Class-Average Differential Uniformity and Fixed-Instance Constraints in Modular and Rounded Feistel Components

Yu Morishima[†,§]

Masahiro Kaminaga[†]

† Tohoku Gakuin University, Miyagi, Japan.

[§]morishima@mail.tohoku-gakuin.ac.jp

Abstract

This paper analyzes differential probabilities in one Feistel round with binary-encoded modular-linear maps or rounded linear maps as internal components. For any fixed round key and any nonzero input difference chosen independently of the matrix instance, the output-difference distribution is uniform after averaging over the instance distribution. For a fixed instance, the 2-adic valuations of the relevant matrix entries deterministically fix certain output-difference bits. We consider an adversary that selects a low-weight input difference maximizing the number of fixed bits and derive lower bounds on the restricted maximum differential probability and its average over instances. The matrix-averaged lower bound is expressed in terms of the exponential moment of the maximum fixed-bit count and is compared with the generic lower bound induced by differential-pair symmetry.

1 Introduction

In block-cipher design, resistance to differential cryptanalysis is assessed through differential probabilities from input differences to output differences, with the maximum differential probability serving as a basic measure of randomness in round functions and S-boxes[1, 2]. In analyses based on expected differential probability and the Markov-cipher framework, differential characteristics averaged over a cipher ensemble or parameter distribution must be clearly distinguished from those associated with a particular fixed instance[3, 4]. Because class-averaged uniformity does not imply fixed-instance uniformity, we analyze the two settings separately, using lower bounds for the fixed-instance quantities.

Lattice-based problems underpin many modern public-key constructions[5, 6], and related maps have also been studied as components of symmetric-key primitives, including hash functions and pseudorandom functions. In the

modular-linear map of an Ajtai-type short integer solution (SIS) hash, the modular-linear transformation defined by a matrix $A \in \mathbb{Z}_q^{n\times m}$,

$$x \longmapsto Ax \bmod q$$

plays a fundamental role[7]. Learning with rounding (LWR), a modification of learning with errors (LWE), replaces random error with a rounding operation and uses deterministic arithmetic maps defined by matrices to construct pseudorandom functions and symmetric-key encryption schemes[8, 9, 10].

However, the collision resistance of SIS-type maps and the pseudorandomness of LWE/LWR-type maps do not directly characterize their output-difference distributions or maximum differential probabilities when the maps are viewed as bit-string maps. In particular, such matrix-based maps are often compressing or noninvertible. When their properties as deterministic mappings are studied in a block-cipher setting, it is therefore natural to use them as the F function of a Feistel structure, since the Feistel construction does not require F itself to be invertible. The purpose of this paper is neither to derive block-cipher security from the underlying hardness assumptions nor to establish the pseudorandom permutation guarantee of the Luby–Rackoff theorem, which assumes pseudorandom round functions[11]. Therefore, the terms SIS-type and LWR-type are used here only to denote the classes of maps whose input-output difference interactions are analyzed.

The differential characteristics of a single Feistel round provide a basis for evaluating complete multiround constructions. Accordingly, we analyze the class-average differential probability for input differences fixed independently of the matrix and derive lower bounds on the restricted maximum differential probability when an input difference may be selected after observing the public matrix. The effects of modular arithmetic on output-difference distributions have also been studied in addition–rotation–XOR (ARX) differential analysis[12, 13]. For conventional ARX constructions, differential analysis focuses primarily on carry propagation in modular addition and on dependencies among multiple additions. By contrast, SIS- and LWR-type maps are constructed from matrix-vector products, making the relation between the support of an input difference and the corresponding matrix entries central to the analysis. We derive conditions under which the 2-adic valuations of those entries deterministically fix output-difference bits. In what follows, we refer to each such condition as a 2-adic fixed-bit certificate, or simply as a certificate. The number of bits fixed by these conditions is then used to derive differential-probability bounds.

The main results of this paper are summarized as follows.

Our first result concerns output-difference uniformity under averaging over the matrix-generation distribution. For any fixed round key, if the matrix A is sampled uniformly and the result is averaged over the matrix-generation distribution, then for every nonzero input difference fixed independently of A, the output-difference distributions of the F functions based on the SIS-type and LWR-type maps are uniform.

Our second result concerns fixed-instance differential behavior when the input difference may be selected after the public matrix is known. The 2-adic val-

uation conditions determine fixed output-difference bits and yield lower bounds on the restricted maximum differential probability. We further derive a distributional lower bound for the maximum fixed-bit count and compare the resulting bound with the generic lower bound for fixed functions.

2 Feistel round functions from modular-linear and rounded components

Throughout, let $\mathbb{F}_2 = \{0,1\}$ and $\mathbb{Z}_q = \mathbb{Z}/q\mathbb{Z}$, and let $\oplus$ denote bitwise XOR. We take the standard representatives of residue classes modulo 2^r to be $\{0,1,\ldots,2^r-1\}$. For $u \in \mathbb{Z}_{2^r}$, write its binary representation as

$$u = \sum_{b=0}^{r-1} u_b 2^b, \qquad u_b \in \{0,1\},$$

and define

$$\mathcal{E}_r(u) = (u_0, u_1, \ldots, u_{r-1}) \in \mathbb{F}_2^r,$$

where bit position $b = 0$ denotes the least significant bit. When $\mathcal{E}_r(u)$ is applied to a vector $u = (u_1, \ldots, u_n) \in \mathbb{Z}_{2^r}^n$, the encodings $\mathcal{E}_r$ of the individual components are concatenated.

Consider a Feistel-type round function with block length m, and write the state in round i as

$$(L_i, R_i) \in \mathbb{F}_2^{m/2} \times \mathbb{F}_2^{m/2},$$

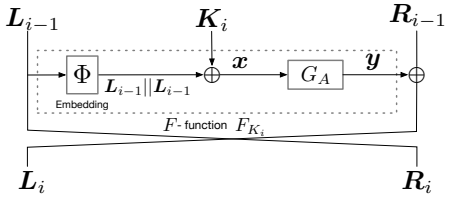


Figure 1: Structure of a Feistel round using F_{A,K_i}.

and update the round as follows using a round key $K_i \in \mathbb{F}_2^m$ and an F function F_{K_i}:

$$R_i = L_{i-1}, \qquad L_i = R_{i-1} \oplus F_{K_i}(L_{i-1}),$$

where m is even.

Figure 1 shows the structure of the round function. The value L_{i-1} input to F_{K_i} is embedded through a fixed injective $\mathbb{F}_2$-linear map $\Phi : \mathbb{F}_2^{m/2} \longrightarrow \mathbb{F}_2^m$ and then XORed with the round key K_i to obtain $x \in \mathbb{F}_2^m$. The internal function $G_A : \mathbb{F}_2^m \to \mathbb{F}_2^{m/2}$ depends on the matrix instance A and maps the input x to

an output $y \in \mathbb{F}_2^{m/2}$. To make the dependence on the matrix instance explicit, we write the F function as

$$F_{A,K_i}(L_{i-1}) = G_A(K_i \oplus \Phi(L_{i-1})).$$

For a fixed round key K_i, as L_{i-1} ranges over $\mathbb{F}_2^{m/2}$, the inputs to G_A range over the coset $K_i \oplus \mathrm{Range}(\Phi)$. Thus, the fixed-instance differential characteristics may generally depend not only on the matrix A, but also on the coset $K_i \oplus \mathrm{Range}(\Phi)$ determined by the fixed key K_i.

Remark 1. *To provide the Feistel half-block L_{i-1} as input to the internal function $G_A : \mathbb{F}_2^m \to \mathbb{F}_2^{m/2}$, we use the following linear embedding:*

$$\Phi(L_{i-1}) = L_{i-1} \| L_{i-1}.$$

This simple design places a Feistel half-block in the input space of a compressing modular-linear map from a long binary input to a shorter residue vector, a structure characteristic of Ajtai-type SIS hashes. Alternatively, the matrix dimensions could be chosen to match the half-block input, allowing L_{i-1} to be used directly. The resulting internal map, however, would be non-compressing, with equal input and output dimensions. To retain Ajtai's compressing internal map G_A while using it as a Feistel F function, we embed the half-block into an m-bit internal input.

The 2-adic fixed-bit certificate conditions and the corresponding differential-probability lower bounds for a fixed matrix apply to any injective $\mathbb{F}_2$-linear embedding through $\Delta x = \Phi(\Delta L_{i-1})$. In contrast, the candidate-difference set and distributional evaluation in Sec. 5.2 use $\Phi(L) = L \| L$; for a different embedding, the distribution formulas must be reevaluated according to $\mathrm{Supp}(\Phi(\Delta L))$.

Rather than presenting a complete multiround attack, we analyze the differential characteristics of one Feistel round through its round function F_{A,K_i}. In the Feistel update, XOR with the right-half input difference only relabels the possible left-half output differences and hence preserves their probabilities. We therefore consider the output y and its difference when a difference is applied to L_{i-1}.

For a nonzero input difference ΔL_{i-1}, let the input pair be L_{i-1}, L'_{i-1}, the corresponding output difference and output pair be Δy and y, y', and the difference at the input to the internal function G_A be Δx. Then the following relations hold:

$$\begin{aligned} \Delta L_{i-1} &= L_{i-1} \oplus L'_{i-1}, \\ \Delta y &= y \oplus y', \\ \Delta x &= \Phi(\Delta L_{i-1}). \end{aligned}$$

Differential cryptanalysis exploits biases in the occurrence probabilities of output differences for a given input difference. For a fixed matrix A and a fixed round key K_i, define the differential probability (DP) of the round-function

instance by

$$\begin{aligned}&\mathrm{DP}_{A,K_i}(\Delta L_{i-1}, \Delta y)\\ &= \Pr_{L_{i-1}\leftarrow\mathbb{F}_2^{m/2}}\Big[F_{A,K_i}(L_{i-1}) \oplus F_{A,K_i}(L_{i-1}\oplus\Delta L_{i-1})\\ &\qquad\qquad = \Delta y\Big],\end{aligned} \tag{1}$$

and define the maximum output differential probability for a given nonzero input difference ΔL_{i-1} by

$$p_{\max,A,K_i}(\Delta L_{i-1}) := \max_{\Delta y\in\mathbb{F}_2^{m/2}} \mathrm{DP}_{A,K_i}(\Delta L_{i-1}, \Delta y).$$

When the matrix A is public, an adversary can choose an advantageous difference from multiple nonzero input differences after observing A. Let the set of input differences available to the adversary be

$$\mathcal{D} \subseteq \mathbb{F}_2^{m/2} \setminus \{0\},$$

and define the maximum differential probability restricted to input differences in $\mathcal{D}$, namely the restricted maximum differential probability (MDP), by

$$\mathrm{MDP}^{\mathcal{D}}_{A,K_i} := \max_{\Delta L_{i-1}\in\mathcal{D}} p_{\max,A,K_i}(\Delta L_{i-1}).$$

When $\mathcal{D} = \mathbb{F}_2^{m/2} \setminus \{0\}$, this coincides with the global MDP. To evaluate fixed-instance characteristics by input-difference weight, we consider restricted MDPs over sets of low-weight input differences.

For a fixed round key K_i, define the ensemble expected differential probability (EDP), obtained by averaging the differential probability over the matrix-generation distribution, by

$$\mathrm{EDP}_{K_i}(\Delta L_{i-1}, \Delta y) := \mathbb{E}_A\left[\mathrm{DP}_{A,K_i}(\Delta L_{i-1}, \Delta y)\right].$$

For any function $F : \mathbb{F}_2^{m/2} \to \mathbb{F}_2^{m/2}$ with a fixed input-output correspondence and any nonzero difference ΔL_{i-1}, define

$$D_{F,\Delta L_{i-1}}(L_{i-1}) = F(L_{i-1}) \oplus F(L_{i-1} \oplus \Delta L_{i-1}).$$

Then

$$D_{F,\Delta L_{i-1}}(L_{i-1}) = D_{F,\Delta L_{i-1}}(L_{i-1} \oplus \Delta L_{i-1})$$

always holds. Consequently, in the output-difference distribution of a fixed function, the inputs in each pair $\{L_{i-1}, L_{i-1} \oplus \Delta L_{i-1}\}$ produce the same difference. Therefore, for a fixed function, the maximum output differential probability for a given input difference is at least $2^{1-m/2}$.

Accordingly, the uniform differential probability $2^{-m/2}$ after ensemble averaging, derived in Sec.3, is obtained by taking the matrix average first for each output difference. It is not a baseline for the maximum output differential probability of a uniformly random function sampled once and then fixed. Comparisons with a fixed function must distinguish this value from the lower bound $2^{1-m/2}$ induced by the pair symmetry above.

2.1 SIS-type modular-linear component

This subsection defines G_A^{SIS}, which uses an SIS-type modular-linear map as the internal function G_A. The map is based on the compressing transformation used in Ajtai-type SIS hash functions. Given a matrix $A \in \mathbb{Z}_q^{n_S \times m}$, it maps a binary vector $x \in \mathbb{F}_2^m$ by

$$h_A^{\mathrm{SIS}}(x) = Ax \bmod q$$

into $\mathbb{Z}_q^{n_S}$[7, 6]. Each bit of x is interpreted as either 0 or 1 over $\mathbb{Z}_q$. If distinct inputs (x, x') give the same output under this function, then

$$A(x - x') \equiv 0 \pmod{q}.$$

Thus, finding a collision amounts to finding a nonzero short vector $(x - x')$ satisfying this relation, which corresponds to the SIS problem. The function is parameterized by n_S, m, q. We take $q = 2^\ell$, while the dimension-matching condition $n_S\ell = m/2$ is imposed for the present Feistel construction.

Construct the internal function G_A^{SIS} as follows. First, $h_A^{\mathrm{SIS}}(x)$ maps the internal input x to an element of $\mathbb{Z}_q^{n_S}$. The output $\tilde{y} \in \mathbb{Z}_q^{n_S}$ is then binary-encoded componentwise using ℓ bits by $\mathcal{E}_\ell : \mathbb{Z}_q^{n_S} \to \mathbb{F}_2^{m/2}$, yielding y. Thus, the SIS-type internal function is

$$G_A^{\mathrm{SIS}}(x) = \mathcal{E}_\ell(h_A^{\mathrm{SIS}}(x)).$$

Under the parameter assumptions $q = 2^\ell$ and $n_S\ell = m/2$, $\mathcal{E}_\ell$ is a bijection from $\mathbb{Z}_q^{n_S}$ to $\mathbb{F}_2^{m/2}$.

2.2 LWR-type rounded component

LWR[9] replaces the addition of an error term in LWE with an operation that rounds the output of a linear map over $\mathbb{Z}_q$ into $\mathbb{Z}_p$. Consequently, LWR requires no random-error generation, and the resulting maps consist of linear operations and deterministic rounding. Motivated by constructions that combine rounded linear operations with a binary vector as one operand, such as the construction in Ref. [14], we consider the following simplified map.

Let the parameters be $q = 2^\ell$, $p = 2^k$, $s = \ell - k \geq 0$, and $1 \leq k \leq \ell$, and assume $n_L k = m/2$. The main analysis uses nearest rounding with ties rounded down for $s \geq 1$. Define the offset

$$c_{\mathrm{td}} = \begin{cases} 0, & s = 0, \\ 2^{s-1} - 1, & s \geq 1. \end{cases}$$

When $s = 0$, we have $p = q$ and define $\rho_{q,q}(u) = u$ as the identity map. For $s \geq 1$ and $u \in \mathbb{Z}_q$, define

$$\rho_{q,p}(u) = \lfloor ((u + c_{\mathrm{td}}) \bmod q)/2^s \rfloor \in \mathbb{Z}_p.$$

If $u = a2^s + r$ with $0 \le r < 2^s$, this rule rounds to a when $0 \le r \le 2^{s-1}$ and to $a+1$ when $2^{s-1} < r < 2^s$, with the result interpreted in $\mathbb{Z}_p$. For later use, also consider

$$2^s\mathbb{Z}_q = \{\, 2^s u \bmod q : u \in \mathbb{Z}_q \,\} \subset \mathbb{Z}_q$$

and define the quotient map from this set to $\mathbb{Z}_p$,

$$\pi_s : 2^s\mathbb{Z}_q \longrightarrow \mathbb{Z}_p$$

as follows. For $D \in 2^s\mathbb{Z}_q$, let

$$D \equiv 2^s v \pmod{q}$$

and, for the $v \in \mathbb{Z}_p$ satisfying this relation, set

$$\pi_s(D) = v.$$

The element v is uniquely determined in $\mathbb{Z}_p$. In particular, when $s = 0$, π_0 is the identity map.

For a binary vector $x \in \mathbb{F}_2^m$ and a matrix $A \in \mathbb{Z}_q^{n_L \times m}$, define the LWR-type map $h_A^{\mathrm{LWR}} : \mathbb{F}_2^m \to \mathbb{Z}_p^{n_L}$ by

$$h_A^{\mathrm{LWR}}(x) = \rho_{q,p}(Ax \bmod q),$$

where rounding is applied independently to each row component of $Ax \bmod q$. Using the bijection $\mathcal{E}_k : \mathbb{Z}_p^{n_L} \to \mathbb{F}_2^{m/2}$ that binary-encodes each component with k bits, define the LWR-type internal function by

$$G_A^{\mathrm{LWR}}(x) = \mathcal{E}_k(h_A^{\mathrm{LWR}}(x)).$$

We record a basic property of the rounding map that will be used later. Suppose that $u, u' \in \mathbb{Z}_q$ satisfy

$$u - u' \in 2^s\mathbb{Z}_q.$$

Let the standard representatives in $\{0, \ldots, q-1\}$ of $u + c_{\mathrm{td}}$ and $u' + c_{\mathrm{td}}$ be

$$U = (u + c_{\mathrm{td}}) \bmod q, \qquad U' = (u' + c_{\mathrm{td}}) \bmod q.$$

Then

$$U - U' \equiv u - u' \pmod{q}.$$

Since the residue class of $U - U'$ modulo q belongs to $2^s\mathbb{Z}_q$, we have

$$U \equiv U' \pmod{2^s}.$$

Hence, there exist $a, a' \in \{0, \ldots, p-1\}$ and $r \in \{0, \ldots, 2^s - 1\}$ such that

$$U = 2^s a + r, \qquad U' = 2^s a' + r.$$

By construction, $\rho_{q,p}(u) = a$ and $\rho_{q,p}(u') = a'$. Since $U - U' = 2^s(a - a')$ and $U - U' \equiv u - u' \pmod{q}$, the definition of π_s gives

$$\rho_{q,p}(u) - \rho_{q,p}(u') \equiv \pi_s(u - u') \pmod{2^k}.$$

Remark 2. *Each instance of the component classes defined in this section is determined by a matrix A. In an application to a Feistel cipher, one may use the same A throughout the cipher instance or a different A in each round. Because our main results concern local differential characteristics within one round, they apply under either choice.*

We evaluate the class average for A sampled independently and uniformly from the matrix space, while in the fixed-instance setting the same A is fixed as a public parameter. The round key K_i is a secret value independent of A, and the adversary selects an input difference from the candidate set $\mathcal{D}$ after observing A. The certificate conditions do not depend on the generation distribution of A, whereas the distributional analysis of certificate counts assumes that the matrix entries are independent and uniform.

3 Uniformity of output differences in the class average

In this section, we fix a nonzero input difference ΔL_{i-1} independently of the matrix A and analyze the EDP averaged over the matrix-generation distribution. In Eq. (1), denote the differential probabilities of the SIS-type and LWR-type F functions by $\mathrm{DP}^{\mathrm{SIS}}_{A,K_i}(\Delta L_{i-1}, \Delta y)$ and $\mathrm{DP}^{\mathrm{LWR}}_{A,K_i}(\Delta L_{i-1}, \Delta y)$, respectively. Define their averages over A by

$$\mathrm{EDP}^{\mathrm{SIS}}_{K_i}(\Delta L_{i-1}, \Delta y) = \mathbb{E}_A[\mathrm{DP}^{\mathrm{SIS}}_{A,K_i}(\Delta L_{i-1}, \Delta y)],$$

and

$$\mathrm{EDP}^{\mathrm{LWR}}_{K_i}(\Delta L_{i-1}, \Delta y) = \mathbb{E}_A[\mathrm{DP}^{\mathrm{LWR}}_{A,K_i}(\Delta L_{i-1}, \Delta y)].$$

The following theorem then holds for the output difference in the class average.

Theorem 3. *Fix a round key K_i and a nonzero difference $\Delta L_{i-1} \in \mathbb{F}_2^{m/2}$. Let the input difference to the internal function be*

$$\Delta x = \Phi(\Delta L_{i-1}).$$

Then $\Delta x \neq 0$.

For the SIS-type map, take $q = 2^\ell$ and $n_S \ell = m/2$ so that $\mathcal{E}_\ell$ is bijective. If the entries of $A \leftarrow \mathbb{Z}_q^{n_S \times m}$ are sampled independently and uniformly, then for every $\Delta y \in \mathbb{F}_2^{m/2}$,

$$\mathrm{EDP}^{\mathrm{SIS}}_{K_i}(\Delta L_{i-1}, \Delta y) = 2^{-m/2}$$

holds.

For the LWR-type map, take $q = 2^\ell$, $p = 2^k$, $s = \ell - k$, and $n_L k = m/2$ so that $\mathcal{E}_k$ is bijective, and use the nearest-rounding-with-ties-rounded-down map defined above as $\rho_{q,p}$. If the entries of $A \leftarrow \mathbb{Z}_q^{n_L \times m}$ are sampled independently and uniformly, then for every $\Delta y \in \mathbb{F}_2^{m/2}$,

$$\mathrm{EDP}^{\mathrm{LWR}}_{K_i}(\Delta L_{i-1}, \Delta y) = 2^{-m/2}$$

holds.

Proof. For a fixed L_{i-1}, the inputs to the internal function G_A are

$$\begin{aligned} x &= K_i \oplus \Phi(L_{i-1}) \\ x' &= K_i \oplus \Phi(L_{i-1} \oplus \Delta L_{i-1}) \\ &= x \oplus \Phi(\Delta L_{i-1}), \end{aligned}$$

and $\Delta x = \Phi(\Delta L_{i-1})$ is nonzero.

Let x_u and Δx_u denote the bits of x and Δx at position u, respectively, and define the set of nonzero bit positions of Δx by

$$\mathcal{I}(\Delta x) := \{u : \Delta x_u = 1\}.$$

Let a_j be the jth row of A and $a_{j,i}$ its ith entry, and let $\tilde{y} = h_A^{\mathrm{SIS}}(x)$ and $\tilde{y}' = h_A^{\mathrm{SIS}}(x')$. Then the output components $\tilde{y}_j, \tilde{y}'_j$ corresponding to the jth row of A can be written as follows:

$$\begin{aligned} \tilde{y}_j &= \sum_{i \in \mathcal{I}(\Delta x)} a_{j,i} x_i + \sum_{i \notin \mathcal{I}(\Delta x)} a_{j,i} x_i \pmod q, \\ \tilde{y}'_j &= \sum_{i \in \mathcal{I}(\Delta x)} a_{j,i} (x_i \oplus 1) + \sum_{i \notin \mathcal{I}(\Delta x)} a_{j,i} x_i \pmod q. \end{aligned}$$

For any $\tau \in \mathcal{I}(\Delta x)$, we have $x_\tau \neq x'_\tau$ and $x_\tau, x'_\tau \in \{0,1\}$. Thus, the coefficient $a_{j,\tau}$ occurs in exactly one of $\tilde{y}_j$ and $\tilde{y}'_j$. After fixing all coefficients in row j other than $a_{j,\tau}$, we can write

$$(\tilde{y}_j, \tilde{y}'_j) = (a_{j,\tau} + c, c') \quad \text{or} \quad (c, a_{j,\tau} + c') \pmod q,$$

where $c, c' \in \mathbb{Z}_q$ are fixed and $a_{j,\tau}$ is uniform over $\mathbb{Z}_q$.

Consequently, the output difference in row j is

$$\mathcal{E}_\ell(a_{j,\tau} + c) \oplus \mathcal{E}_\ell(c') \quad \text{or} \quad \mathcal{E}_\ell(c) \oplus \mathcal{E}_\ell(a_{j,\tau} + c').$$

Addition over $\mathbb{Z}_q$, the map $\mathcal{E}_\ell : \mathbb{Z}_q \to \mathbb{F}_2^\ell$, and XOR with a fixed vector are all bijections. Hence this output difference is uniform over $\mathbb{F}_2^\ell$. The conclusion holds for every fixed assignment of the coefficients other than $a_{j,\tau}$, so uniformity is preserved after removing the conditioning on those coefficients. If $\{a_{j,i} : i \neq \tau\}$ is fixed simultaneously in every row of A, then $a_{1,\tau}, \ldots, a_{n_S,\tau}$ are mutually independent and uniform over $\mathbb{Z}_q$. Therefore, all n_S encoded row differences are independent and uniform, and the probability of every output difference is $2^{-n_S \ell} = 2^{-m/2}$.

The LWR-type component is handled similarly. When $s = 0$, $p = q$ and $\rho_{q,q}$ is the identity map, so the argument reduces to that for the SIS-type component with output encoding width $k = \ell$. Each row's encoded output difference is therefore uniform over $\mathbb{F}_2^k$. Assume henceforth that $s \geq 1$. The prerounding value of the output component corresponding to row j of A is computed exactly

as in the SIS-type component. Under the same conditioning used in the SIS-type case, for every row j, fixing one prerounding value leaves the other uniform over $\mathbb{Z}_q$.

In the rounding operation, addition of the fixed offset c_{td} is a permutation of $\mathbb{Z}_q$ and therefore preserves uniformity. The map $\rho_{q,p}$ sends $\mathbb{Z}_q$ to $\mathbb{Z}_p$, and every fiber has the same size 2^s. Thus, applying $\rho_{q,p}$ to a uniform prerounding value over $\mathbb{Z}_q$ produces a uniform rounded value over $\mathbb{Z}_p$. Since $\mathcal{E}_k$ is bijective, the XOR difference with the other fixed output is uniform over $\mathbb{F}_2^k$. The coefficients are independent across rows, so the row output differences are also independent, and the probability of any output difference is $(2^{-k})^{n_L} = 2^{-m/2}$.

This argument holds for every fixed input L_{i-1} and every fixed round key K_i. The key K_i only applies the same XOR translation to the two internal inputs and does not affect the internal difference $\Delta x = \Phi(\Delta L_{i-1})$. Therefore, taking probability over L_{i-1} and then averaging over A gives

$$\begin{aligned}\mathrm{EDP}^{\mathrm{SIS}}_{K_i}(\Delta L_{i-1}, \Delta y) &= \mathrm{EDP}^{\mathrm{LWR}}_{K_i}(\Delta L_{i-1}, \Delta y)\\ &= 2^{-m/2},\end{aligned}$$

which proves the theorem. Because this equality holds for every fixed K_i, further averaging over any key distribution independent of A leaves the value equal to $2^{-m/2}$. □

Thus, for the SIS-type and LWR-type F functions, the differential probabilities averaged over the matrix-generation distribution are uniform. This does not imply that the output-difference distribution is uniform for every fixed instance A. As shown in the following sections, a fixed A can induce certificate conditions on output-difference bits. Uniformity in the class average must therefore be distinguished from differential uniformity of a fixed instance.

4 Certificate conditions for fixed instances

In this section, with the matrix A fixed, we derive output-difference bits that are fixed independently of the input value x by conditions on 2-adic valuations. The bit positions fixed by certificates, their values, and the certificate count are the same for every fixed round key K_i. For an output-difference component Δy_j, let $(\Delta y_j)_b$ denote its bit at position b.

As in the proof of Theorem 3, consider a nonzero input difference Δx to the internal function and its set of nonzero bit positions $\mathcal{I}(\Delta x)$, and define its weight by $d = |\mathcal{I}(\Delta x)|$. For $a \in \mathbb{Z}_q$ represented by an integer in $\{0, 1, \ldots, 2^\ell - 1\}$, let

$$\nu_2(a) = \max\{r : 0 \le r \le \ell,\ 2^r \mid a\}.$$

This is the ordinary 2-adic valuation of the standard representative, truncated at ℓ; in particular, for $a = 0$, set $\nu_2(0) = \ell$.

For the jth row of A, define

$$r_j(A, \Delta x) = \min_{i \in \mathcal{I}(\Delta x)} \nu_2(a_{j,i}).$$

This is the largest integer r such that every entry $a_{j,i}$ in row j corresponding to a nonzero component of Δx is divisible by 2^r.

4.1 The SIS-type map

First consider the SIS-type map. Among the bits of the output difference Δy_j corresponding to row j, define the positions whose values are fixed by the 2-adic valuation condition as

$$\mathcal{S}^{\mathrm{SIS}}_{A,j}(\Delta x) = \{\, b \in \{0, \ldots, \ell - 1\} : b \le r_j(A, \Delta x) \,\},$$

and let

$$\mathcal{S}^{\mathrm{SIS}}_{A}(\Delta x) = \{(j, b) : 1 \le j \le n_S,\ b \in \mathcal{S}^{\mathrm{SIS}}_{A,j}(\Delta x)\}.$$

Also define the certificate count in row j by

$$c^{\mathrm{SIS}}_{A,j}(\Delta x) = |\mathcal{S}^{\mathrm{SIS}}_{A,j}(\Delta x)| = \min\{\ell, r_j(A, \Delta x) + 1\}$$

and the certificate count over the entire output difference by

$$C^{\mathrm{SIS}}_{A}(\Delta x) = |\mathcal{S}^{\mathrm{SIS}}_{A}(\Delta x)| = \sum_{j=1}^{n_S} c^{\mathrm{SIS}}_{A,j}(\Delta x).$$

For $(j, b) \in \mathcal{S}^{\mathrm{SIS}}_{A}(\Delta x)$, we have $b \le r_j(A, \Delta x)$, and therefore

$$2^b \mid a_{j,i} \qquad (i \in \mathcal{I}(\Delta x))$$

holds. Consequently,

$$\sum_{i \in \mathcal{I}(\Delta x)} a_{j,i}$$

is divisible by 2^b. Under this condition, define the fixed value of bit b of the output difference by

$$\gamma^{\mathrm{SIS}}_{j,b}(A, \Delta x) = \left(\frac{\sum_{i \in \mathcal{I}(\Delta x)} a_{j,i}}{2^b} \right) \bmod 2.$$

In particular, if $r_j(A, \Delta x) > b$, then $\gamma^{\mathrm{SIS}}_{j,b}(A, \Delta x) = 0$.

Theorem 4. *Fix a row a_j of A, a round key K_i, and a nonzero input difference ΔL_{i-1}. Let $\Delta x = \Phi(\Delta L_{i-1})$, and let Δy_j be the jth output-difference component of the SIS-type F function. Then, for every $(j, b) \in \mathcal{S}^{\mathrm{SIS}}_{A}(\Delta x)$,*

$$(\Delta y_j)_b = \gamma^{\mathrm{SIS}}_{j,b}(A, \Delta x)$$

holds for every input $L_{i-1} \in \mathbb{F}_2^{m/2}$.

Proof. Consider the fixed row a_j of A and the input pair $x, x' = x \oplus \Delta x$ to the internal function. The corresponding pre-encoding components of the SIS-type internal map can be written as

$$\tilde{y}_j \equiv \sum_{i=1}^{m} a_{j,i} x_i \pmod{q}, \ \tilde{y}'_j \equiv \sum_{i=1}^{m} a_{j,i} x'_i \pmod{q}.$$

The corresponding arithmetic difference is

$$\tilde{y}_j - \tilde{y}'_j \equiv \sum_{i \in \mathcal{I}(\Delta x)} a_{j,i}(x_i - x'_i) \pmod{q}.$$

For the indices $i \in \mathcal{I}(\Delta x)$ that occur in this sum, $x_i - x'_i \in \{1, -1\}$. Moreover, for $(j, b) \in \mathcal{S}_A^{\mathrm{SIS}}(\Delta x)$, every $a_{j,i}$ is divisible by 2^b, and

$$a_{j,i} \equiv -a_{j,i} \pmod{2^{b+1}}$$

holds. Thus, modulo 2^{b+1}, the sign $x_i - x'_i \in \{1, -1\}$ disappears, giving

$$\tilde{y}_j - \tilde{y}'_j \equiv \sum_{i \in \mathcal{I}(\Delta x)} a_{j,i} \pmod{2^{b+1}},$$

and hence

$$\tilde{y}_j - \tilde{y}'_j \equiv 2^b \gamma_{j,b}^{\mathrm{SIS}}(A, \Delta x) \pmod{2^{b+1}}$$

holds.

This congruence means that the lower b bits of $\tilde{y}_j$ and $\tilde{y}'_j$ are equal and that the difference at bit position b is fixed to $\gamma_{j,b}^{\mathrm{SIS}}(A, \Delta x)$. Therefore, after binary encoding, the output difference also satisfies

$$(\Delta y_j)_b = \gamma_{j,b}^{\mathrm{SIS}}(A, \Delta x)$$

for every input L_{i-1}. □

It follows from the theorem that, if $r_j(A, \Delta x) = r < \ell$ in row j, then bits $0, \ldots, r-1$ of the SIS-type output difference are fixed to 0, while bit r is fixed to $\gamma_{j,r}^{\mathrm{SIS}}(A, \Delta x)$. The value $\gamma_{j,r}^{\mathrm{SIS}}(A, \Delta x)$ may itself be 0, so the statement does not imply that bit r is necessarily nonzero. If $r = \ell$, all ℓ bits of the output symbol are fixed to 0.

4.2 The LWR-type map

Next consider the LWR-type map. Here $q = 2^\ell$, $p = 2^k$, and $s = \ell - k$, and nearest rounding with ties rounded down computes $\lfloor U/2^s \rfloor$ from the value

$$U = (u + c_{\mathrm{td}}) \bmod q$$

obtained after adding the fixed offset c_{td}. Under the analogous arithmetic-difference calculation, bit b after rounding corresponds to bit $s + b$ over $\mathbb{Z}_q$ after

the offset addition. Among the bits of the output difference corresponding to row j, define the output-bit positions fixed after rounding by the 2-adic valuation condition as

$$\mathcal{S}_{A,j}^{\mathrm{LWR}}(\Delta x) = \{\, b \in \{0, \ldots, k-1\} : s + b \le r_j(A, \Delta x) \,\},$$

and let

$$\mathcal{S}_{A}^{\mathrm{LWR}}(\Delta x) = \{(j, b) : 1 \le j \le n_L,\ b \in \mathcal{S}_{A,j}^{\mathrm{LWR}}(\Delta x)\}.$$

Also define the certificate count in row j by

$$c_{A,j}^{\mathrm{LWR}}(\Delta x) = |\mathcal{S}_{A,j}^{\mathrm{LWR}}(\Delta x)|$$

and the certificate count over the entire output difference by

$$C_{A}^{\mathrm{LWR}}(\Delta x) = |\mathcal{S}_{A}^{\mathrm{LWR}}(\Delta x)| = \sum_{j=1}^{n_L} c_{A,j}^{\mathrm{LWR}}(\Delta x).$$

Then the equality $c_{A,j}^{\mathrm{LWR}}(\Delta x) = 0$ holds when $r_j(A, \Delta x) < s$, and

$$c_{A,j}^{\mathrm{LWR}}(\Delta x) = \min\{k, r_j(A, \Delta x) - s + 1\}$$

holds when $r_j(A, \Delta x) \ge s$.

For $(j, b) \in \mathcal{S}_{A}^{\mathrm{LWR}}(\Delta x)$, we have $s + b \le r_j(A, \Delta x)$, and therefore

$$2^{s+b} \mid a_{j,i} \qquad (i \in \mathcal{I}(\Delta x))$$

holds. Thus, define the fixed bit value after rounding by

$$\gamma_{j,b}^{\mathrm{LWR}}(A, \Delta x) = \left(\frac{\sum_{i \in \mathcal{I}(\Delta x)} a_{j,i}}{2^{s+b}} \right) \bmod 2.$$

Theorem 5. *Fix a row a_j of A, a round key K_i, and a nonzero input difference ΔL_{i-1}. Let $\Delta x = \Phi(\Delta L_{i-1})$, and let Δy_j be the jth output-difference component of the LWR-type F function. Then, for every $(j, b) \in \mathcal{S}_{A}^{\mathrm{LWR}}(\Delta x)$,*

$$(\Delta y_j)_b = \gamma_{j,b}^{\mathrm{LWR}}(A, \Delta x)$$

holds for every input $L_{i-1} \in \mathbb{F}_2^{m/2}$.

Proof. Consider the fixed row a_j of A and the input pair $x, x' = x \oplus \Delta x$ to the internal function.

Write the prerounding values in row j as

$$\tilde{y}_j \equiv \sum_{i=1}^{m} a_{j,i} x_i \pmod q,\ \tilde{y}'_j \equiv \sum_{i=1}^{m} a_{j,i} x'_i \pmod q,$$

and define their arithmetic difference by

$$D_j \equiv \tilde{y}_j - \tilde{y}'_j \pmod q,$$

where D_j is a residue class in $\mathbb{Z}_q$. As in the proof of Theorem 4, considering the congruence modulo 2^{s+b+1} gives

$$D_j \equiv 2^{s+b}\gamma_{j,b}^{\mathrm{LWR}}(A, \Delta x) \pmod{2^{s+b+1}}.$$

In particular, $D_j \in 2^s\mathbb{Z}_q$. For the addition of the offset c_{td} in the rounding operation, let

$$U_j = (\tilde{y}_j + c_{\mathrm{td}}) \bmod q, \qquad U'_j = (\tilde{y}'_j + c_{\mathrm{td}}) \bmod q.$$

Then, for the fixed offset,

$$U_j - U'_j \equiv D_j \pmod{q}$$

holds. Because $D_j \in 2^s\mathbb{Z}_q$, the lower s bits of U_j and U'_j are equal. By the basic property of $\rho_{q,p}$,

$$\rho_{q,p}(\tilde{y}_j) - \rho_{q,p}(\tilde{y}'_j) \equiv \pi_s(D_j) \pmod{2^k}.$$

Dividing the congruence for D_j modulo 2^{s+b+1} by 2^s, equivalently using the definition of π_s, gives

$$\pi_s(D_j) \equiv 2^b\gamma_{j,b}^{\mathrm{LWR}}(A, \Delta x) \pmod{2^{b+1}},$$

and therefore

$$\rho_{q,p}(\tilde{y}_j) - \rho_{q,p}(\tilde{y}'_j) \equiv 2^b\gamma_{j,b}^{\mathrm{LWR}}(A, \Delta x) \pmod{2^{b+1}}.$$

This congruence means that the lower b bits of the two rounded values agree and that bit b either agrees or flips according to $\gamma_{j,b}^{\mathrm{LWR}}(A, \Delta x)$. Consequently, the binary-encoded output difference also satisfies

$$(\Delta y_j)_b = \gamma_{j,b}^{\mathrm{LWR}}(A, \Delta x)$$

for every L_{i-1}. □

In an LWR-type component, the lower $s = \ell - k$ bits do not appear in the output because of rounding. Hence fixed bits arising from the 2-adic condition above appear in the rounded output difference only when the prerounding arithmetic difference is a multiple of 2^s as a residue class. Specifically, if $r_j(A, \Delta x) < s$ in row j, the preceding argument yields no fixed output-difference bit. If $s \le r_j(A, \Delta x) < \ell$, bits $0, \ldots, r_j(A, \Delta x) - s - 1$ of the rounded output difference are fixed to 0, while bit $r_j(A, \Delta x) - s$ is fixed to

$$\gamma_{j,r_j(A,\Delta x)-s}^{\mathrm{LWR}}(A, \Delta x).$$

If $r_j(A, \Delta x) = \ell$, however, all output bits are fixed to 0.

The preceding theorems imply the following properties.

Corollary 6. *For the SIS-type map, every fixed matrix A, round key K_i, and nonzero input difference ΔL_{i-1} satisfy*

$$C_A^{\mathrm{SIS}}(\Phi(\Delta L_{i-1})) \geq n_S.$$

Thus, every fixed SIS-type instance has at least one fixed bit in the output-difference component Δy_j corresponding to each row.

Proof. For the SIS type, $r_j(A, \Delta x) \geq 0$ holds for every row j. Hence $0 \in \mathcal{S}_{A,j}^{\mathrm{SIS}}(\Delta x)$, so every row has at least one certificate. Therefore, $C_A^{\mathrm{SIS}}(\Delta x) \geq n_S$. □

Corollary 7. *For the output-difference component Δy_j corresponding to each row of the LWR-type component, the condition under which the 2-adic valuation criterion yields at least one certificate is*

$$r_j(A, \Delta x) \geq s.$$

This is equivalent to requiring every matrix entry $a_{j,i}$ in row j for which $\Delta x_i = 1$ to be a multiple of 2^s. Thus, if $r_j(A, \Delta x) < s$, the 2-adic valuation condition yields no certificate; only when $r_j(A, \Delta x) \geq s$ does it yield a certificate in the rounded output difference.

Proof. For the LWR type, the condition under which the 2-adic valuation fixes the least significant bit of the rounded output difference is $s \leq r_j(A, \Delta x)$. Therefore, row j has at least one certificate in the rounded output difference if and only if $r_j(A, \Delta x) \geq s$. □

In summary, the SIS type retains a certificate in every row for every nonzero difference, whereas the LWR type produces certificates only in rows satisfying $r_j(A, \Delta x) \geq s$.

5 Fixed-instance certificate bounds and statistical evaluation

The 2-adic valuation condition is sufficient but not necessary and therefore does not completely characterize all fixed bits of an output difference. The certificate count should therefore be viewed as a lower bound on the total number of fixed output-difference bits. In this section, we use the certificate count to derive lower bounds on the maximum output differential probability for a fixed input difference and on the restricted MDP of a fixed instance, and statistically evaluate the resulting lower bound over the matrix-generation distribution.

We now unify the certificate counts of the SIS and LWR types in a single notation. For row j of A, write $r_j = r_j(A, \Delta x)$. An SIS-type output difference can be viewed as n_S symbols represented with ℓ bits each. An LWR-type output difference consists of n_L symbols represented with k bits each. For the LWR

type, each rounded output symbol is obtained from the offset-adjusted representative $U_j = (\tilde{y}_j + c_{\mathrm{td}}) \bmod q$ by extracting the k bits beginning at position s. Let w be the number of bits per row symbol, t the extraction bit position in the corresponding residue representative, and n the number of rows, and write

$$(w, t, n) = \begin{cases} (\ell, 0, n_S), & \text{SIS type,} \\ (k, s, n_L), & \text{LWR type} \end{cases}$$

for the SIS and LWR types.

The per-row certificate counts for the SIS and LWR types are

$$c_{A,j}^{\mathrm{SIS}}(\Delta x) = \min\{\ell, r_j + 1\},$$
$$c_{A,j}^{\mathrm{LWR}}(\Delta x) = \max\{0, \min\{k, r_j - s + 1\}\}.$$

Using w, t, n, write these as

$$c_{A,j}^{(w,t)}(\Delta x) = \max\{0, \min\{w, r_j - t + 1\}\},$$

so that the certificate count over the entire output difference is

$$C_A^{(w,t)}(\Delta x) = \sum_{j=1}^{n} c_{A,j}^{(w,t)}(\Delta x).$$

Let hw denote Hamming weight, and let the Hamming weight d of Δx be an even integer with $2 \le d \le m$. Under the embedding $\Phi(L_{i-1}) = L_{i-1} \| L_{i-1}$, define the set of candidate input differences available to the adversary and having embedded Hamming weight d by

$$\mathcal{D}_d := \{\Delta L_{i-1} \in \mathbb{F}_2^{m/2} : \mathrm{hw}(\Delta L_{i-1}) = d/2\}.$$

Following the preceding definition, write the restricted MDP over this candidate set as

$$\mathrm{MDP}_{A,K_i}^{\mathcal{D}_d} := \max_{\Delta L_{i-1} \in \mathcal{D}_d} p_{\max,A,K_i}(\Delta L_{i-1}).$$

For a matrix A fixed as a public parameter, define the maximum certificate count attainable when the adversary selects an input difference from $\mathcal{D}_d$ by

$$C_A^{\star}(d) := \max_{\Delta L_{i-1} \in \mathcal{D}_d} C_A^{(w,t)}(\Phi(\Delta L_{i-1})),$$

and let

$$\Delta L_A^{\star} \in \arg\max_{\Delta L_{i-1} \in \mathcal{D}_d} C_A^{(w,t)}(\Phi(\Delta L_{i-1}))$$

be a maximizing input difference.

For fixed A and Δx, the quantities $c_{A,j}^{(w,t)}(\Delta x)$ and $C_A^{(w,t)}(\Delta x)$ are uniquely determined. In Sec. 5.1, A is treated as a fixed value; in Sec. 5.2, it is treated as a random variable induced by sampling A.

5.1 Certificate counts and lower bounds on differential probability

Fix a public matrix A and a round key K_i. We first derive a lower bound on $p_{\max,A,K_i}(\Delta L_{i-1})$ for each fixed input difference ΔL_{i-1}, and then maximize over $\mathcal{D}_d$ to obtain a lower bound on the restricted MDP.

For fixed $A, K_i, \Delta L_{i-1}$, define the set of output differences Δy that can occur, namely the support of the map $\Delta y \mapsto \mathrm{DP}_{A,K_i}(\Delta L_{i-1}, \Delta y)$, by

$$\begin{aligned}&\mathrm{Supp}_{A,K_i}(\Delta L_{i-1})\\&= \{\Delta y \in \mathbb{F}_2^{m/2} : \mathrm{DP}_{A,K_i}(\Delta L_{i-1}, \Delta y) > 0\}.\end{aligned}$$

The certificate conditions from the preceding section and the pair symmetry of a fixed function yield an upper bound on the size of this support and a lower bound on $p_{\max,A,K_i}(\Delta L_{i-1})$ for fixed $A, K_i, \Delta L_{i-1}$.

Corollary 8. *For every fixed instance of either an SIS-type or an LWR-type component,*

$$\begin{aligned}&|\,\mathrm{Supp}_{A,K_i}(\Delta L_{i-1})|\\&\le \min\left\{2^{m/2-1}, 2^{m/2-C_A^{(w,t)}(\Phi(\Delta L_{i-1}))}\right\}.\end{aligned}$$

Consequently,

$$p_{\max,A,K_i}(\Delta L_{i-1}) \ge 2^{-m/2+\max\{1,C_A^{(w,t)}(\Phi(\Delta L_{i-1}))\}}.$$

Proof. By the certificate conditions in the preceding section, $C_A^{(w,t)}(\Phi(\Delta L_{i-1}))$ bits of the output difference are fixed independently of the input. These pointwise constraints remain valid when the inputs are restricted to the coset $K_i \oplus \mathrm{Range}(\Phi)$ determined by the fixed key K_i. Therefore, at most $2^{m/2-C_A^{(w,t)}(\Phi(\Delta L_{i-1}))}$ output differences are possible. In addition, the pair symmetry of a fixed function gives

$$D_{F,\Delta L_{i-1}}(L_{i-1}) = D_{F,\Delta L_{i-1}}(L_{i-1} \oplus \Delta L_{i-1}),$$

so only difference values generated by the $2^{m/2-1}$ input pairs can occur. Hence the support size is also at most $2^{m/2-1}$. The two support-size bounds and the pigeonhole principle yield the lower bound on $p_{\max,A,K_i}$. □

In particular, for the SIS-type map, every fixed instance satisfies $C_A^{(w,t)}(\Phi(\Delta L_{i-1})) \ge n_S$, and

$$p_{\max,A,K_i}(\Delta L_{i-1}) \ge 2^{-m/2+n_S} > 2^{-m/2}.$$

These lower bounds hold for every fixed input difference ΔL_{i-1}. Selecting, for the public A, a difference attaining $C_A^\star(d)$ and maximizing this lower bound over $\mathcal{D}_d$ gives the following corollary.

Corollary 9. *For every public matrix A and every fixed round key K_i,*

$$\mathrm{MDP}_{A,K_i}^{\mathcal{D}_d} \ge 2^{-m/2+\max\{1,C_A^\star(d)\}}.$$

Proof. Applying Corollary 8 to $\Delta L_A^\star \in \mathcal{D}_d$ gives the inequalities

$$\begin{aligned}\mathrm{MDP}_{A,K_i}^{\mathcal{D}_d} &\geq p_{\max,A,K_i}(\Delta L_A^\star) \\ &\geq 2^{-m/2+\max\{1,C_A^\star(d)\}},\end{aligned}$$

and proves the claim. □

The fixed-bit relations supplied by certificates can be used not only to lower-bound differential probabilities but also to distinguish the component used in one round from a uniformly random function.

Corollary 10. *Fix a public matrix A and a nonzero input difference ΔL, and suppose that $C \geq 1$ certificates are obtained. Let S be their set of output-bit positions and γ_S the vector of fixed values. Both S and γ_S are computable from $A, \Delta L$ and are independent of K_i. Choose $X_1, \ldots, X_T \in \mathbb{F}_2^{m/2}$ such that the $T \leq 2^{m/2-1}$ input pairs $\{X_r, X_r \oplus \Delta L\}$, $r = 1, \ldots, T$, are pairwise disjoint, and make $2T$ oracle queries. Consider the distinguisher that declares the function to be the constructed function if, for every $r = 1, \ldots, T$,*

$$\big(F(X_r) \oplus F(X_r \oplus \Delta L)\big)\big|_S = \gamma_S.$$

The acceptance probability for the constructed function $F = F_{A,K_i}$ is 1 for every K_i. By contrast, for a uniformly random function $F = R$ chosen independently of A, the acceptance probability is 2^{-CT}.

Proof. For the constructed function, the claim follows from the certificate conditions. For a uniformly random function, the $2T$ inputs are distinct, so the T output differences are independent and uniform, and each difference satisfies the C specified bit conditions with probability 2^{-C}. □

This corollary is a distinguishing result in an oracle model permitting direct queries to the component F_{A,K_i} used in one round. It does not immediately yield a distinguisher or key-recovery attack against the complete multiround Feistel permutation.

5.2 Distribution and numerical evaluation of the maximum certificate count

In this subsection, we sample A from the matrix-generation distribution and statistically evaluate the fixed-instance lower bound obtained in the preceding subsection. We first fix one candidate difference Δx and determine the distribution of $C_A^{(w,t)}(\Delta x)$ over A, and then take the maximum across candidate differences on the same sampled matrix A. Finally, we apply the distribution of this maximum to the restricted-MDP lower bound from the preceding subsection.

As the first step, sample $A \leftarrow \mathbb{Z}_q^{n\times m}$ uniformly, with independent entries. Fix a nonzero difference Δx of even Hamming weight $d \geq 2$, and assume

$$0 \leq t, \qquad 1 \leq w, \qquad t + w \leq \ell.$$

Below, abbreviate $c_j := c_{A,j}^{(w,t)}(\Delta x)$.

For $v \in \{t, \ldots, t+w-1\}$, the condition $r_j(A, \Delta x) \geq v$ means that certificates are obtained through bit $v - t$ of the output difference. If an entry a of A is sampled uniformly over $\mathbb{Z}_q$, the probability that a satisfies this condition is

$$\Pr_{a \leftarrow \mathbb{Z}_q}[\nu_2(a) \geq v] = 2^{-v}$$

and the probability that it holds for all d coefficients in row j on which the fixed Δx acts is

$$\Pr_A[r_j \geq v] = 2^{-vd}.$$

At this stage, $\Delta x, w, t, d$ are fixed and row j of A is sampled at random. Hence c_j is an integer-valued random variable determined by row j of A, and the following probability generating function and moments are obtained.

Theorem 11. *For fixed Δx, the probability generating function of the certificate count c_j determined by row j of A is*

$$\begin{aligned}\mathcal{P}_{w,t,d}(z) &:= \mathbb{E}_A[z^{c_j}] \\ &= 1 - 2^{-dt} \\ &\quad + \sum_{u=1}^{w-1} \left(2^{-d(t+u-1)} - 2^{-d(t+u)}\right) z^u \\ &\quad + 2^{-d(t+w-1)} z^w.\end{aligned}$$

Moreover, because the rows of A are independent, for every $c \in \{0, \ldots, nw\}$ the probability of the total certificate count $C_A^{(w,t)}(\Delta x)$ is

$$\Pr_A[C_A^{(w,t)}(\Delta x) = c] = [z^c]\mathcal{P}_{w,t,d}(z)^n.$$

Proof. By definition,

$$c_j = \max\{0, \min\{w, r_j - t + 1\}\}.$$

For $1 \leq u \leq w$, the event $c_j \geq u$ is equivalent to $r_j \geq t+u-1$, and its probability is $2^{-d(t+u-1)}$. Therefore,

$$\Pr_A[c_j \geq u] = 2^{-d(t+u-1)}.$$

In particular,

$$\Pr_A[c_j = 0] = 1 - \Pr_A[c_j \geq 1] = 1 - 2^{-dt}.$$

For $1 \leq u < w$, moreover,

$$\begin{aligned}\Pr_A[c_j = u] &= \Pr_A[c_j \geq u] - \Pr_A[c_j \geq u+1] \\ &= 2^{-d(t+u-1)} - 2^{-d(t+u)},\end{aligned}$$

and

$$\Pr_A[c_j = w] = \Pr_A[c_j \geq w] = 2^{-d(t+w-1)}.$$

Combining these probabilities yields the stated probability generating function. Because the rows of A are independent,

$$C_A^{(w,t)}(\Delta x) = \sum_{j=1}^{n} c_j$$

has probability generating function $\mathcal{P}_{w,t,d}(z)^n$, from which the probability distribution of $C_A^{(w,t)}(\Delta x)$ follows. □

Corollary 12. *The expectation and the base-2 exponential moment of the certificate count admit the following closed-form expressions:*

$$\mathbb{E}_A[C_A^{(w,t)}(\Delta x)] = n\, 2^{-dt} \frac{1 - 2^{-dw}}{1 - 2^{-d}},$$

and

$$\mathbb{E}_A\left[2^{C_A^{(w,t)}(\Delta x)}\right] = \left(1 + 2^{-dt} \sum_{r=0}^{w-1} 2^{(1-d)r}\right)^n.$$

Proof. Since c_j is an integer-valued random variable between 0 and w, the tail-sum formula gives

$$\begin{aligned}
\mathbb{E}_A[c_j] &= \sum_{u=1}^{w} \Pr_A[c_j \geq u] \\
&= \sum_{u=1}^{w} 2^{-d(t+u-1)} \\
&= 2^{-dt} \frac{1 - 2^{-dw}}{1 - 2^{-d}}.
\end{aligned}$$

Therefore,

$$\mathbb{E}_A[C_A^{(w,t)}(\Delta x)] = n\mathbb{E}_A[c_j].$$

By independence of the rows, moreover,

$$\begin{aligned}
\mathbb{E}_A\left[2^{C_A^{(w,t)}(\Delta x)}\right] &= \mathcal{P}_{w,t,d}(2)^n \\
&= \left(1 + 2^{-dt} \sum_{r=0}^{w-1} 2^{(1-d)r}\right)^n.
\end{aligned}$$

□

We next evaluate the maximum certificate count when A is sampled from the matrix-generation distribution and a candidate difference is selected from $\mathcal{D}_d$ as a function of each public realization of A. As a first step, we determine the probability that a sampled matrix A admits at least one candidate difference

in $\mathcal{D}_d$ having a certificate. Because each realization of A is public, whenever such a candidate difference exists it can be selected from $\mathcal{D}_d$ as a function of A. This probability $P^{\star}_{+,\mathrm{pub}}(d)$ is given as follows.

Proposition 13. *Let $h = d/2$. For a fixed row j and bit position u, define the probability that both corresponding matrix coefficients $a_{j,u}$ and $a_{j,m/2+u}$ have 2-adic valuation at least t by*

$$p_t := 2^{-2t}.$$

Then

$$\begin{aligned} P^{\star}_{+,\mathrm{pub}}(d) &:= \Pr_A[C^{\star}_A(d) > 0] \\ &= 1 - \left[\sum_{r=0}^{h-1} \binom{m/2}{r} p_t^r (1-p_t)^{m/2-r}\right]^n. \end{aligned}$$

Proof. Consider each row j and each bit position $u \in \{1, \ldots, m/2\}$ of the input difference ΔL_{i-1} before embedding. Under the embedding Φ, the matrix coefficients corresponding to bit u are $a_{j,u}$ and $a_{j,m/2+u}$. The probability that both coefficients have 2-adic valuation at least t is

$$p_t = 2^{-2t}$$

as claimed. These events are independent across bit positions u. Therefore, the number of bit positions satisfying the condition in row j follows a binomial distribution with $m/2$ trials and success probability p_t.

Each candidate difference in $\mathcal{D}_d$ selects $h = d/2$ bit positions. Hence the probability that no candidate difference yields a certificate in row j is

$$\sum_{r=0}^{h-1} \binom{m/2}{r} p_t^r (1-p_t)^{m/2-r}.$$

This is the probability that a given row contains no certificate-producing candidate difference. Since the rows are independent, we compute the probability that no candidate difference yields a certificate in any of the n rows, and take the complementary event to obtain the result. □

To obtain the full distribution of $C^{\star}_A(d)$, however, one must handle dependencies among candidate differences determined by the same matrix A. Over the full candidate set $\mathcal{D}_d$, distinct candidate differences may refer to the same columns of A, so their certificate counts are generally not independent. To calculate a maximum distribution exactly, we therefore use a subset $\mathcal{D}_d^{\perp}$ whose members have pairwise disjoint supports. The maximum over this subset gives a lower bound on $C^{\star}_A(d)$ over the full candidate set. Let $h = d/2$, and construct $\mathcal{D}_d^{\perp} \subseteq \mathcal{D}_d$ as a maximum-cardinality subset whose members have pairwise disjoint supports. Each candidate difference uses $h = d/2$ bit positions and no two candidates share these positions. Thus, the maximum number of candidate differences with pairwise disjoint supports is

$$M_d^{\perp} := |\mathcal{D}_d^{\perp}| = \left\lfloor \frac{m/2}{h} \right\rfloor = \left\lfloor \frac{m}{d} \right\rfloor.$$

Define the maximum certificate count over this candidate subset by

$$\widetilde{C}_A^\star(d) := \max_{\Delta L \in \mathcal{D}_d^\perp} C_A^{(w,t)}(\Phi(\Delta L)).$$

Because the matrix entries are independent and uniform, for a fixed internal difference of weight d, the distribution of the total certificate count $C_A^{(w,t)}(\Delta x)$ depends only on d, not on the particular support of Δx. Based on Theorem 11, define the cumulative distribution function for $c \in \{0, \ldots, nw\}$ by

$$\begin{aligned} F_d(c) &:= \Pr_A\left[C_A^{(w,t)}(\Delta x) \le c\right] \\ &= \sum_{r=0}^{c} [z^r]\mathcal{P}_{w,t,d}(z)^n, \end{aligned}$$

and, for convenience, set

$$F_d(-1) := 0.$$

This completes the definition of F_d.

The certificate counts corresponding to the differences $\Delta L \in \mathcal{D}_d^\perp$ depend on pairwise disjoint sets of columns of A, so they are mutually independent and each has cumulative distribution function F_d. The distribution of their maximum is therefore as follows.

Theorem 14. *The cumulative distribution function and probability mass function of $\widetilde{C}_A^\star(d)$ are given by*

$$\Pr_A[\widetilde{C}_A^\star(d) \le c] = F_d(c)^{M_d^\perp},$$

and

$$\Pr_A[\widetilde{C}_A^\star(d) = c] = F_d(c)^{M_d^\perp} - F_d(c-1)^{M_d^\perp},$$

respectively. Moreover,

$$\widetilde{C}_A^\star(d) \le C_A^\star(d).$$

When $d = 2$, $\mathcal{D}_2^\perp = \mathcal{D}_2$ can be chosen, and equality holds.

Proof. The quantity $C_A^{(w,t)}(\Phi(\Delta L_{i-1}))$ depends only on the columns of the matrix corresponding to $\mathrm{Supp}(\Phi(\Delta L_{i-1}))$. The candidate differences in $\mathcal{D}_d^\perp$ have pairwise disjoint supports, and because $\Phi(L_{i-1}) = L_{i-1}\|L_{i-1}$, the sets of matrix columns referenced by distinct candidates are also disjoint. Independence of the matrix entries implies that the certificate counts of the candidates are independent, and each has the cumulative distribution function F_d. Therefore, the cumulative distribution function of the maximum is $F_d(c)^{M_d^\perp}$. The point probabilities follow by taking differences of adjacent cumulative probabilities. Furthermore, since $\mathcal{D}_d^\perp \subseteq \mathcal{D}_d$, we have $\widetilde{C}_A^\star(d) \le C_A^\star(d)$. When $d = 2$, all candidates have distinct supports of weight one, so we may take $\mathcal{D}_2^\perp = \mathcal{D}_2$. □

From the distribution in Theorem 14, we obtain the mean of the auxiliary maximum $\widetilde{C}_A^\star(d)$ and the exponential moment required to average the restricted-MDP lower bound from the preceding subsection over matrices. This mean is given by

$$\begin{aligned}\mu_{\mathrm{pub}}^{\perp}(d) &:= \mathbb{E}_A\left[\widetilde{C}_A^\star(d)\right] \\ &= \sum_{c=1}^{nw}\left(1 - F_d(c-1)^{M_d^{\perp}}\right).\end{aligned}$$

We also define

$$\begin{aligned}\Lambda_{\mathrm{pub}}^{\perp}(d) &:= \log_2 \mathbb{E}_A\left[2^{\max\{1,\widetilde{C}_A^\star(d)\}}\right] \\ &= \log_2 \sum_{c=0}^{nw} 2^{\max\{1,c\}}\left(F_d(c)^{M_d^{\perp}} - F_d(c-1)^{M_d^{\perp}}\right).\end{aligned}$$

The exponent 1 in $\max\{1, \widetilde{C}_A^\star(d)\}$ is already implied by the differential-pair symmetry common to every fixed deterministic function. To isolate the additional exponent of the certificate-based lower bound relative to this generic lower bound, define

$$\Gamma_{\mathrm{pub}}^{\perp}(d) := \Lambda_{\mathrm{pub}}^{\perp}(d) - 1.$$

Since $\max\{1, \widetilde{C}_A^\star(d)\} \geq 1$, we have $\Gamma_{\mathrm{pub}}^{\perp}(d) \geq 0$.

The value $\widetilde{C}_A^\star(d)$ obtained in Theorem 14 is a lower bound on $C_A^\star(d)$. Substituting this lower bound into the restricted-MDP lower bound that holds for each fixed matrix instance and averaging over A yields the following result.

Theorem 15. *Under the independent-uniform matrix model of this subsection, fix an even d with $2 \leq d \leq m$, and let $\mathcal{D}_d$ and $\mathcal{D}_d^{\perp}$ be the candidate sets defined above. For every fixed round key K_i, suppose that a difference is selected from $\mathcal{D}_d$ after observing the public matrix A. Then*

$$\begin{aligned}\mathbb{E}_A[\mathrm{MDP}_{A,K_i}^{\mathcal{D}_d}] &\geq 2^{-m/2+\Lambda_{\mathrm{pub}}^{\perp}(d)} \\ &= 2^{1-m/2+\Gamma_{\mathrm{pub}}^{\perp}(d)} \\ &= 2^{\Gamma_{\mathrm{pub}}^{\perp}(d)}\, 2^{1-m/2}.\end{aligned}$$

Thus, the certificate-based matrix-averaged lower bound is no smaller than the generic fixed-function lower bound $2^{1-m/2}$. Moreover, if

$$\Pr_A[\widetilde{C}_A^\star(d) \geq 2] > 0,$$

then the inequality relative to this generic lower bound is strict.

Proof. By Corollary 9 and

$$\widetilde{C}_A^\star(d) \leq C_A^\star(d),$$

for every realization of A we have

$$\begin{aligned}\mathrm{MDP}_{A,K_i}^{\mathcal{D}_d} &\geq 2^{-m/2+\max\{1,C_A^\star(d)\}} \\ &\geq 2^{-m/2+\max\{1,\widetilde{C}_A^\star(d)\}}.\end{aligned}$$

Therefore,

$$\begin{aligned}\mathbb{E}_A[\mathrm{MDP}_{A,K_i}^{\mathcal{D}_d}] &\geq 2^{-m/2}\mathbb{E}_A\left[2^{\max\{1,\widetilde{C}_A^\star(d)\}}\right] \\ &= 2^{-m/2+\Lambda_{\mathrm{pub}}^\perp(d)} \\ &= 2^{1-m/2+\Gamma_{\mathrm{pub}}^\perp(d)}.\end{aligned}$$

The last two equalities follow from the definitions of $\Lambda_{\mathrm{pub}}^\perp(d)$ and $\Gamma_{\mathrm{pub}}^\perp(d)$. Since $2^{\max\{1,\widetilde{C}_A^\star(d)\}} \geq 2$ for every A, we have $\Gamma_{\mathrm{pub}}^\perp(d) \geq 0$. If $\widetilde{C}_A^\star(d) \geq 2$ with positive probability, then the inequality relative to the generic lower bound is strict. □

Theorem 15 does not compare the class-average EDP with the differential distribution of a fixed instance. Instead, it compares two lower bounds on the same quantity $\mathbb{E}_A[\mathrm{MDP}_{A,K_i}^{\mathcal{D}_d}]$, namely the generic lower bound $2^{1-m/2}$ common to fixed functions and the certificate-based lower bound. Accordingly, $\Gamma_{\mathrm{pub}}^\perp(d)$ quantifies, in bits, the base-2 logarithm of the ratio of the certificate-based lower bound to the generic lower bound. It does not represent the ratio of the actual restricted MDPs, a gap between individual instances, or a security loss.

We compute the numerical values of the following three complementary statistics. The quantity $P_{+,\mathrm{pub}}^\star(d)$ is the probability that the full candidate set contains at least one difference with a certificate. The quantity $\mu_{\mathrm{pub}}^\perp(d)$ is the mean maximum certificate count selected over a candidate set with pairwise disjoint supports. The quantity $\Gamma_{\mathrm{pub}}^\perp(d)$ is the base-2 logarithmic difference between the certificate-based matrix-averaged restricted-MDP lower bound and the generic lower bound induced by differential-pair symmetry. It is computed directly from the preceding definitions as $\Gamma_{\mathrm{pub}}^\perp(d) = \Lambda_{\mathrm{pub}}^\perp(d) - 1$. We consider $m \in \{128, 256\}$ and $d = 2, 4$, corresponding to input-difference weights $\mathrm{hw}(\Delta L) = 1, 2$.

Table 1 presents the mean maximum certificate count and the difference in base-2 exponents between the two lower bounds, $\Gamma_{\mathrm{pub}}^\perp(d)$, for SIS-type parameters $\ell \in \{4, 8, 16\}$.

Table 1: Maximum certificate statistics and lower-bound exponent differences for public SIS-type components.

m	$(\ell, n_S, d, M_d^{\perp})$	$\mu_{\mathrm{pub}}^{\perp}$	$\Gamma_{\mathrm{pub}}^{\perp}$
128	$(4, 16, 2, 64)$	28.16086	28.43184
128	$(4, 16, 4, 32)$	19.80029	19.15428
128	$(8, 8, 2, 64)$	16.26428	16.39599
128	$(8, 8, 4, 32)$	10.59017	9.83349
128	$(16, 4, 2, 64)$	9.60590	9.38824
128	$(16, 4, 4, 32)$	5.81658	4.99992
256	$(4, 32, 2, 128)$	53.07022	54.32950
256	$(4, 32, 4, 64)$	38.41849	37.93870
256	$(8, 16, 2, 128)$	29.78353	30.57196
256	$(8, 16, 4, 64)$	20.29136	19.62990
256	$(16, 8, 2, 128)$	17.07530	17.19879
256	$(16, 8, 4, 64)$	10.98232	10.21742

Table 2: Certificate occurrence and lower-bound exponent differences for public LWR-type components.

(a) Certificate occurrence probability over the full candidate set.

m	(k, s, n_L, d)	$P_{+,\mathrm{pub}}^{\star}$
128	$(2, 2, 32, 2)$	$1 - 3.95525 \times 10^{-58}$
128	$(2, 2, 32, 4)$	$1 - 4.85647 \times 10^{-35}$
128	$(4, 4, 16, 2)$	0.981827
128	$(4, 4, 16, 4)$	0.346225
256	$(2, 2, 64, 2)$	$1 - 2.44735 \times 10^{-230}$
256	$(2, 2, 64, 4)$	$1 - 1.14911 \times 10^{-167}$
256	$(4, 4, 32, 2)$	0.999999891
256	$(4, 4, 32, 4)$	0.950937

(b) Mean maximum certificate count and lower-bound exponent difference over a disjoint-support candidate set.

m	$(k, s, n_L, d, M_d^{\perp})$	$\mu_{\mathrm{pub}}^{\perp}$	$\Gamma_{\mathrm{pub}}^{\perp}$
128	$(2, 2, 32, 2, 64)$	7.71272	7.49696
128	$(2, 2, 32, 4, 32)$	1.38024	0.52761
128	$(4, 4, 16, 2, 64)$	1.96886	1.34173
128	$(4, 4, 16, 4, 32)$	0.00830360	0.00080
256	$(2, 2, 64, 2, 128)$	12.91312	13.18404
256	$(2, 2, 64, 4, 64)$	2.23211	1.40172
256	$(4, 4, 32, 2, 128)$	3.09922	2.44647
256	$(4, 4, 32, 4, 64)$	0.0328553	0.00322

For the SIS type, every output symbol has a certificate for every nonzero

input difference, so $P^{\star}_{+,\mathrm{pub}} = 1$. Here $M^{\perp}_{d} = \lfloor m/d \rfloor$. For $d = 2$, $\mathcal{D}^{\perp}_{2} = \mathcal{D}_{2}$, and the tabulated values are exact over the full candidate set. For $d = 4$, the tabulated values are exact for the disjoint-support maximum $\widetilde{C}^{\star}_{A}(d)$; since $\widetilde{C}^{\star}_{A}(d) \le C^{\star}_{A}(d)$ pointwise, they provide lower bounds on the corresponding full-candidate-set statistics.

For the parameter sets in Table 1, both $\mu^{\perp}_{\mathrm{pub}}$ and $\Gamma^{\perp}_{\mathrm{pub}}$ decrease as ℓ or d increases. Increasing d makes a shared high 2-adic valuation less likely and reduces $M^{\perp}_{d} = \lfloor m/d \rfloor$, the size of the disjoint-support candidate subset; it does not in general reduce the full candidate set $\mathcal{D}_{d}$. The dependence on ℓ results from the simultaneous increase in $w = \ell$ and decrease in $n_S = m/(2\ell)$. For the parameter sets considered here, their combined effect produces the decreasing trend shown in Table 1. For example, when $m = 128$ and $d = 2$, increasing ℓ from 4 to 16 reduces $\Gamma^{\perp}_{\mathrm{pub}}$ from 28.43184 to 9.38824. The corresponding certificate-based lower bound changes from $2^{-34.56816}$ to $2^{-53.61176}$, whereas the generic fixed-function lower bound is 2^{-63} in both cases.

Table 2 presents the certificate occurrence probabilities and the corresponding certificate-count and lower-bound statistics for the LWR type.

The value $P^{\star}_{+,\mathrm{pub}}$ in Table 2(a) is the probability that the full candidate set $\mathcal{D}_{d}$ contains at least one candidate with a certificate. This occurrence probability is relevant to the availability of a certificate-based distinguisher, but it does not by itself imply that the certificate-based restricted-MDP lower bound is strictly larger than the generic pair-symmetry lower bound. When the certificate count is one, the differential-probability lower bound obtained from the support size coincides with the generic pair-symmetry lower bound. For $d = 2$, $\mathcal{D}^{\perp}_{2} = \mathcal{D}_{2}$, and the values $\mu^{\perp}_{\mathrm{pub}}$ and $\Gamma^{\perp}_{\mathrm{pub}}$ in part (b) are exact over the full candidate set. For $d = 4$, the tabulated values are exact for the disjoint-support maximum $\widetilde{C}^{\star}_{A}(d)$; since $\widetilde{C}^{\star}_{A}(d) \le C^{\star}_{A}(d)$ pointwise, they provide lower bounds on the corresponding full-candidate-set statistics.

This distinction is clear for $m = 128$, $(k, s, n_L) = (4, 4, 16)$, and $d = 4$. Although $P^{\star}_{+,\mathrm{pub}} = 0.346225$, we have $\Gamma^{\perp}_{\mathrm{pub}} = 0.00080$, so the base-2 logarithmic difference between the two lower bounds is nearly zero. For $d = 2$ and $(k, s, n_L) = (4, 4, 16)$,

$$(P^{\star}_{+,\mathrm{pub}}, \mu^{\perp}_{\mathrm{pub}}, \Gamma^{\perp}_{\mathrm{pub}}) = (0.981827, 1.96886, 1.34173),$$

so the certificate-based matrix-averaged lower bound is approximately $2^{1.34173} \approx 2.53$ times the generic lower bound. For $d = 2$ and $(k, s, n_L) = (2, 2, 32)$, the corresponding tuple is

$$(1 - 3.95525 \times 10^{-58}, 7.71272, 7.49696),$$

and the certificate-based matrix-averaged lower bound is approximately $2^{7.49696} \approx 180.6$ times the generic lower bound.

These statistics do not give exact restricted MDPs. Rather, they characterize certificate occurrence, certificate counts, and the base-2 logarithmic difference between the certificate-based and generic lower bounds. The distributional

evaluation also assumes independent uniform matrix entries and therefore does not directly apply to random matrices drawn from other distributions or to structured matrices.

6 Conclusion

We analyzed differential characteristics in one Feistel round whose internal components are binary-encoded modular-linear maps or rounded linear maps. For any fixed round key and any nonzero input difference chosen independently of the matrix, the output-difference distribution obtained by averaging over the matrix-generation distribution is uniform. For a fixed public matrix, certificate conditions restrict the output-difference space and yield lower bounds on the restricted MDP. Using the probability generating function for one input difference, we derived the distribution of the maximum certificate count over a disjoint-support subset and quantified the resulting matrix-averaged lower bound and its difference from the generic pair-symmetry lower bound. The analysis is limited to differential behavior within one round and does not directly provide a distinguisher or key-recovery attack against the complete multiround Feistel permutation. Extending the constraints across multiple rounds and analyzing structured matrices remain topics for future work.

Acknowledgments

This work was supported by JSPS KAKENHI Grant Number 25K07752.

References

[1] E. Biham and A. Shamir. Differential cryptanalysis of DES-like cryptosystems. *J. Cryptol.*, 4:3–72, 1991.

[2] K. Nyberg and L. R. Knudsen. Provable security against a differential attack. *J. Cryptol.*, 8:27–37, 1995.

[3] X. Lai, J. L. Massey, and S. Murphy. Markov ciphers and differential cryptanalysis. In *Advances in Cryptology – EUROCRYPT '91*, volume 547 of *Lecture Notes in Computer Science*, pages 17–38, Berlin, 1991. Springer.

[4] J. Daemen, M. Lamberger, N. Pramstaller, V. Rijmen, and F. Vercauteren. Computational aspects of the expected differential probability of 4-round AES and AES-like ciphers. *Computing*, 85:85–104, 2009.

[5] C. Peikert. A decade of lattice cryptography. *Found. Trends Theor. Comput. Sci.*, 10(4):283–424, 2016.

[6] D. Micciancio and O. Regev. Worst-case to average-case reductions based on Gaussian measures. *SIAM J. Comput.*, 37(1):267–302, 2007.

[7] M. Ajtai. Generating hard instances of lattice problems (extended abstract). In *Proceedings of the 28th Annual ACM Symposium on Theory of Computing (STOC 1996)*, pages 99–108, New York, NY, USA, 1996. Association for Computing Machinery.

[8] O. Regev. On lattices, learning with errors, random linear codes, and cryptography. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC 2005)*, pages 84–93, New York, NY, USA, 2005. Association for Computing Machinery.

[9] A. Banerjee, C. Peikert, and A. Rosen. Pseudorandom functions and lattices. In *Advances in Cryptology – EUROCRYPT 2012*, volume 7237 of *Lecture Notes in Computer Science*, pages 719–737, Berlin, 2012. Springer.

[10] J. Alwen, S. Krenn, K. Pietrzak, and D. Wichs. Learning with rounding, revisited. In *Advances in Cryptology – CRYPTO 2013, Part I*, volume 8042 of *Lecture Notes in Computer Science*, pages 57–74, Berlin, 2013. Springer.

[11] M. Luby and C. Rackoff. How to construct pseudorandom permutations from pseudorandom functions. *SIAM J. Comput.*, 17(2):373–386, 1988.

[12] H. Lipmaa, J. Wallén, and P. Dumas. On the additive differential probability of exclusive-or. In *Fast Software Encryption – FSE 2004*, volume 3017 of *Lecture Notes in Computer Science*, pages 317–331, Berlin, 2004. Springer.

[13] S. Paul and B. Preneel. Solving systems of differential equations of addition. In *Information Security and Privacy – ACISP 2005*, volume 3574 of *Lecture Notes in Computer Science*, pages 75–88, Berlin, 2005. Springer.

[14] A. Davidson, A. Deo, and L. Tremblay Thibault. Pool: A practical OT-based OPRF from learning with rounding. In *Proceedings of the 2025 ACM SIGSAC Conference on Computer and Communications Security (CCS 2025)*, pages 1038–1052, New York, NY, USA, 2025. Association for Computing Machinery.