_i' \\ S \leftarrow R(S, M_i, M_i') \end{cases} \quad (3)$$

- **Finalization:** the round function $R(S, |AD|, |M|)$ is iterated 16 times. Then, the 256-bit tag is computed as:

$$T = (S[0] \oplus S[1] \oplus S[2] \oplus S[3]) \| (S[4] \oplus S[5] \oplus S[6]) \ . \quad (4)$$

Security. The increased tag size allows the authors of Rocca-S to claim 256 bits of security against forgery, state and key-recovery attacks (nonce-respecting). In the quantum setting, they claim 128 bits of security against nonce-respecting forgery and key-recovery attacks. However, like Rocca, they did not consider attacks in the Q2 setting and did not make security claims in this model. To the best of our knowledge, Rocca-S remains secure in the Q1 setting.

3 Tools

In this section we give the main algorithmic tools of our attacks. These tools are gathered from previous works in quantum cryptanalysis [BBC+21] and quantum computing [ORR13] and adapted here to our setting. To the best of our knowledge, the case of "smaller correlation" (Theorem 2) is new.

We assume basic knowledge of the quantum circuit model [NC02] (Toffoli / CNOT / Hadamard gates, ket $|\cdot\rangle$ notations). As is commonly done in previous works [Bon17, CHLS20, KLLN16a], we query AE schemes using a *standard oracle*. However, in our main quantum algorithm, we need a *phase oracle*.

Definition 1 (Standard oracle). For $f : \{0,1\}^n \to \{0,1\}^m$, a *standard oracle* for f is a quantum circuit O_f that maps $|x\rangle |y\rangle$ to $|x\rangle |y \oplus f(x)\rangle$.

Definition 2 (Phase oracle). For $f : \{0,1\}^n \to \{0,1\}^m$, a *phase oracle* for f is a quantum circuit O_f' that maps $|x\rangle |y\rangle$ to $(-1)^{y \cdot f(x)} |x\rangle |y\rangle$.

Both oracles are equivalent by composing with a Hadamard transform. Also, if one knows a classical circuit that implements f, both oracles are easy to construct.

These AE schemes are nonce-based. While the nonce can be chosen by the adversary, it cannot be repeated between two queries. Since Q2 queries are merely an extension of classical queries, the same can be said in the quantum setting. Therefore, we impose that *each of the Q2 queries is answered using a different, classical nonce.* Using a classical nonce or randomness is common in proofs of quantum security for encryption and AE modes [ATTU16, BBC+21].

That is, the adversary has access to a family of oracles: $O_{N,m}$ for different nonces N and message lengths m (we assume that the AD is empty), and they cannot make two queries with the same nonce.

Each oracle encrypts several (pairs of) message blocks (M_i, M_i'), depending on the selected length, and returns the corresponding (pairs of) ciphertexts (C_i, C_i'), and the tag:

$$\begin{aligned} O_{N,m} \ : \ & |M_0, M_0', \dots, M_{m-1}, M_{m-1}'\rangle \ |y_0, y_0', \dots, y_{m-1}, y_{m-1}', y\rangle \\ & \mapsto |M_0, M_0', \dots, M_{m-1}, M_{m-1}'\rangle \\ & \quad |y_0 \oplus C_0, y_0' \oplus C_0', \dots, y_{m-1} \oplus C_{m-1}, y_{m-1}' \oplus C_{m-1}', y \oplus T\rangle \ . \end{aligned}$$

Quantum Search.   Grover's exhaustive search algorithm [Gro96] is a procedure to find a "good" element in a search space of size 2^n in $\left\lfloor \frac{\pi}{4} 2^{n/2} \right\rfloor$ iterates; each iterate queries a phase oracle that flips only the phase of this good element. *Amplitude amplification* [BHMT02] generalizes this to any algorithm $\mathcal{A}$ (even a quantum algorithm) that outputs a good element with probability p. It then makes about $\frac{\pi}{4} \frac{1}{\sqrt{p}}$ iterates, with two calls to $\mathcal{A}$ and one query to the oracle per iterate, to succeed with overwhelming probability.

Grover Search Cost Estimates.   All AE schemes studied in this paper are based on the AES round function. Quantum attacks on them require to implement AES components. Since the scope of this paper is only to demonstrate the existence of attacks, we will use approximate quantum gate and query counts (by a factor 2 at best). For example, Table 4 in [JBS$^+$22] gives a count $12240 = 2^{13.58}$ Toffoli gates for a full (10-round) AES-128. We use this to assume that a single round of AES can be implemented with 2^{10} Toffoli gates (we focus only on Toffoli counts for simplicity).

For all four schemes, implementing Grover's exhaustive key search requires to recompute the initialization of the scheme. The number of iterates depends on the key size (128 or 256 bits) and the cost of the Grover iterate is dominated by this initialization function which, in turn, can be estimated using the number of AES rounds it contains. These estimates are summarized in Table 3.

At some point in our algorithms, we also need to solve AES S-Box differential equations of the form: $S(x \oplus \Delta) \oplus S(x) = \Delta'$. This can be done using a small Grover search on x, costing 2^6 S-Boxes at most, i.e., 4 rounds of AES, or 2^{12} Toffoli gates.

Table 3: Toffoli count of Grover's key search for studied schemes. As the exponent is rounded to the nearest integer, the Toffoli gate counts for some schemes can appear identical even though the number of AES rounds necessary to compute an output differs between them.

Cipher	Key length	AES rounds	Toffoli count
Rocca	256	$4 \times 20 = 80$	2^{145}
Rocca-S	256	$6 \times 16 = 96$	2^{145}
Tiaoxin	128	$6 \times 15 = 90$	2^{81}
AEGIS-128L	128	$8 \times 10 = 80$	2^{81}

Toffoli Counts of Arithmetic Operations.   Since the complexities of our attacks will be clearly below those of exhaustive search, we give only imprecise upper bounds on the cost of quantum circuits for arithmetic operations. Using the addition circuit of [CDKM04], an n-bit addition costs $2n$ Toffoli gates, and a controlled variant can be implemented with $4n$ Toffoli gates. Using a simple implementation as a series of controlled additions, an n-bit product can be implemented with $4n^2$ Toffoli gates. A table lookup circuit, implementing $|i\rangle |0\rangle \mapsto |i\rangle |c_i\rangle$ where the c_i are classically stored values, takes $4 \times 2^n \times m$ Toffoli and CNOT gates when c_i is on m bits and i is on n bits. Finally, a Euclidean division of an n-bit integer by an m-bit integer costs about $4nm$ Toffoli gates using a sequence of n conditional subtractions.

3.1   Linear Post-processing

The generic approach to post-process the output of an oracle requires two identical calls, due to the reversibility of quantum computations. This is not doable in our case since we only query the oracle once. Fortunately, truncations [HS18] and more generally linear functions [BBC$^+$21] can be computed from a single call.

For our purposes, we need to separate the linear function in two parts, one of which goes directly into the phase. This can be obtained using [BBC⁺21, Lemma 2] as a black-box, but we give the whole proof (adapted directly from [BBC⁺21]) to be self-contained.

Lemma 1 (Extended linear post-processing, adapted from [BBC⁺21]). *Let $f : \{0,1\}^n \to \{0,1\}^m$ be a function, let O_f be a standard oracle for f: $|x\rangle |0\rangle \mapsto |x\rangle |f(x)\rangle$. Let $g : \{0,1\}^m \to \{0,1\}^\ell$ and $h : \{0,1\}^m \to \{0,1\}$ be two linear functions, i.e., $\forall x, y, g(x \oplus y) = g(x) \oplus g(y), h(x \oplus y) = h(x) \oplus h(y)$, with standard oracles O_g and O_h. Then there exists a circuit implementing the operator:*

$$|x\rangle |y\rangle \mapsto (-1)^{h(f(x))} |x\rangle |y \oplus g(f(x))\rangle \quad ,$$

which makes a single query to O_f, two queries to O_g and two queries to O_h.

Proof. On input $|x\rangle |y\rangle$, create the uniform superposition over outputs z and append a qubit in the state $H |1\rangle = \frac{1}{\sqrt{2}} (|0\rangle - |1\rangle)$:

$$|x\rangle |y\rangle \frac{1}{2^{m/2}} \sum_{z \in \mathbb{F}_2^m} |z\rangle (H |1\rangle)$$

Compute O_h with register z as input and the last qubit as output; compute O_g with register z as input and y as output:

$$|x\rangle \frac{1}{2^{m/2}} \sum_{z \in \mathbb{F}_2^m} |y \oplus g(z)\rangle |z\rangle (-1)^{h(z)} (H |1\rangle)$$

Notice that the result of h appears in the phase now, because we used $H |1\rangle$ as its output register. Now, apply O_f with register x as input and z as output:

$$|x\rangle \frac{1}{2^{m/2}} \sum_{z \in \mathbb{F}_2^m} (-1)^{h(z)} |y \oplus g(z)\rangle |z \oplus f(x)\rangle (H |1\rangle)$$

Redo the computations of O_h and O_g:

$$|x\rangle \frac{1}{2^{m/2}} \sum_{z \in \mathbb{F}_2^m} (-1)^{h(z)+h(z \oplus f(x))} |y \oplus g(z) \oplus g(z \oplus f(x))\rangle |z \oplus f(x)\rangle (H |1\rangle)$$

Erase the qubit $(H |1\rangle)$ and use the linearity of g, h to rewrite:

$$|x\rangle \frac{1}{2^{m/2}} \sum_{z \in \mathbb{F}_2^m} (-1)^{h(f(x))} |y \oplus g(f(x))\rangle |z \oplus f(x)\rangle$$

$$= (-1)^{h(f(x))} |x\rangle |y \oplus g(f(x))\rangle \frac{1}{2^{m/2}} \sum_{z \in \mathbb{F}_2^m} |z \oplus f(x)\rangle \quad .$$

The last register becomes disentangled and always contains a uniform superposition over $\{0,1\}^m$, which we can erase, leading to the result. □

In particular, we can truncate the output of a stream cipher and separate it in two parts, one that remains in the computational basis state, and one that goes into the phase.

3.2   Properties of the Walsh Transform

Let $f : \{0,1\}^n \to \{-1,1\}$ be a function. The Walsh-Hadamard transform of f is defined as: $\widehat{f}(y) = \sum_{x \in \{0,1\}^n} (-1)^{x \cdot y} f(x)$. It corresponds to the Fourier transform in the group $\mathbb{Z}_2^n$. The quantum Hadamard transform $H^{\otimes n}$ computes a (normalized) Walsh-Hadamard transform on the amplitudes of its n-qubit input state. That is:

$$\frac{1}{2^{n/2}} \sum_x f(x) |x\rangle \xrightarrow{H^{\otimes n}} \frac{1}{2^n} \sum_y \widehat{f}(y) |y\rangle \quad .$$

In the following, we need the following important properties of the Walsh transform.

Proposition 1.   *1. (Shift) Let $s \in \{0,1\}^n$ be a constant and $g : x \mapsto f(x \oplus s)$. Then for all x, $\widehat{g}(x) = (-1)^{x \cdot s} \widehat{f}(x)$.*

2. (Convolution theorem) Let $f, g : \{0,1\}^n \to \{-1,1\}$. Then for all x, $\widehat{\widehat{f}\widehat{g}}(x) = 2^n \sum_y f(y) g(x \oplus y)$.

3. (Product) Let $g_1, \ldots, g_n$ be functions of codomain $\{-1,1\}$. Then the Walsh transform of $(x_1, \ldots, x_n) \mapsto \prod_i g_i(x_i)$ is $(x_1, \ldots, x_n) \mapsto \prod_i \widehat{g}_i(x_i)$.

3.3   Amplitude Transduction

Quantum rejection sampling is the process of transforming a quantum state into another one, by modifying its amplitudes – in a way similar to classical rejection sampling which transforms probability distributions.

Suppose that we have a quantum state of the form: $\sum_x u_x |x\rangle |\alpha_x\rangle$, where $0 \le \alpha_x < 2^n$ is an integer. (Therefore $|\alpha_x\rangle$ is indeed a basis state). We want to transform this into a state $\sum_x u_x \frac{\alpha_x}{2^n} |x\rangle |\alpha_x\rangle$, i.e., move α_x into the amplitude (up to renormalization).

A typical way to do this is to append a qubit register starting in state $|0\rangle$, which is transformed into a superposition of the form: $\frac{\alpha_x}{2^n} |0\rangle + |\psi_{\alpha_x}\rangle$ where $|\psi\rangle$ is a superposition of non-zero basis states. This step is called *amplitude transduction*. Then, the state becomes: $\sum_x u_x |x\rangle |\alpha_x\rangle \left(\frac{\alpha_x}{2^n} |0\rangle + |\psi_{\alpha_x}\rangle \right)$. Measuring $|0\rangle$ in the last register collapses the state on the wanted superposition.

We use the amplitude transduction algorithm of [SLSB19].

Lemma 2 ([SLSB19]). *There exists a quantum algorithm that, on input $|\alpha\rangle |0^{n+1}\rangle$ where $0 \le \alpha < 2^n$ is an integer, returns $|\alpha\rangle \left(\frac{\alpha}{2^n} |0^{n+1}\rangle + |\psi_\alpha\rangle \right)$ where $|\psi_\alpha\rangle$ is a superposition of non-zero basis states. This algorithm uses $\mathcal{O}(n)$ basis gates.*

Proof. The algorithm runs as follows. First, we apply a Hadamard transform on n qubits:

$$|\alpha\rangle \left(\frac{1}{2^{n/2}} \sum_{0 \le y \le 2^n - 1} |y\rangle \right) |0\rangle \quad .$$

We perform a comparison between y and α, which costs $\mathcal{O}(n)$ gates, and write the result in the last qubit:

$$|\alpha\rangle \frac{1}{2^{n/2}} \left(\sum_{0 \le y < \alpha} |y\rangle |0\rangle + \sum_{\alpha \le y < 2^n} |y\rangle |1\rangle \right) \quad .$$

We apply a Hadamard transform on the register holding y, obtaining:

$$|\alpha\rangle \frac{1}{2^n} \left(\sum_{0 \leq y < \alpha} \sum_z (-1)^{y \cdot z} |z\rangle |0\rangle + \sum_{\alpha \leq y < 2^n} \sum_z (-1)^{y \cdot z} |z\rangle |1\rangle \right)$$

$$= |\alpha\rangle \left(\frac{\alpha}{2^n} |0^n\rangle |0\rangle + \frac{1}{2^n} \sum_{z \neq 0} |z\rangle \left(\left(\sum_{0 \leq y < \alpha} (-1)^{y \cdot z} \right) |0\rangle + \left(\sum_{\alpha \leq y < 2^n} (-1)^{y \cdot z} \right) |1\rangle \right) \right) ,$$

where we recover a state with the form claimed. The exact form of $|\psi_\alpha\rangle$ depends only on the value of α and is not relevant for the rest of our study. $\qquad\square$

Approximation. In the context of this paper, the value $\frac{\alpha_x}{2^n}$ will be a fixed-point approximation of the amplitude that we actually want. Since the approximation error will be at most 2^{-n}, if n is large enough, the resulting quantum state will be close to our target state, and the algorithm will run without failure.

In particular, let α_x' be the "exact" amplitude and assume that our approximation satisfies: $\alpha_x = \lfloor 2^n \alpha_x' \rfloor \implies 2^n \alpha_x' - 1 < \alpha_x \leq 2^n \alpha_x'$. Let $|\psi_i\rangle$ be the "ideal" state after transduction success and $|\psi_r\rangle$ the "real" one, respectively:

$$|\psi_i\rangle := \frac{1}{\sqrt{\sum_x |\alpha_x u_x 2^{-n}|^2}} \sum_x \alpha_x 2^{-n} u_x |x\rangle , \quad |\psi_r\rangle := \frac{1}{\sqrt{\sum_x |\alpha_x' u_x|^2}} \sum_x \alpha_x' u_x |x\rangle . \quad (5)$$

Here $\sum_x |\alpha_x u_x 2^{-n}|^2 \leq \sum_x |\alpha_x' u_x|^2$. The Euclidean distance between them can be bounded as follows:

$$\| |\psi_i\rangle - |\psi_r\rangle \|^2 = 2 - 2 \langle \psi_i | \psi_r \rangle$$

$$= 2 - 2 \frac{1}{\sqrt{\sum_x |\alpha_x u_x 2^{-n}|^2} \sqrt{\sum_x |\alpha_x' u_x|^2}} \sum_x \alpha_x 2^{-n} |u_x|^2 \alpha_x'$$

$$\leq 2 - 2 \frac{\sum_x \alpha_x 2^{-n} \alpha_x' |u_x|^2}{\sum_x |\alpha_x' u_x|^2} \text{ using } \alpha_x \leq 2^n \alpha_x'$$

$$\leq \frac{\sum_x 2^{-n+1} \alpha_x' |u_x|^2}{\sum_x |\alpha_x' u_x|^2} \text{ using } 2^n \alpha_x' - \alpha_x < 1$$

$$\leq 2^{-n+1} \frac{\sum_x |u_x|^2}{\sum_x |\alpha_x' u_x|^2} \leq 2^{-n+1} \frac{1}{\sum_x |\alpha_x' u_x|^2} .$$

Let $p = \sum_x |\alpha_x' u_x|^2$ be the "ideal" probability to succeed in the transduction. We have:

$$\| |\psi_i\rangle - |\psi_r\rangle \|^2 \leq 2^{-n+1}/p . \quad (6)$$

Consider an algorithm (e.g., Algorithm 1, that we will define later) that uses transduction once, succeeds here with probability p, does further operations, measures and succeeds with probability p_r' (resp. p_i'). By Lemma 3.6 in [BV97], the total variation distance between the two probability distributions resulting from the "ideal" and "real" states is at most $4\| |\psi_i\rangle - |\psi_r\rangle \|$. Consequently:

$$p_r' \geq p_i' - 4\| |\psi_i\rangle - |\psi_r\rangle \| \geq p_i' - 4\sqrt{2^{-n+1}/p} . \quad (7)$$

If we ensure $p \gg 2^{-n}$, it is enough to study the approximated version. This will be the case in the attacks studied in this paper, as we typically use more than 300 bits of precision to approximate the amplitudes, while the success probability is bigger than 2^{-50}. More generally, while increasing this precision may require more costly arithmetic circuits, we haven't encountered a case where this limits the attacks.

3.4   Quantum Hidden Shift Algorithm with a Single Query

We want to solve the following problem.

Problem 1 (Hidden shift)**.** *Let* $g : \{0,1\}^n \to \{-1,1\}$ *be a function that can be computed in polynomial time. Given access to a quantum oracle for* $f : x \mapsto g(x \oplus s)$*, where* s *is a secret value, find* s*.*

The algorithm that we present here (Algorithm 1) is from [ORR13], and uses quantum rejection sampling. Several special cases have appeared before in cryptanalysis: for example, shifted multiplicative characters [vDHI06] and bent functions [Röt10]. In both cases, the algorithm avoids the rejection sampling by considering a situation in which the Fourier transform of the shifted function is easy to compute: in the former case, it's a multiplicative character, and in the latter, a constant.

In our case, we are interested in the probability to succeed after making a single phase query to the function f.

Ideas of Algorithm 1.   The first step is to query f and to perform a Hadamard transform. This places the Walsh coefficients of f into the amplitudes of the state. Next, we remark that by Proposition 1, these coefficients are actually those of g, multiplied by $(-1)^{x \cdot s}$. If we had the state $\frac{1}{2^{n/2}} \sum_x (-1)^{x \cdot s} |x\rangle$, we could immediately do a Hadamard transform and obtain $|s\rangle$. The Walsh coefficients of g prevent us to do that.

Thus, the next step is to *correct* the amplitudes by multiplying them by $1/\widehat{g}(x)$, using amplitude transduction (Subsection 3.3). Ideally, we would obtain the wanted state $\sum (-1)^{x \cdot s} |x\rangle$. However, the product operation is not possible if $\widehat{g}(x) = 0$. Furthermore, if the smallest values of $\widehat{g}(x)$ are very small compared to the average, the probability to measure 0 (and succeed) gets smaller. Thus, the best strategy, as suggested in [ORR13], is to dismiss the small Walsh coefficients of g. We introduce a bound M in the algorithm and only multiply by $\frac{M}{\widehat{g}(x)}$ if $|\widehat{g}(x)| \geq M$, and otherwise, by 0 (meaning that we eliminate the coordinate).

Theorem 1. *Let* M *be a bound and* $G = \#\{x, |\widehat{g}(x)| \geq M\}$*. Define* $p := \frac{GM^2}{2^{2n}}$ *and* $p' := \frac{G}{2^n}$*. Then:* • *the probability to measure* 0 *in Step 11 of Algorithm 1 is bigger than* $p - 2^{-n/2+1}$*;* • *the probability to measure* s *at the end of the algorithm is bigger than* $pp' - 2^{-n/2+1}$*. The algorithm makes one phase query to* f*, two computations of* $\widehat{g}$*, and uses* $\mathcal{O}(n^2)$ *additional Toffoli gates.*

Proof. Following Algorithm 1 until Step 5, we obtain the state:

$$\frac{1}{2^n} \sum_x (-1)^{x \cdot s} \widehat{g}(x) |x\rangle |\alpha_x\rangle .$$

Here $\alpha_x = \lfloor 2^n M/|\widehat{g}|(x) \rfloor$ if $|\widehat{g}(x)| \geq M$ and 0 otherwise. Recall that $|\widehat{g}(x)|$ is, by definition, an integer between 0 and 2^n. We first compute it, then compare the result with M, and perform a Euclidean division of $2^n M$ (a constant) by $|\widehat{g}(x)|$. This costs $\mathcal{O}(n^2)$ gates.

Since we have computed $\widehat{g}(x)$, we know its sign, and we can handle it immediately. We perform a controlled phase flip by $sgn(\widehat{g}(x))$, which will cancel the sign of $\widehat{g}(x)$ in the phase, obtaining the state:

$$\frac{1}{2^n} \sum_x (-1)^{x \cdot s} |\widehat{g}(x)| |x\rangle |\alpha_x\rangle .$$

The next steps realize amplitude transduction following Lemma 2. In the $n + 1$-qubit ancillary register, the amplitude on 0^{n+1} is equal to $\frac{\alpha_x}{2^n}$, for all x. This includes the cases

where $\alpha_x = 0$, where there is simply no amplitude on 0^{n+1}. Therefore, the probability to measure 0^{n+1} at Step 10 is equal to:

$$\sum_x \left(\frac{1}{2^n} (-1)^{x \cdot s} |\widehat{g}(x)| \frac{\alpha_x}{2^n} \right)^2 \geq \frac{1}{2^{2n}} \sum_{x, |\widehat{g}(x)| \geq M} \left(|\widehat{g}(x)| \left(\frac{M}{|\widehat{g}(x)|} - 2^{-n} \right) \right)^2$$

$$\geq \frac{1}{2^{2n}} \sum_{x, |\widehat{g}(x)| \geq M} \frac{M^2}{|\widehat{g}(x)|^2} |\widehat{g}(x)|^2 - \frac{1}{2^{2n}} \sum_{x, |\widehat{g}(x)| \geq M} |\widehat{g}(x)| M 2^{-n+1} \ .$$

While the first term is equal to $\frac{GM^2}{2^{2n}} = p$, the second can be shown to be negligible. First, we have $|\widehat{g}(x)| \leq 2^n$ by definition, so the sum can be bounded as:

$$\frac{1}{2^{2n}} \sum_{x, |\widehat{g}(x)| \geq M} |\widehat{g}(x)| M 2^{-n+1} \leq \frac{MG}{2^{2n-1}} \ .$$

Furthermore, we know that $G \leq 2^n$ and:

$$GM^2 \leq \sum_{x, |\widehat{g}(x)| \geq M} \widehat{g}(x)^2 \leq \sum_x \widehat{g}(x)^2 = 2^n \sum_x g(x)^2 = 2^{2n} \ .$$

Consequently, $M^2 G^2 \leq 2^{3n} \implies MG \leq 2^{3n/2}$ which bounds the second term by $2^{-n/2+1}$.

Assuming that we succeeded at Step 11 (i.e., we measure 0^{n+1}), the state collapses and becomes proportional to $\sum_x (-1)^{x \cdot s} |\widehat{g}(x)| \alpha_x |x\rangle$.

Following the discussion in Subsection 3.3, the state is close to:

$$\frac{1}{\sqrt{G}} \sum_{x, |\widehat{g}(x)| \geq M} (-1)^{x \cdot s} |x\rangle \ ,$$

as long as $p \gg 2^{-n}$ (which is the case here since $p > 2^{-n/2}$). We then apply H:

$$\frac{1}{\sqrt{2^n G}} \sum_y \sum_{x, |\widehat{g}(x)| \geq M} (-1)^{x \cdot s + x \cdot y} |y\rangle \ .$$

Afterwards, the probability to measure $y = s$ is:

$$p' := \frac{1}{2^n G} \times G^2 = \frac{G}{2^n} = \Pr_x(|\widehat{g}(x)| \geq M) \ . \tag{8}$$

All in all, the total probability to succeed is: $pp' - 2^{-n/2+1}$ where $pp' = M^2 G^2 2^{-3n}$, completing the proof. $\qquad\square$

Remark 1. The condition $p \gg 2^{-n/2}$ might appear as a strong limitation of Theorem 1. However, the values of n encountered in this paper range from 384 to 640, since we are recovering large hidden shifts, while the mere condition of having a valid attack imposes us $p > 2^{-128}$.

In order to use Theorem 1, we need an efficient algorithm to compute $\widehat{g}$. In order to maximize the success probability, we need to know the distribution of the Walsh coefficients to choose M appropriately. Both will be possible in the cases we are interested in, because g will be the product of many small-range independent functions. Then $\widehat{g}$ is easy to compute by taking the product of Walsh coefficients (see Proposition 1).

Remark 2 (Global phase). If we have access to $\pm g(x \oplus s)$, where the leading sign is not known, it turns into a global phase that is irrelevant for the algorithm. At the final step, we will still measure s.

Algorithm 1 Quantum hidden shift with rejection sampling and the technique of [SLSB19].

Input: Quantum access to $f(x) = g(x \oplus s)$ for a known g, a bound M
Output: s, with probability pp'

1: Start from n qubits initialized to 0 $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad \triangleright |0^n\rangle$
2: Apply $H^{\otimes n}$ $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad \triangleright \frac{1}{2^{n/2}} \sum_x |x\rangle$
3: Query f in the phase $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad \triangleright \frac{1}{2^{n/2}} \sum_x f(x) |x\rangle$
4: Apply $H^{\otimes n}$ $\qquad\qquad\qquad \triangleright \frac{1}{2^n} \sum_x \widehat{f}(x) |x\rangle = \frac{1}{2^n} \sum_x (-1)^{x \cdot s} \widehat{g}(x) |x\rangle$
5: Compute the amplitude multiplier:

$$
\alpha_x := \begin{cases} \lfloor 2^n M / |\widehat{g}(x)| \rfloor & \text{if } |\widehat{g}(x)| \geq M \\ 0 & \text{otherwise} \end{cases} \tag{9}
$$

$\qquad$ where $0 \leq \alpha_x \leq 1$ in an additional register $\qquad \triangleright \frac{1}{2^n} \sum_x (-1)^{x \cdot s} \widehat{g}(x) |x\rangle |\alpha_x\rangle$
6: Compute the sign $sgn(\widehat{g}(x))$ in the phase $\qquad \triangleright \frac{1}{2^n} \sum_x (-1)^{x \cdot s} |\widehat{g}(x)| |x\rangle |\alpha_x\rangle$
7: Append an ancilla register $|0^n\rangle$ and apply $H^{\otimes n}$ on it

$$
\frac{1}{2^n} \sum_x (-1)^{x \cdot s} |\widehat{g}(x)| |x\rangle |\alpha_x\rangle \frac{1}{2^{n/2}} \sum_y |y\rangle \tag{10}
$$

8: Perform a comparison between y and α_x and store the result in a new ancilla qubit

$$
\frac{1}{2^n} \sum_x (-1)^{x \cdot s} |\widehat{g}(x)| |x\rangle |\alpha_x\rangle \frac{1}{2^{n/2}} \left(\sum_{0 \leq y < \alpha_x} |y\rangle |0\rangle + \sum_{\alpha_x \leq y < 2^n} |y\rangle |1\rangle \right) \tag{11}
$$

9: Apply $H^{\otimes n}$ on the register holding y: the amplitude on the $|0^{n+1}\rangle$ component is $\frac{\alpha_x}{2^n}$
10: Erase $|\alpha_x\rangle$
11: Measure the last register. If the obtained value is different from 0^{n+1}, abort
12: Otherwise, the state has collapsed and is close to:

$$
\frac{1}{\sqrt{G}} \sum_{x, |\widehat{g}(x)| \geq M} (-1)^{x \cdot s} |x\rangle \tag{12}
$$

13: Apply $H^{\otimes n}$, measure and return the result.

Remark 3 (Self-correlation). A technique similar to this algorithm appeared also in [Sch23], where instead of dividing by the Walsh coefficient, one multiplies by it. This would compute the discrete convolution: $(f * g)(y) = \sum_x f(x \oplus y)g(x)$ in the amplitudes of the state, and lead to a similar result since $(f * g)(y)$ is greater for $y = s$. However the analysis when cutting off the small Walsh coefficients is more difficult, so we settled for the easier method.

Related Quantum algorithms. Problem 1 is very similar to Simon's problem [Sim97]. Indeed, it would be possible to solve it with Simon's algorithm, which is now a fairly standard approach in quantum cryptanalysis. The main issue is that Simon's algorithm requires $\mathcal{O}(n)$ queries to the function, while we can only afford one query to it. Hidden shift attacks can also refer to the approach of [BN18]. The problem there is slightly different, as the shift is with a modular addition and not an XOR. Moreover, as with Simon's approach, multiple queries must be performed.

3.5 Hidden Shift with Smaller Correlation

For the attack on AEGIS-128L (Subsection 4.5), we need to solve a more difficult variant of Problem 1, in which the function that we query is multiplied by a highly biased function h, which is unknown. We model this function as selected uniformly at random among Boolean functions of the same Hamming weight.

Problem 2 (Correlated hidden shift). *Let $g : \{0,1\}^n \to \{-1,1\}$ be a function that can be computed in polynomial time. Let $h : \{0,1\}^n \to \{-1,1\}$ be a function selected uniformly at random in $\mathcal{H}_c = \left\{ h \in \{0,1\}^n \to \{-1,1\} \big| \Pr_x[h(x) = -1] = \frac{1}{2}(1+c) \right\}$ for some $0 < c \le 1$. Given access to a quantum oracle for $f : x \mapsto h(x)g(x \oplus s)$, where s is a secret value; find s.*

It can be noticed that Problem 1 corresponds to the case $c = 1$. When c is 0, we cannot hope to recover the secret s as the function queried will be completely random. However, if c is large enough, we can still use Algorithm 1.

Theorem 2. *Consider the setting of Problem 2. On average over h, applying Algorithm 1 on f with g as the known function will recover s with probability greater than $pc^2p' - 2^{-n/2}$, where $pc^2 = \frac{M^2Gc^2}{2^{2n}}$ is the probability to measure 0 at Step 11 and $p' = \frac{G}{2^n}$ is the probability to succeed in the second step.*

Proof. By similar bounds as in the proof of Theorem 1, in the following we can assume that the quantum rejection sampling works exactly, by subtracting a term $2^{-n/2}$ in the probability of success.

Following Algorithm 1, the state after Step 10 is:

$$|\psi\rangle |1\rangle + \frac{1}{2^n} \sum_{x,|\widehat{g}(x)|\ge M} \frac{M}{\widehat{g}(x)} \widehat{f}(x) |x\rangle |0\rangle \quad .$$

If we postpone the measurement of Step 11 at the end of the algorithm, we have the state:

$$H |\psi\rangle |1\rangle + \frac{M}{2^{3n/2}} \sum_y \sum_{x,|\widehat{g}(x)|\ge M} (-1)^{x \cdot y} \frac{\widehat{f}(x)}{\widehat{g}(x)} |y\rangle |0\rangle \quad . \tag{13}$$

We will now estimate the amplitude of $|s\rangle |0\rangle$. We start by rewriting $\widehat{f}(x)$ using the convolution theorem:

$$\widehat{f}(x) = 2^{-n} \sum_z (-1)^{z \cdot s} \widehat{g}(z) \widehat{h}(x \oplus z) \quad .$$

Note that if h is constant and equal to 1, it has a single nonzero Walsh coefficient in 0 ($z = x$), equal to 2^n, and we recover the equality $\widehat{f}(x) = (-1)^{x \cdot s} \widehat{g}(x)$ (and the rest of the proof of Theorem 1). The amplitude is

$$\frac{M}{2^{5n/2}} \sum_{x, |\widehat{g}(x)| \geq M} (-1)^{x \cdot s} \frac{1}{\widehat{g}(x)} \left(\sum_z (-1)^{z \cdot s} \widehat{g}(z) \widehat{h}(x \oplus z) \right) . \tag{14}$$

We separate the term $z = x$ from the rest, noticing that $\widehat{h}(0) = \sum_x h(x) = 2^n c$ by our definition of c:

$$\frac{M}{2^{5n/2}} \left(\sum_{x, |\widehat{g}(x)| \geq M} (-1)^{x \cdot s} \frac{1}{\widehat{g}(x)} \left(\sum_{z \neq x} (-1)^{z \cdot s} \widehat{g}(z) \widehat{h}(x \oplus z) \right) + Gc2^n \right) . \tag{15}$$

where $G := \#\{x, |\widehat{g}(x)| \geq M\}$.

Now, we can use the fact that the terms that depend on h in this amplitude (and the Walsh-Hadamard transform) are linear in h, meaning that the average over h of this amplitude is the amplitude for the average function $h^*(x) = \frac{1}{|\mathcal{H}_c|} \sum_{h \in \mathcal{H}_c} h(x)$. Now, as $\mathcal{H}_c$ is a symmetric distribution over the input values, h^* will be a constant function. This means that for all $x \neq 0$, $\widehat{h^*}(x) = 0$. Thus, the average amplitude is simply the isolated part, $\frac{MGc}{2^{3n/2}}$.

Note that we need to estimate the probability, that is, the average of the square of the amplitude. We use the well-known fact that this is always greater than the square of the average (the gap between the two being the variance), and obtain that the probablity to measure $|s\rangle |0\rangle$ is, on average over h, greater than $\frac{(MGc)^2}{2^{3n}}$.

For the probability to measure 0 at Step 11, it has the expression:

$$\frac{M^2}{2^{4n}} \sum_{x, |\widehat{g}(x)| \geq M} \left| 2^n c + \frac{1}{\widehat{g}(x)} \sum_{z \neq x} (-1)^{z \cdot s} \widehat{g}(z) \widehat{h}(x \oplus z) \right|^2 . \tag{16}$$

We can use the same argument: the average over h is bigger than $\frac{GM^2 c^2}{2^{2n}}$.

Finally, taking into account the failure probability of rejection sampling, we obtain the desired probabilities. $\square$

4   Applications

Our attack combines Algorithm 1 with linear post-processing to recover the internal state. Recall that the nonce and key are fixed classical values, which means that after initialization, in all targeted designs, the internal state is a fixed value. We want to recover it (or part of it).

4.1   State-recovery on Rocca: Hidden Shifts

Assume that we encrypt a couple of fixed message blocks (e.g., 0), then the internal state S remains a fixed value. Our goal is to recover this S. We encrypt 5 pairs of message blocks in superposition and unroll several of the corresponding ciphertexts. Some ciphertexts are linear in the message, and thus directly give a constant that depend on the initial state (denoted by E_i). The important part of these ciphertexts are places that contain both a linear combination of the input messages (denoted by X_i) and a function of some state values (denoted by V_i).

$$C_0 = M_0 \oplus \underbrace{S[5] \oplus A(S[1])}_{:=E_0}$$

$$C_0' = M_0' \oplus \underbrace{S[2] \oplus A(S[0] \oplus S[4])}_{:=E_1}$$

$$C_1 = M_1 \oplus \underbrace{S[3] \oplus A(S[4]) \oplus A(S[7] \oplus A(S[0]))}_{:=E_2}$$

$$C_1' = M_1' \oplus S[1] \oplus S[6] \oplus A(\underbrace{M_0 \oplus M_0'}_{:=X_0} \oplus \underbrace{S[3] \oplus S[7]}_{:=V_0})$$

$$C_2' = M_2' \oplus S[4] \oplus S[7] \oplus A(S[0]) \oplus A(S[5])$$
$$\oplus A(\underbrace{M_1 \oplus M_1'}_{:=X_1} \oplus \underbrace{S[0] \oplus S[1] \oplus S[6] \oplus A(S[2])}_{:=V_1})$$

$$C_4 = M_4 \oplus S[0] \oplus S[6] \oplus A(\underbrace{M_0}_{:=X_2} \oplus \underbrace{S[7]}_{:=V_2}) \oplus A\left[\underbrace{M_2'}_{:=X_3} \oplus \underbrace{S[7] \oplus A(S[0]) \oplus A(S[1] \oplus S[6])}_{:=V_3}\right]$$

$$\oplus A\left[S[4] \oplus S[7] \oplus A(S[0]) \oplus A(S[5])\right] \oplus A\left[\underbrace{M_1 \oplus M_0'}_{:=X_4}\right.$$

$$\underbrace{\left. \oplus S[0] \oplus S[3] \oplus S[6] \oplus A(\underbrace{M_0 \oplus M_2}_{=0} \oplus S[4] \oplus S[7] \oplus A(S[5])) \oplus A(S[3] \oplus A(S[4]))\right]}_{:=V_4} .$$

Note that V_4 depends on $M_0 \oplus M_2$. As M_2 is a free variable (it is not involved in any X_i), we can ensure that $M_0 \oplus M_2 = 0$, so that V_4 only depends on initial state variables.

From these formulas, we can see that accessing C_0, C_0', C_1 directly gives E_0, E_1 and E_2, while the other ciphertexts are, up to constants and a plaintext block, sums of functions of the form $A(X_i \oplus V_i)$.

We now describe the quantum oracle we will construct from the query oracle. The inputs will be the 128-bit variables $X_0, \ldots, X_4$ and the message blocks for the query depend on them as follows (the others are simply put to 0):

$$M_0 = X_2, \qquad M_0' = X_0 \oplus X_2, \qquad M_1 = X_0 \oplus X_2 \oplus X_4$$
$$M_1' = X_0 \oplus X_2 \oplus X_4 \oplus X_1, \qquad M_2 = X_2, \qquad M_2' = X_3$$

From these equations, it is easy to see that we have

$$M_0 \oplus M_0' = X_0, \qquad M_1 \oplus M_1' = X_1, \qquad M_0 = X_2$$
$$M_2' = X_3, \qquad M_1 \oplus M_0' = X_4, \qquad M_0 \oplus M_2 = 0 ,$$

which is what we need.

Because the message blocks are either constant, or linear functions of the X_i variables, we can add M_i/M_i' to the ciphertexts C_i/C_i'. Next, we use a linear post-processing (Lemma 1) in order to construct the following oracle:

$$|X_0, \ldots, X_4\rangle |0\rangle$$
$$\mapsto |X_0, \ldots, X_4\rangle \underbrace{|C_0 \oplus M_0}_{E_0}, \underbrace{C_0' \oplus M_0'}_{E_1}, \underbrace{C_1 \oplus M_1\rangle}_{E_2} (-1)^{F(C_1' \oplus M_1', C_2' \oplus M_2', C_4 \oplus M_4)} ,$$

where F is a linear function.

Remark 4. The E_i are values that are also available in classical attack scenarios. The quantum advantage comes from the ability to retrieve the V_i to recover the state.

Table 4: Number of occurrences of Walsh coefficients with given absolute value for a column of the AES S-Box's LAT.

|LAT coefficient|	0	2^2	2×2^2	3×2^2	4×2^2	5×2^2	6×2^2	7×2^2	8×2^2
Occurences	17	48	36	40	34	24	36	16	5

Hidden Shift Problem. Now we define the function F. Recall that A is a single AES round, of the form: $A = \mathsf{MC} \circ \mathsf{SR} \circ \mathsf{SB}$. In order to transform the output into a single bit, we will take the dot-product with an appropriate 128-bit mask; we construct this mask with 16 copies of a single 8-bit mask β.

From now on, we choose an arbitrary β. While there is no particular constraint in the case of Rocca, the choice of the mask is more important in the cases of Tiaoxin and AEGIS (it will also be different for them).

On input a 128-bit AES state $Z = (z_0, \ldots, z_{15})$, we define the function: $L(Z) = \sum_i \beta \cdot z_i$. Then, the function F is simply: $F(Z_1, Z_2, Z_3) = \bigoplus_i L \circ \mathsf{MC}^{-1}(Z_i)$. In other words, it removes the last MC layer, uses a linear mask on each byte and XORs them all. By definition:

$$F(C_1' \oplus M_1', C_2' \oplus M_2', C_4 \oplus M_4) = L \circ \mathsf{MC}^{-1}(S[1] \oplus S[6]) \oplus L \circ \mathsf{MC}^{-1} \circ A(X_0 \oplus V_0)$$
$$\oplus\, L \circ \mathsf{MC}^{-1}(A(S[0]) \oplus A(S[5]) \oplus S[4] \oplus S[7]) \oplus L \circ \mathsf{MC}^{-1} \circ A(X_1 \oplus V_1)$$
$$\oplus\, L \circ \mathsf{MC}^{-1}(S[0] \oplus S[6] \oplus A(S[4] \oplus S[7]) \oplus A(S[0]) \oplus A(S[5]))$$
$$\oplus\, L \circ \mathsf{MC}^{-1} \circ A(X_2 \oplus V_2) \oplus L \circ \mathsf{MC}^{-1} \circ A(X_3 \oplus V_3) \oplus L \circ \mathsf{MC}^{-1} \circ A(X_4 \oplus V_4) \ .$$

Notice that $L \circ \mathsf{MC}^{-1} \circ A(X) = L \circ \mathsf{SB}(X)$ since L is invariant by permutation of the bytes. Next, we define the functions g and f:

$$\begin{cases} g(X_0, \ldots, X_4) := (-1)^{\sum_{i<5} L(\mathsf{SB}(X_i))} \\ f(X_0, \ldots, X_4) := (-1)^{F(C_1' \oplus M_1', C_2' \oplus M_2', C_4 \oplus M_4)} = \pm g(X_0 \oplus V_0, \ldots, X_4 \oplus V_4) \ , \end{cases}$$

where f has a leading unknown bit depending on the constant terms.

We will now retrieve the hidden shift $V_0, \ldots, V_4$ using Algorithm 1. We rename the individual bytes of $X_0, \ldots, X_4$ as $x_0, \ldots, x_{79}$ and rewrite g as:

$$g(x_0, \ldots, x_{79}) = \prod_{i=0}^{79} (-1)^{\beta \cdot \mathsf{SBox}(x_i)} \ . \tag{17}$$

In particular, f is still a shifted version of this function. Now, to bound the runtime and success probability of Algorithm 1, we need to analyze the Walsh coefficients of g.

Analysis of $\widehat{g}$. Since g is the product of 80 individual functions of one byte: $g_\beta : x \mapsto (-1)^{\beta \cdot \mathsf{SBox}(x)}$, we can use Proposition 1 and compute $\widehat{g}$ as a product of $\widehat{g_\beta}$. By definition, $\widehat{g_\beta}(x)$ is, up to a constant, the coefficient at column β and row x in the Linear Approximation Table of the SBox. Thus, $\widehat{g_\beta}$ corresponds to one column of the LAT. Moreover, we are interested only in the distribution of Walsh coefficients, and for the AES SBox, all non-zero columns are equivalent. Thus, any non-zero mask β gives the same result. The distribution is given in Table 4.

It could be a priori difficult to compute the Walsh spectrum of g, since it has a 640-bit input. However, by representing the distribution of Walsh coefficients as a table like in Table 4, we can compute the exact distribution, which is actually quite sparse. For 80 S-Boxes, the table contains approximately 7.5 million non-zero coefficients.

To run Algorithm 1, we need to select the threshold M maximizing the success probability. Recall that it is the product pp', where $G := \#\{x, |\widehat{g}(x)| \geq M\}$, $p = \frac{M^2}{2^{2n}}G$ is the success in the first step (which we can detect) and $p' = \frac{G}{2^n}$ is the success in the second step. Since we know the entire Walsh spectrum of g, we select M to maximize $pp' = M^2 G^2 2^{-3n}$:

$$M := 2^{326.23}, \quad G = 2^{610.60}, \quad p = 2^{-16.94}, \quad p' = 2^{-29.40}, \quad pp' = 2^{-46.34}.$$

These parameters will minimize the query complexity of the attack, however they might not be the best if we want to minimize the time complexity, as we will see below.

Quantum Arithmetic. Finally, we must design a quantum circuit that computes $\widehat{g}(x)$, compares $|\widehat{g}(x)|$ with M and computes $\lfloor 2^n M/|\widehat{g}(x)| \rfloor$. First, we notice that we can compute $|\widehat{g}(x)|$ exactly. The computation of each $\widehat{g_\beta}$ requires a circuit with approximately $2^8 \times 8 \times 4$ Toffoli and CNOT gates that, for each i, compares its input with i, and writes the corresponding output value. We do this 80 times in parallel. Overall, this costs $\leq 2^{20}$ Toffoli gates. Afterwards, we take the product of all coefficients, two by two: the bit-length of the numbers that we multiply increases at each product. This step costs $\leq 2^{17}$ Toffoli gates.

In the end $|\widehat{g}(x)|$ is an integer between 0 and $2^{80 \times 5} = 2^{400}$. The comparison with M is done with about 2×400 Toffoli gates (see Section 3). The computation of $\lfloor 2^n M/|\widehat{g}(x)| \rfloor$, which is required for amplitude transduction, is a Euclidean division of a 980-bit constant number by a 400-bit one, which is done with about $4 \times 980 \times 400 = 2^{20.6}$ Toffoli gates. In total, the overhead in Algorithm 1 with respect to the query of f can be upper bounded by 2^{22} Toffoli gates (we did not count the additional CNOT gates, but their numbers are of the same order).

4.2 State-recovery on Rocca: Recovering the State

When Algorithm 1 succeeds in both steps (rejection sampling and final measurement), we obtain the values for all hidden shifts $V_0, \ldots, V_4$, byte by byte, which we combine with the E_i that we can directly measure. We have the knowledge of:

$$\begin{cases} S[5] \oplus A(S[1]) \\ S[2] \oplus A(S[0] \oplus S[4]) \\ S[3] \oplus A(S[4]) \oplus A(S[7] \oplus A(S[0])) \\ S[3] \oplus S[7] \\ S[0] \oplus S[1] \oplus S[6] \oplus A(S[2]) \\ S[7] \\ S[7] \oplus A(S[0]) \oplus A(S[1] \oplus S[6]) \\ S[0] \oplus S[3] \oplus S[6] \oplus A(S[4] \oplus S[7] \oplus A(S[5])) \oplus A(S[3] \oplus A(S[4])) \end{cases}$$

This gives us a system of equations that we need to solve.

Preliminaries on AES-like Equation Systems. It is important to recall here that $A = \mathsf{MC} \circ \mathsf{SR} \circ \mathsf{SB}$ is a keyless AES round, where the SR operation shifts the bytes in row i by i positions left. This is represented in Figure 4. In particular, if we know a *diagonal* of S, then we can deduce a *column* of $A(S)$ (and the converse). However, knowing a column of S only allows to deduce an *antidiagonal* of $\mathsf{SR} \circ \mathsf{SB}(S)$.

Figure 4: Mapping of known bytes by A.

Solving the System: Step 1.   First, we obtain directly $S[3]$ and $S[7]$. We then consider the smaller system:

$$\begin{cases} (E1) & S[5] \oplus A(S[1]) \\ (E2) & S[2] \oplus A(S[0] \oplus S[4]) \\ (E3) & A(S[4]) \oplus A(S[7] \oplus A(S[0])) \\ (E4) & S[0] \oplus S[1] \oplus S[6] \oplus A(S[2]) \\ (E5) & A(S[0]) \oplus A(S[1] \oplus S[6]) \\ (E6) & S[0] \oplus S[6] \oplus A(S[4] \oplus S[7] \oplus A(S[5])) \oplus A(S[3] \oplus A(S[4])) \end{cases}$$

Focusing on $(E2)$ to $(E5)$, we deduce the following, where $*$ are known values:

$$\begin{cases} S[2] \oplus A(S[0] \oplus S[4]) = * \\ A(S[4]) \oplus A(* \oplus A(S[0])) = * \\ A(S[0]) \oplus A(S[0] \oplus A(S[2]) \oplus *) = * \end{cases} \text{ i.e. } \begin{cases} S[2] \oplus A(S[0] \oplus S[4]) = * \\ \mathsf{SB}(S[4]) \oplus \mathsf{SB}(* \oplus A(S[0])) = * \\ \mathsf{SB}(S[0]) \oplus \mathsf{SB}(S[0] \oplus A(S[2]) \oplus *) = * \end{cases}$$

In particular, the third inequality is obtained by replacing $S[1] \oplus S[6]$ in $(E5)$ by $A(S[2]) \oplus S[0] \oplus *$ from $(E4)$. This implies:

$$\begin{cases} S[4] = \mathsf{SB}^{-1}\left(\mathsf{SB}(* \oplus A(S[0])) \oplus *\right) \\ S[2] = * \oplus A\left[S[0] \oplus S[4]\right] \\ \mathsf{SB}(S[0]) \oplus \mathsf{SB}\left(S[0] \oplus A(S[2]) \oplus *\right) = * \end{cases} \tag{18}$$

This sub-system in $S[0], S[2], S[4]$ admits on average one solution, and we can solve it in time 2^{96} by the following strategy:

- Guess two columns and two diagonals of $S[0]$. Obtain two columns of $S[4]$ by the first equation (see Figure 5)

- Deduce two columns of $S[0] \oplus S[4]$.

- Obtain two columns and two diagonals of $A(S[2])$ by the third equation (see Figure 5)

- Deduce two diagonals of $S[2]$.

- Using the second equation, solve the obtained linear system in the 2 remaining diagonals of $S[2]$; obtain the whole $S[2]$ (on average one solution)

- Using the third equation, obtain $S[0]$. Each S-Box equation of the form $\mathsf{SBox}(* \oplus x) \oplus \mathsf{SBox}(* \oplus x) = *$ has on average one solution; half of the time they have zero solutions and half of the time, they have two solutions. So, if one of these equations has no solution, we backtrack.

- Otherwise, we have found 2^{16} possibilities for $S[0]$. We use the first equation to compute $S[4]$ and we check that all equations are satisfied.

Though we need to examine 2^{16} solutions for $S[0]$, this will be done only 2^{96-16} times, so overall the time to solve the sub-system is 2^{96}. For each guess there are a few AES rounds to compute and 16 S-Box differential equations to solve.

Figure 5: Representation of the first (top) and third (bottom) equations in Equation 18, and the bytes that we guessed. (*) denotes a known state.

Solving the System: Step 2. Having obtained $S[0], S[2], S[3], S[4], S[7]$, three equations remain:

$$\begin{cases} S[5] \oplus A(S[1]) = * \\ S[1] \oplus S[6] = * \\ S[6] \oplus A(* \oplus A(S[5])) = * \end{cases} \implies \begin{cases} S[5] \oplus A(S[1]) = * \\ S[1] \oplus A(* \oplus A(S[5])) = * \end{cases} \tag{19}$$

We solve this remaining sub-system as follows. We guess two diagonals of $S[5]$ (i.e., two columns of $A(S[5])$) and two diagonals of $A(S[5])$, for a total of 12 bytes, which are represented in Figure 6.

Next, we solve a linear system in $Y := \mathsf{SR} \circ \mathsf{SB}(S[1])$. Indeed, by the first line in Equation 19, we have two diagonals of $A(S[1]) = \mathsf{MC}(Y)$, e.g., the bytes $(0, 5, 10, 15, 4, 9, 14, 3)$ if we follow the pattern of the figure. By the second line, we have two columns of $S[1]$, e.g., the bytes $(0, 1, 2, 3, 4, 5, 6, 7)$, which give the bytes $(0, 7, 10, 13, 1, 4, 11, 14)$ of Y. It appears that for each column of Y, we know two bytes before and after the MC operation. Though the positions of these bytes differ for each column, thanks to the MDS property of the MC matrix, we can always express the two unknown bytes of Y as a linear combination of the four known ones.

Having obtained Y, we deduce $S[1]$, and check if both equations are satisfied. The time complexity of this step is therefore slightly smaller than the first one, since it does not require to solve S-Box differential equations.


Figure 6: Representation of Equation 19, and the bytes that we guessed. (*) denotes a know state.

Summary: Hybrid Attack. So far we are using a classical algorithm for the state-recovery part. With the selection of M that minimizes the number of superposition queries, the adversary queries its oracle for Rocca a total of $2^{46.34}$ times on average. After each query, they perform the amplitude product, costing 2^{22} Toffoli gates, and succeed in the first step $2^{29.40}$ times on average. For each of these successes, they retrieve a candidate value for the hidden shift and solve the equation system. Once the system is solved, the candidate internal state can be tested by computing backwards a few rounds and checking the ciphertexts.

Overall, this first hybrid attack costs $2^{46.34}$ superposition queries, $2^{96+29.4} = 2^{125.4}$ classical time to solve the system, and $2^{68.34}$ additional Toffoli gates.

Quantum Attack. We can accelerate the attack by using quantum search to speedup the system solving. To solve the first subsystem in $S[0], S[2], S[4]$, we proceed as follows: we create a quantum algorithm that samples a valid $S[0]$, i.e., a value of $S[0]$ that passes the S-Box differential equation, then tests if one of the 2^{16} possibilities solves the entire system. As seen in Section 3, the S-Box differential equation can be solved in 2^{12} Toffoli gates, and we have 16 of them to solve. Computing the remaining AES rounds costs less than 2^{16}.

Then this algorithm is a sequence of two Grover searches with Toffoli count:

$$\left(\frac{\pi}{2} 2^{16/2} + \frac{\pi}{2} 2^{16/2}\right) 2^{16} \simeq 2^{26} \ .$$

On the output of this algorithm, we use amplitude amplification [BHMT02]. By design, the probability that one of the possibilities for $S[0]$ solves the system is 2^{16-96}, so there are around $\frac{\pi}{4} 2^{(96-16)/2}$ iterates to make, and the total time is:

$$\frac{\pi}{2} 2^{(96-16)/2} \times 2^{26} \simeq 2^{67} \ .$$

At this point, our quantum attack requires $2^{46.34}$ superposition queries and $2^{67+29.4} = 2^{96.4}$ Toffoli gates. We can optimize this by noticing how the (average) Toffoli count depends on the probabilities p and p' to succeed in both steps of Algorithm 1:

$$\frac{1}{p'} \left[\frac{1}{p} \left(2^{22}\right) + 2^{67}\right] \ . \tag{20}$$

Since we know the entire distribution of the Walsh coefficients, we can solve this minimization problem on M and G, and we adopt:

$$M = 2^{304.56}, \quad G = 2^{625.89}, \quad p = \frac{M^2}{2^{2n}} G = 2^{-45.00}, \quad p' = \frac{G}{2^n} = 2^{-14.11} \tag{21}$$

which gives a complexity of $1/(pp') = 2^{59.11}$ Q2 encryption queries and

$$2^{14.11} \left(2^{22+45} + 2^{67}\right) \simeq 2^{81}$$

Toffoli gates. If we count that Q2 queries should have at least the same Toffoli cost as quantum implementations of Rocca, the gate count is comparable to the generic forgery attack in 2^{64} Q2 queries, though we do not require decryption queries anymore.

4.3 State-recovery on Rocca-S

The attack on Rocca-S is very similar to the one on Rocca. We have the same strategy: combine Algorithm 1 with linear post-processing to recover enough information on the internal state with a single query, and obtain this internal state by solving a simple system of equations.

Starting from an internal state S, we encrypt several message blocks and focus on the

following outputs:

$$C_0 = M_0 \oplus \underbrace{S[0] \oplus A\left(S[3] \oplus S[5]\right)}_{:=E_0}$$

$$C_0' = M_0' \oplus \underbrace{S[2] \oplus A\left(S[4] \oplus S[6]\right)}_{:=E_1}$$

$$C_1 = M_1 \oplus \underbrace{S[1] \oplus S[6] \oplus A\left(S[3] \oplus S[6] \oplus A\left(S[2]\right) \oplus A\left(S[4]\right)\right)}_{:=E_2}$$

$$C_1' = M_1' \oplus \underbrace{S[0] \oplus A\left(M_0' \oplus S[4] \oplus A\left(S[3]\right) \oplus A\left(S[5]\right)\right) \oplus A\left(S[1]\right)}_{:=E_3}$$

$$C_2' = M_2' \oplus S[1] \oplus S[6] \oplus A\Big(\underbrace{M_0}_{:=X_0} \oplus \underbrace{A\left(S[0]\right)}_{:=V_0}\Big) \oplus$$

$$A\Big(\underbrace{M_0' \oplus M_1'}_{:=X_1} \oplus \underbrace{A\left(S[3] \oplus A\left(S[4]\right)\right) \oplus A\left(S[3]\right) \oplus A\left(S[6] \oplus A\left(S[2]\right)\right)}_{:=V_1}\Big)$$

$$C_3' = M_3' \oplus M_0 \oplus S[4] \oplus A\Big(\underbrace{M_1}_{:=X_2} \oplus \underbrace{A\left(S[1] \oplus S[6]\right)}_{:=V_2}\Big) \oplus A\left(S[0]\right) \oplus A\left(S[5]\right) \oplus$$

$$A\Big[\underbrace{M_1' \oplus M_2'}_{=0} \oplus A\left(S[4] \oplus A\left(S[0] \oplus A\left(S[1]\right)\right) \oplus A\left(S[5]\right)\right) \oplus$$

$$A\Big(S[6] \oplus A\Big(\underbrace{M_0'}_{=0} \oplus A\left(S[3]\right)\Big) \oplus A\left(S[2]\right)\Big) \oplus A\left(S[6] \oplus A\left(S[2]\right)\right)\Big]$$

Similarly as before, we set 3 input variables X_0, X_1, X_2 such that:

$$M_0' = 0, \qquad\qquad M_1' \oplus M_2' = 0$$
$$X_0 = M_0, \qquad\qquad X_1 = M_0' \oplus M_1', \qquad\qquad X_2 = M_1 \ ,$$

and the other plaintext blocks are fixed to 0. We also define the three corresponding hidden shifts:

$$V_0 := A(S[0])$$
$$V_1 := A\left(S[3] \oplus A\left(S[4]\right)\right) \oplus A\left(S[3]\right) \oplus A\left(S[6] \oplus A\left(S[2]\right)\right)$$
$$V_2 := A\left(S[1] \oplus S[6]\right)$$

With this input, the expression of $C_0 \oplus M_0, C_0' \oplus M_0', C_1 \oplus M_1$ and $C_1' \oplus M_1'$ becomes constant, so we can immediately retrieve 4 values depending on the state S:

$$\begin{cases} E_0 := S[0] \oplus A\left(S[3] \oplus S[5]\right) \\ E_1 := S[2] \oplus A\left(S[4] \oplus S[6]\right) \\ E_2 := S[1] \oplus S[6] \oplus A\left(S[3] \oplus S[6] \oplus A\left(S[2]\right) \oplus A\left(S[4]\right)\right) \\ E_3 := S[0] \oplus A\left(S[4] \oplus A\left(S[3]\right) \oplus A\left(S[5]\right)\right) \oplus A\left(S[1]\right) \end{cases} \tag{22}$$

Alongside, we compute a linear function of $(C_2' \oplus M_2', C_3' \oplus M_3' \oplus M_0$ which inverts MixColumns, multiplies each S-Box by an arbitrary mask β, and XORs the results. The situation is similar to Rocca except that we only combine $3 \times 16 = 48$ S-Boxes instead of 80. This reduces somewhat the gate count overhead for arithmetic operations, which we can still upper bound at 2^{22} Toffoli gates. More importantly, it modifies the values of M, G, p and p'.

When Algorithm 1 succeeds in both steps, we obtain the values of V_0, V_1, V_2. The first one gives $S[0]$ immediately, so now we know:

$$
\begin{cases}
S[0] \\
S[3] \oplus S[5] \\
S[2] \oplus A\left(S[4] \oplus S[6]\right) \\
S[3] \oplus S[6] \oplus A\left(S[2]\right) \oplus A\left(S[4]\right) \\
A\left(S[4] \oplus A\left(S[3]\right) \oplus A\left(S[5]\right)\right) \oplus A\left(S[1]\right) \\
A\left(S[3] \oplus A\left(S[4]\right)\right) \oplus A\left(S[3]\right) \oplus A\left(S[6] \oplus A\left(S[2]\right)\right) \\
S[1] \oplus S[6]
\end{cases}
$$

We will solve this system in about 2^{64} quantum (or 2^{128} classical) computations. First, we guess $S[3]$ and deduce $S[5]$. We use the fourth and sixth equations of above, where $*$ denotes a known value:

$$
\begin{cases}
S[6] \oplus A(S[2]) \oplus A(S[4]) = * \\
A(* \oplus A(S[4])) \oplus A(S[6] \oplus A(S[2])) = *
\end{cases}
\tag{23}
$$

which implies $A(* \oplus A(S[4])) \oplus A(* \oplus A(S[4])) = *$. We can compose by the inverse of the AES linear layer to retrieve 16 independent differential S-Box equations of the form: $\mathsf{SB}(* \oplus x) \oplus \mathsf{SB}(* \oplus x) = *$, where the variable to recover is $A(S[4])$ (byte by byte). Like before, these equations have two solutions half of the time, and otherwise zero. So we need to try on average 2^{16} values of $S[3]$ until all 16 equations have solutions, and in that case, we need to check the 2^{16} different obtained values of $A(S[4])$. We will deduce the whole internal state and check the equations.

Again, checking $S[3]$ costs about 2^{16} Toffoli gates and checking a given solution costs less. We use amplitude amplification [BHMT02] over an algorithm that: finds a valid $S[3]$, then, searches through the corresponding solutions using a Grover search. We will find the internal state in time:

$$
\frac{\pi}{2} 2^{(128-16)/2} \left(\frac{\pi}{2} 2^{16/2} + \frac{\pi}{2} 2^{16/2} \right) 2^{16} \simeq 2^{82.3} \ .
$$

Summary of the Attack and Optimization.   We propose two optimizations of this attack: one that minimizes the number of Q2 queries (i.e., maximizes pp'), and one that minimizes the total time complexity. In the first case, we set:

$$
M = 2^{195.40}, \quad G = 2^{365.62}, \quad p = 2^{-11.58}, \quad p' = 2^{-18.37}, \quad pp' = 2^{-29.95} \ .
\tag{24}
$$

The adversary queries its oracle for Rocca-S a total of $2^{29.95} \simeq 2^{30}$ times on average before encountering a success. The quantum arithmetic requires an additional time of $2^{29.95} \times 2^{22} \simeq 2^{52}$ Toffoli gates. Solving the equation system happens only if the first step (amplitude product) succeeded, so $2^{18.37}$ times. We can solve it quantumly, for a total Toffoli count $2^{18.37} \times 2^{82.3} \simeq 2^{101}$.

However, the average Toffoli count can be expressed as:

$$
\frac{1}{p'} \left[\frac{1}{p}(2^{22}) + 2^{82.3} \right] \ .
\tag{25}
$$

By minimizing this expression instead, we obtain the following choice:

$$
M = 2^{164.34}, \quad G = 2^{379.02}, \quad p = 2^{-60.30}, \quad p' = 2^{-4.98}, \quad pp' = 2^{-65.28} \ .
\tag{26}
$$

which gives a complexity of $\simeq 2^{65}$ Q2 encryption queries and $\simeq 2^{87}$ Toffoli gates.

4.4 Key-recovery on Tiaoxin

Our method allows to recover the state T_3 at some point of the encryption phase. Afterwards, we can invert the round function on T_3, and the initialization phase, and recover the key which was loaded in the initial state.

Let us fix the state $T_3[0, 1, 2]$ at the beginning of the encryption phase and unroll a few ciphertexts:

$$C_0 = M_0 \oplus T_3[0] \oplus T_3[1] \oplus A(T_3[2]) \oplus A(T_4[0]) \oplus \text{AND}(T_6[2], T_4[2])$$
$$C_0' = M_0 \oplus M_0' \oplus T_4[1] \oplus T_6[0] \oplus A(T_3[0]) \oplus A(T_6[5]) \oplus \text{AND}(T_6[4], T_3[1])$$
$$C_1' = M_0 \oplus M_1 \oplus M_0' \oplus M_1' \oplus T_6[0] \oplus A(\underbrace{M_0}_{:=X_0} \oplus \underbrace{T_3[0] \oplus A(T_3[2])}_{:=V_0}) \oplus$$
$$A(T_6[4]) \oplus A(T_6[5]) \oplus A(T_4[0]) \oplus \text{AND}(T_6[3], A(T_3[0]))$$
$$C_3' = M_0 \oplus M_1 \oplus M_2 \oplus M_3 \oplus M_0' \oplus M_1' \oplus M_2' \oplus M_3' \oplus$$
$$T_6[0] \oplus A(T_6[3]) \oplus A(T_6[4]) \oplus A(T_6[5]) \oplus A(T_6[2])$$
$$A\big[\underbrace{M_0 \oplus M_1 \oplus M_2}_{:=X_2} \oplus \underbrace{T_3[0] \oplus A(T_3[1]) \oplus A(T_3[2]) \oplus A(A(T_3[0]))}_{:=V_2}\big] \oplus$$
$$A\big[M_0' \oplus M_1' \oplus T_4[0] \oplus A(T_4[2]) \oplus A(T_4[3])\big] \oplus$$
$$\text{AND}(T_6[1], A(\underbrace{M_0 \oplus M_1}_{:=X_1} \oplus \underbrace{T_3[0] \oplus A(T_3[1]) \oplus A(T_3[2])}_{:=V_1}))$$

We set the following as variables: $X_0 = M_0, X_1 = M_0 \oplus M_1, X_2 = M_0 \oplus M_1 \oplus M_2$. The rest is fixed. We focus on C_1' and C_3', and define the shift values:

$$\begin{cases} V_0 := T_3[0] \oplus A(T_3[2]) \\ V_1 := T_3[0] \oplus A(T_3[1]) \oplus A(T_3[2]) \\ V_2 := T_3[0] \oplus A(T_3[1]) \oplus A(T_3[2]) \oplus A(A(T_3[0])) \end{cases} \tag{27}$$

We then observe the XOR of C_1' and C_3'. More precisely, let L be the function that selects one bit in each column of the state and XORs them. We assume that on these 4 bits, $T_6[1] = 1$.

Remark 5. Notice that the outputs of the S-Boxes are processed with different masks than before. While the choice of mask was inconsequential for Rocca and Rocca-S, here it becomes quite important, as we have to ensure that $T_6[1] = 1$ at each bit position selected by the mask. A similar constraint occurs for AEGIS-128L in the next section.

Then we have:

$$L\left[\text{AND}(T_6[1], A(M_0 \oplus M_1 \oplus T_3[0] \oplus A(T_3[1]) \oplus A(T_3[2])))\right]$$
$$= L \circ A(M_0 \oplus M_1 \oplus T_3[0] \oplus A(T_3[1]) \oplus A(T_3[2])) ,$$

and we define the function:

$$F(C_1' \oplus M_0 \oplus M_1 \oplus M_0' \oplus M_1', C_3' \oplus M_0 \oplus M_1 \oplus M_2 \oplus M_3 \oplus M_0' \oplus M_1' \oplus M_2' \oplus M_3')$$
$$= b \oplus L \circ A(M_0 \oplus V_0) \oplus L \circ A(M_0 \oplus M_1 \oplus V_1) \oplus L \circ A(M_0 \oplus M_1 \oplus M_2 \oplus V_2) ,$$

where b is an unknown bit depending on T. We have $L \circ A = (L \circ \text{MC} \circ \text{SR}) \circ \text{SB}$, so column by column, we have a one-bit function of the S-Box outputs, which can be rewritten as: $(x_0, x_1, x_2, x_3) \mapsto \bigoplus_i \alpha_i \cdot \text{SB}(x_i)$ for well-chosen masks $\alpha_0, \ldots, \alpha_3$.

The situation is thus the same as before, since the distribution of Walsh coefficients is independent of the mask α_i (as long as it's nonzero). Recovering the entire state

$T_3[0, 1, 2]$ from the shifts V_0, V_1, V_2 is trivial and costs only a few AES rounds. Afterwards, we compute backwards through the 15 rounds of initialization on the state T_3 (30 AES rounds), and obtain a candidate key K that we can immediately check. All of this can be done classically.

Since there are $16 \times 3 = 48$ S-Boxes in the function, we optimize the probability similarly to Rocca-S and obtain $pp' \simeq 2^{-30}$. The Toffoli cost of the entire attack is roughly $2^{30} \times 2^{22} = 2^{52}$ and it contains 2^{30} Q2 queries. Note that we need to multiply these numbers by 2^4, since we assumed to have guessed correctly 4 bits of $T_6[1]$.

4.5   State-recovery on AEGIS-128L

Contrary to the rest of this section, the attack on AEGIS-128L uses a function of smaller correlation: we use Theorem 2.

Starting from an initial state S, we encrypt pairs of message blocks (M_i, M'_i) with $M'_i = 0$. To simplify the notations, we will express the ciphertext blocks in function of T, the state after one update, that is, $T[i] = S[i] \oplus A(S[i-1])$. Our aim is to recover T. As they cannot be expressed from T, we ignore the first pair of ciphertext blocks and focus on:

$$C_1 = M_1 \oplus T[1] \oplus T[6] \oplus \mathrm{AND}\,(T[2], T[3])$$
$$C'_1 = M'_1 \oplus T[2] \oplus T[5] \oplus \mathrm{AND}\,(T[6], T[7])$$
$$C'_2 = M'_2 \oplus T[2] \oplus T[5] \oplus A\,(M'_0 \oplus T[4]) \oplus A\,(T[1])$$
$$\qquad \oplus \mathrm{AND}\,(T[6] \oplus A\,(T[5])\,, T[7] \oplus A\,(T[6]))$$
$$C_6 = M_6 \oplus A\Big[M_0 \oplus M_1 \oplus M_2 \oplus M_3 \oplus M_4 \oplus T[0] \oplus A\,(T[7] \oplus A\,(T[6]))$$
$$\qquad \oplus A\,(T[7]) \oplus A\,(T[7] \oplus A\,(T[6] \oplus A\,(T[5])) \oplus A\,(T[6]))$$
$$\qquad \oplus A\Big(T[7] \oplus A\,(T[6] \oplus A\,(T[5] \oplus A\,(M'_0 \oplus T[4])) \oplus A\,(T[5]))$$
$$\qquad\quad \oplus A\,(T[6] \oplus A\,(T[5])) \oplus A\,(T[6])\Big)\Big]$$
$$\oplus A\Big[M_0 \oplus M_1 \oplus M_2 \oplus M_3 \oplus T[0] \oplus A\,(T[7] \oplus A\,(T[6])) \oplus A\,(T[7])$$
$$\qquad \oplus A\,(T[7] \oplus A\,(T[6] \oplus A\,(T[5])) \oplus A\,(T[6]))\Big]$$
$$\oplus A\Big[M_0 \oplus M_1 \oplus M_2 \oplus T[0] \oplus A\,(T[7] \oplus A\,(T[6])) \oplus A\,(T[7])\Big]$$
$$\oplus A\Big[M_0 \oplus M_1 \oplus T[0] \oplus A\,(T[7])\Big]$$
$$\oplus A\,(M_0 \oplus T[0]) \oplus Y \oplus \mathrm{AND}(h'(M_0, M_1, M_2, M_3), h''(M_0, M_1, M_2))\ ,$$

where h' and h'' are two functions whose exact expression is irrelevant here, and Y is a constant (an expression in which only T and M'_i intervene). Similarly to Tiaoxin, we use a linear post-processing which truncates C_6 to only 4 bits. Therefore, though it is completely unknown (and depends on the unknown state T), the AND term will become a function h with correlation 2^{-4}. Heuristically, we model this function (that will change for each query) as a random one, and we use Theorem 2.

By making M_0 to M_4 vary, we obtain 5 shifts which give us:

$T[0]$	from the M_0 shift
$T[7]$	from the M_1 shift, knowing $T[0]$
$T[6]$	from the M_2 shift, knowing $T[0, 7]$
$T[5]$	from the M_3 shift, knowing $T[0, 6, 7]$
$T[4]$	from the M_4 shift, knowing $T[0, 5, 6, 7]$

Next, we focus on the ciphertext blocks C_1 to C_2' which we have also obtained. Thanks to C_1' and all the state registers that we know, we obtain $T[2]$. Next, thanks to C_2', we obtain $T[1]$. The only register of T that we are missing is $T[3]$. We can find half of it using the expression of C_1: indeed, from the known $T[i]$ we can compute $\mathrm{AND}(T[2], T[3])$. We can expect half the bits of $T[2]$ to be one, which gives us the the corresponding bits of $T[3]$.

State-recovery Attack. After performing this partial state-recovery, we can do a Grover search on the remaining 64 bits of the state. However, checking if we have obtained a valid internal state is not trivial. Indeed, the round function of AEGIS is not invertible, so we cannot compute backwards and check with previous ciphertexts. In fact, there do not seem to be other ciphertext equations that we can exploit (either we have already used them, or they depend on the varying M_i).

Consequently, we do a Grover search using *superposition decryption queries* to test our guess of the state. That is, starting from the recovered internal state, we compute the tag (approximately 6×8 AES rounds, i.e., $\leq 2^{16}$ Toffoli gates) and we try to decipher with an oracle. If the internal state is guessed correctly, the oracle will accept. This operation requires approximately: $\frac{\pi}{2} 2^{32} \times 2^{16} \leq 2^{49}$ gates and 2^{33} decryption queries.

Since the number of S-Boxes is the same as in Rocca, the hidden shift algorithm is actually the same. We keep the same p and p', but introduce the correlation $c = 2^{-4}$. The Toffoli count and the number of queries are respectively:

$$\frac{2}{p'}\left(\frac{1}{p}\frac{1}{c^2}2^{22} + 2^{49}\right) \qquad \text{and} \qquad \frac{2}{p'}\left(\frac{1}{p}\frac{1}{c^2} + 2^{33}\right) \ . \tag{28}$$

If we optimize the Toffoli count, we get the following parameters:

$$M = 2^{324.23}, G = 2^{612.54}, p = 2^{-19.00}, p' = 2^{-27.46} \ , \tag{29}$$

which give $2^{77.46}$ Toffolis, smaller than the cost of Grover search (2^{81}), and $\frac{2 \times 2^{33}}{p'} \leq 2^{62}$ decryption queries, which is also smaller than a forgery attack using a Grover search. We also use $\frac{2}{pp'c^2} \leq 2^{56}$ encryption queries.

5 Discussion

In all instances of our attack, the AE scheme (Rocca, Rocca-S, Tiaoxin, AEGIS) is believed to be secure regarding guess-and-determine attacks that aim at recovering the state. Indeed, when one only observes the ciphertext blocks, the obtained system of equations is intractable.

Our quantum attack works because we can observe hidden shifts in addition to the ciphertexts. This allows us to reduce the state-recovery to a simpler system of equations (the simplest being Tiaoxin-346 which only relies on three shifts). However, there are limitations to this approach. Notably, if we have a ciphertext $C = S_0 \oplus A(S_1 \oplus M_1)$, we have only two choices: either make $M_1 = 0$ a constant, and observe $S_0 \oplus A(S_1)$, or make M_1 a variable, and observe S_1. In the latter case, S_0 is lost. Besides, we can only use one variable for one shift, i.e., if we have $C = A(S_1 \oplus M_1)$ and $C' = A(S_2 \oplus M_1)$, we must drop one of the ciphertext blocks. Another problematic case is when we observe $A(S_0 \oplus A(S_1 \oplus M_1))$. Though we do have a shifted function, the function is now unknown (it depends on S_0) and more complex (two rounds of AES instead of one). The attack can proceed by guessing enough bits of S_0, but becomes more difficult.

In our examples, the choice of the shifts was done by hand, trying to obtain the simplest equation system. More clever choices might still exist. Conversely, making the schemes

secure against this attack means ensuring that none of the equation systems resulting from a combination of ciphertexts and shifts can be tractable.

Previously, some Q2 quantum attacks have been linked to more efficient Q1 attacks [BHN+19]. However such methods do not appear to work in this scenario, as classical queries will have different nonces, and cannot be brought together to emulate a single quantum query. To the best of our knowledge, all the schemes studied in this paper remain secure against Q1 attacks.

As a final remark, we note that the attacks presented in this paper have time and query complexities below those of exhaustive search, without taking parallelization into account. However, the generic attacks are instances of quantum search, while the trials in our attacks can be parallelized perfectly. As a consequence, there might exist other targets than those given in this paper, on which an advantage against exhaustive search is reached under some depth constraint.

Acknowledgments

References

[ABKM22] Gorjan Alagic, Chen Bai, Jonathan Katz, and Christian Majenz. Post-quantum security of the Even-Mansour cipher. In *EUROCRYPT (3)*, volume 13277 of *Lecture Notes in Computer Science*, pages 458–487. Springer, 2022.

[AI23] Ravi Anand and Takanori Isobe. Quantum security analysis of Rocca. *Quantum Information Processing*, 22(4):164, 2023.

[ATTU16] Mayuresh Vivekanand Anand, Ehsan Ebrahimi Targhi, Gelo Noel Tabia, and Dominique Unruh. Post-quantum security of the CBC, CFB, OFB, CTR, and XTS modes of operation. In *PQCrypto*, volume 9606 of *Lecture Notes in Computer Science*, pages 44–63. Springer, 2016.

[BBC+21] Ritam Bhaumik, Xavier Bonnetain, André Chailloux, Gaëtan Leurent, María Naya-Plasencia, André Schrottenloher, and Yannick Seurin. QCB: efficient quantum-secure authenticated encryption. In *ASIACRYPT (1)*, volume 13090 of *Lecture Notes in Computer Science*, pages 668–698. Springer, 2021.

[BHMT02] Gilles Brassard, Peter Hoyer, Michele Mosca, and Alain Tapp. Quantum amplitude amplification and estimation. *Contemporary Mathematics*, 305:53–74, 2002.

[BHN+19] Xavier Bonnetain, Akinori Hosoyamada, María Naya-Plasencia, Yu Sasaki, and André Schrottenloher. Quantum attacks without superposition queries: The offline Simon's algorithm. In *ASIACRYPT (1)*, volume 11921 of *Lecture Notes in Computer Science*, pages 552–583. Springer, 2019.

[BHT98] Gilles Brassard, Peter Høyer, and Alain Tapp. Quantum cryptanalysis of hash and claw-free functions. In *LATIN*, volume 1380 of *Lecture Notes in Computer Science*, pages 163–169. Springer, 1998.

[BLNS21]   Xavier Bonnetain, Gaëtan Leurent, María Naya-Plasencia, and André Schrottenloher. Quantum linearization attacks. In *ASIACRYPT (1)*, volume 13090 of *Lecture Notes in Computer Science*, pages 422–452. Springer, 2021.

[BN18]   Xavier Bonnetain and María Naya-Plasencia. Hidden shift quantum cryptanalysis and implications. In *ASIACRYPT (1)*, volume 11272 of *Lecture Notes in Computer Science*, pages 560–592. Springer, 2018.

[BNS19]   Xavier Bonnetain, María Naya-Plasencia, and André Schrottenloher. Quantum security analysis of AES. *IACR Trans. Symmetric Cryptol.*, 2019(2):55–93, 2019.

[Bon17]   Xavier Bonnetain. Quantum key-recovery on full AEZ. In *SAC*, volume 10719 of *Lecture Notes in Computer Science*, pages 394–406. Springer, 2017.

[BSS22]   Xavier Bonnetain, André Schrottenloher, and Ferdinand Sibleyras. Beyond quadratic speedups in quantum attacks on symmetric schemes. In *EURO-CRYPT (3)*, volume 13277 of *Lecture Notes in Computer Science*, pages 315–344. Springer, 2022.

[BV97]   Ethan Bernstein and Umesh V. Vazirani. Quantum complexity theory. *SIAM J. Comput.*, 26(5):1411–1473, 1997.

[CDKM04]   Steven A Cuccaro, Thomas G Draper, Samuel A Kutin, and David Petrie Moulton. A new quantum ripple-carry addition circuit. *arXiv preprint quant-ph/0410184*, 2004.

[CHLS20]   Carlos Cid, Akinori Hosoyamada, Yunwen Liu, and Siang Meng Sim. Quantum cryptanalysis on contracting feistel structures and observation on related-key settings. In *INDOCRYPT*, volume 12578 of *Lecture Notes in Computer Science*, pages 373–394. Springer, 2020.

[ENP19]   Maria Eichlseder, Marcel Nageler, and Robert Primas. Analyzing the linear keystream biases in AEGIS. *IACR Trans. Symmetric Cryptol.*, 2019(4):348–368, 2019.

[Gro96]   Lov K. Grover. A fast quantum mechanical algorithm for database search. In *STOC*, pages 212–219. ACM, 1996.

[HII+22]   Akinori Hosoyamada, Akiko Inoue, Ryoma Ito, Tetsu Iwata, Kazuhiko Mimematsu, Ferdinand Sibleyras, and Yosuke Todo. Cryptanalysis of rocca and feasibility of its security claim. *IACR Trans. Symmetric Cryptol.*, 2022(3):123–151, 2022.

[HS18]   Akinori Hosoyamada and Yu Sasaki. Quantum demiric-selçuk meet-in-the-middle attacks: Applications to 6-round generic feistel constructions. In *SCN*, volume 11035 of *Lecture Notes in Computer Science*, pages 386–403. Springer, 2018.

[HS20]   Akinori Hosoyamada and Yu Sasaki. Finding hash collisions with quantum computers by using differential trails with smaller probability than birthday bound. In *EUROCRYPT (2)*, volume 12106 of *Lecture Notes in Computer Science*, pages 249–279. Springer, 2020.

[JBS+22]   Kyungbae Jang, Anubhab Baksi, Gyeongju Song, Hyunji Kim, Hwajeong Seo, and Anupam Chattopadhyay. Quantum analysis of AES. *IACR Cryptol. ePrint Arch.*, page 683, 2022.

[KEM17]   Daniel Kales, Maria Eichlseder, and Florian Mendel. Note on the robustness of CAESAR candidates. *IACR Cryptol. ePrint Arch.*, page 1137, 2017.

[KLLN16a] Marc Kaplan, Gaëtan Leurent, Anthony Leverrier, and María Naya-Plasencia. Breaking symmetric cryptosystems using quantum period finding. In *CRYPTO (2)*, volume 9815 of *Lecture Notes in Computer Science*, pages 207–237. Springer, 2016.

[KLLN16b] Marc Kaplan, Gaëtan Leurent, Anthony Leverrier, and María Naya-Plasencia. Quantum differential and linear cryptanalysis. *IACR Trans. Symmetric Cryptol.*, 2016(1):71–94, 2016.

[KM10]   Hidenori Kuwakado and Masakatu Morii. Quantum distinguisher between the 3-round feistel cipher and the random permutation. In *ISIT*, pages 2682–2685. IEEE, 2010.

[KM12]   Hidenori Kuwakado and Masakatu Morii. Security on the quantum-type even-mansour cipher. In *ISITA*, pages 312–316. IEEE, 2012.

[LIMS21]   Fukang Liu, Takanori Isobe, Willi Meier, and Kosei Sakamoto. Weak keys in reduced AEGIS and Tiaoxin. *IACR Trans. Symmetric Cryptol.*, 2021(2):104–139, 2021.

[LM17]   Gregor Leander and Alexander May. Grover meets simon - quantumly attacking the fx-construction. In *ASIACRYPT (2)*, volume 10625 of *Lecture Notes in Computer Science*, pages 161–178. Springer, 2017.

[Nat01]   National Institute of Standards and Technology. FIPS 197 advanced encryption standard (AES), 2001.

[NC02]   Michael A Nielsen and Isaac Chuang. Quantum computation and quantum information, 2002.

[NFI]   Yuto Nakano, K. Fukushima, and Takanori Isobe. Encryption algorithm Rocca-S. IETF Draft Standard.

[Nik16]   Ivica Nikolic. Tiaoxin-346 (v2.1). Submission to the CAESAR competition, 2016.

[ORR13]   Maris Ozols, Martin Roetteler, and Jérémie Roland. Quantum rejection sampling. *ACM Trans. Comput. Theory*, 5(3):11:1–11:33, 2013.

[Röt10]   Martin Rötteler. Quantum algorithms for highly non-linear boolean functions. In *SODA*, pages 448–457. SIAM, 2010.

[Sch23]   André Schrottenloher. Quantum linear key-recovery attacks using the QFT. In *CRYPTO (5)*, volume 14085 of *Lecture Notes in Computer Science*, pages 258–291. Springer, 2023.

[Sho94]   Peter W. Shor. Algorithms for quantum computation: Discrete logarithms and factoring. In *FOCS*, pages 124–134. IEEE Computer Society, 1994.

[Sim97]   Daniel R. Simon. On the power of quantum computation. *SIAM J. Comput.*, 26(5):1474–1483, 1997.

[SLN+21]   Kosei Sakamoto, Fukang Liu, Yuto Nakano, Shinsaku Kiyomoto, and Takanori Isobe. Rocca: An efficient AES-based encryption scheme for beyond 5G. *IACR Trans. Symmetric Cryptol.*, 2021(2):1–30, 2021.

[SLN+22] Kosei Sakamoto, Fukang Liu, Yuto Nakano, Shinsaku Kiyomoto, and Takanori Isobe. Rocca: An efficient AES-based encryption scheme for beyond 5G (full version). *IACR Cryptol. ePrint Arch.*, page 116, 2022.

[SLSB19] Yuval R Sanders, Guang Hao Low, Artur Scherer, and Dominic W Berry. Black-box quantum state preparation without arithmetic. *Physical review letters*, 122(2):020502, 2019.

[The24] The Sage Developers. *SageMath, the Sage Mathematics Software System (Version 9.6)*, 2024. `https://www.sagemath.org`.

[vDHI06] Wim van Dam, Sean Hallgren, and Lawrence Ip. Quantum algorithms for some hidden shift problems. *SIAM J. Comput.*, 36(3):763–778, 2006.

[WP13a] Hongjun Wu and Bart Preneel. AEGIS: A fast authenticated encryption algorithm. In *Selected Areas in Cryptography*, volume 8282 of *Lecture Notes in Computer Science*, pages 185–201. Springer, 2013.

[WP13b] Hongjun Wu and Bart Preneel. AEGIS: A fast authenticated encryption algorithm (full version). *IACR Cryptol. ePrint Arch.*, page 695, 2013.

[WP16] Hongjun Wu and Bart Preneel. AEGIS: a fast authenticated encryption algorithm (v1.1). Submission to the CAESAR competition, 2016.