Weak Bijective Quadratic Functions over $\mathbb{F}_p^n$

Applications to MPC-/FHE-/ZK-Friendly Symmetric Primitives

Lorenzo Grassi

Digital Security Group, Radboud University, The Netherlands
l.grassi@science.ru.nl

Abstract. Motivated by new applications such as secure Multi-Party Computation (MPC), Homomorphic Encryption (HE), and Zero-Knowledge proofs (ZK), many MPC-, HE- and ZK-friendly symmetric-key primitives that minimize the number of multiplications over $\mathbb{F}_p$ for a large prime p have been recently proposed in the literature. These symmetric primitive are usually defined via invertible functions, including (i) Feistel and Lai–Massey schemes and (ii) SPN constructions instantiated with invertible non-linear S-Boxes (as invertible power maps $x \mapsto x^d$). However, the "invertibility" property is actually never required in any of the mentioned applications. In this paper, we discuss the possibility to set up MPC-/HE-/ZK-friendly symmetric primitives instantiated with non-invertible *weak bijective* functions. With respect to one-to-one correspondence functions, any output of a weak bijective function admits at most two pre-images. The simplest example of such function is the square map over $\mathbb{F}_p$ for a prime $p \geq 3$, for which $x^2 = (-x)^2$. When working over $\mathbb{F}_p^n$ for $n \gg 1$, a weak bijective function can be set up by re-considering the recent results of Grassi, Onofri, Pedicini and Sozzi as starting point. Given a quadratic local map $F : \mathbb{F}_p^2 \to \mathbb{F}_p$, they proved that the non-linear function over $\mathbb{F}_p^n$ for $n \geq 3$ defined as $\mathcal{S}_F(x_0, x_1, \ldots, x_{n-1}) = y_0 \| y_1 \| \ldots \| y_{n-1}$ where $y_i := F(x_i, x_{i+1})$ is never invertible. Here, we prove that

- the quadratic function $F : \mathbb{F}_p^2 \to \mathbb{F}_p$ that minimizes the probability of having a collision for $\mathcal{S}_F$ over $\mathbb{F}_p^n$ is of the form $F(x_0, x_1) = x_0^2 + x_1$ (or equivalent);
- the function $\mathcal{S}_F$ over $\mathbb{F}_p^n$ defined as before via $F(x_0, x_1) = x_0^2 + x_1$ (or equivalent) is weak bijective.

As concrete applications, we propose modified versions of the MPC–friendly schemes MiMC, HadesMiMC, and (partially of) Hydra, and of the HE–friendly schemes Masta, Pasta, and `Rubato`. By instantiating them with the weak bijective quadratic functions proposed in this paper, we are able to improve the security and/or the performances in the target applications/protocols.

Keywords: Weak Bijective Functions · Local Maps · Quadratic Functions

Contents

1 Introduction

Almost all the symmetric primitives published in the literature – including ciphers, Pseudo-Random Functions/Permutations (PRFs/PRPs), hash functions – are typically designed by iterating an efficiently implementable round function a sufficient number of times such that the resulting composition satisfies the security requirements. Even if not strictly necessary in many scenarios (e.g., stream ciphers, hash functions, and so on), the round function is usually invertible, that is, it is instantiated either via invertible components, or in such a way that, even if the components are not invertible by their own, the overall round function is invertible (as in the case of Feistel or Lai–Massey schemes).

In many cases, this choice is crucial for guarantee security or/and for simplifying the security analysis. As a concrete example, consider the case of a Substitution–Permutation Network (SPN), in which the non-linear layer is instantiated via a concatenation of independent S-Boxes, e.g., $(x_0, x_1, \ldots, x_{n-1}) \mapsto (S(x_0), S(x_1), \ldots, S(x_{n-1}))$ for a certain non-linear function S over a field $\mathbb{F}_q$ for a *small* q (as $q \leq 2^8$). If S is not invertible, finding a collision at the output of any single function S could potentially allow the attacker to break the entire scheme. E.g., Keccak [BDPA13] instantiated with 5-bit *non-invertible* S-Boxes may be easily broken by using a brute force approach, by looking for a collision at the output of the first round set up by considering input messages that active a single S-Box. Similar conclusion holds for e.g. AES [DR02b] instantiated with 8-bit *non-invertible* S-Boxes and used e.g. in counter (CTR) mode.

This problem potentially does not arise in the case of symmetric primitives defined over a field $\mathbb{F}_q$ for a *huge* q, as symmetric primitives targetting Secure Multi-Party Computation (MPC), Homomorphic Encryption (HE), and Zero-Knowledge proofs (ZK), which operate over $\mathbb{F}_p^n$ for a large prime $p \gg 3$ (usually, p is of order $2^{64}, 2^{128}$ or even bigger). However, almost all the MPC-/FHE-/ZK-friendly symmetric cryptographic primitives that have been recently proposed in the literature to minimize the number of field non-linear operations in their natural algorithmic description, often referred to as the *multiplicative complexity*,

Table 1: Current scenario of invertible non-linear round functions that instantiate MPC-/FHE-/ZK-friendly symmetric primitives over $\mathbb{F}_p^n$ proposed in the literature. Note that some primitives are instantiated via several non-linear functions ("Others" include `Horst` schemes and look-up tables).

Symmetric Primitive	Invertible Power Map(s)	Non-Invertible Function(s) in Feistel/Lai–Massey Scheme	Others
MiMC [AGR+16]	✓		
GMiMC [AGP+19]	✓		
HadesMiMC [GLR+20]	✓		
Rescue [AAB+20]	✓		
Poseidon [GKR+21]	✓		
Ciminion [DGGK21]		✓	
Grendel [Sze21]	✓		
Pasta [DGH+21]	✓	✓	
`Reinforced Concrete` [GKL+21]	✓		✓
Neptune [GOPS22]	✓	✓	
Griffin [GHR+22]	✓		✓
Chaghri [AMT22]	✓		
Hydra [GØWS22]	✓	✓	
Anemoi [BBC+22]	✓	✓	

are instantiated via invertible round functions only. (From now on, we also use the term "$\mathbb{F}_p$-multiplication" – or simply, "multiplication" – to refer to a non-linear operation over $\mathbb{F}_p$. Moreover, we do not make any distinction between a $\mathbb{F}_p$-multiplication and a $\mathbb{F}_p$–square operation, since – to the best of our knowledge – they have the same cost in the considered applications/protocols.) The current scenario is indeed summarized in Table 1 (to the best of our knowedlge, the only scheme instantiated via non-invertible components is the HE–friendly scheme Masta [HKC+20], which is based on the Rasta design strategy discussed in Sect. 6). At the same time:

1. *the invertibility property is not required neither in MPC- and HE-applications nor in ZK protocols.* Indeed, MPC and HE applications require a PRF scheme (which is not invertible in general), while ZK protocols requires a hash function (which is not invertible by definition);

2. even if the mentioned block ciphers are invertible, the inverse is actually never used in practice.

Regarding this second point, let's consider the block ciphers MiMC and HadesMiMC, instantiated via the power map $x \mapsto x^d$ defined over $\mathbb{F}_p$ for $d \geq 3$ such that $\gcd(d, p-1) = 1$ in order to guarantee invertibility. Its inverse is again a power map of the form $x \mapsto x^{d'}$, where $d' \geq 3$ is the smallest integer such that $d \cdot d' - 1$ is a multiple of $p - 1$ (we recall that $x^{p-1} = 1$ for each $x \in \mathbb{F}_p \setminus \{0\}$ due to Fermat's little theorem). For small value of d, it follows that d' is of the same order of magnitude of p, or in other words, d' is in general much bigger than d. As a result, computing $x \mapsto x^{d'}$ is in general much more expensive than computing $x \mapsto x^d$, which implies that decrypting is much more expensive than encrypting, a property that is in general **not** desirable in practical use cases. For this reason, MiMC's and HadesMiMC's designers suggest to use such schemes in a mode of operation in which the inverse is not needed (as the CTR mode). Citing MiMC's and HadesMiMC's designers (see e.g. [AGR+16, Sect. 1]), "*[...] decryption is much more expensive than encryption. Using modes where the inverse is not needed is thus advisable.*".

Hence, natural question arise: since the invertibility property is not required for MPC-/HE-/ZK-applications, what happens when considering a symmetric scheme instantiated with non-invertible round functions? Can we still guarantee security, and at the same time decreases the multiplicative complexity? In order to answer these questions, in this paper we start a research regarding quadratic *non-invertible* functions over $\mathbb{F}_p^n$ that can be used as building blocks in MPC-/HE-/ZK-friendly symmetric primitives.

1.1 Weak Bijective Functions constructed via Local Maps

When working with a non-invertible function, an estimation of the probability of the collision event is paramount. From this point of view, it would be desirable that the used non-invertible function admits a fixed *maximum number of preimages for each possible output.*

Weak Bijective Functions. In Sect. 2, we introduce the concept of *"weak bijective" functions*. While a bijective function has a unique preimage for each output, a "weak bijective" function admits at most two preimages for each possible output. The simplest example of weak-bijective function over $\mathbb{F}_p$ is the square map, for which $(+x)^2 = (-x)^2$. Instead, e.g., $x \mapsto x^3 - x$ is not weak bijective (since $0^3 - 0 = 1^3 - 1 = (-1)^3 - (-1) = 0$). In there, we also propose an estimation of the probability that a collision can occur for an *iterated weak bijective functions*, that is, a function defined as the composition of weak bijective functions.

Our Contribution: Weak Bijective Functions over $\mathbb{F}_p^n$. When working over $\mathbb{F}_p^n$, the concatenation of n independent weak bijective function – as the square maps $(x_0, x_1, \dots, x_{n-1}) \mapsto (x_0^2, x_1^2, \dots, x_{n-1}^2)$ – is not a weak bijective function anymore. E.g., for $n = 2$, it is easy to observe that $\big((+x)^2, (+y)^2\big) = \big((-x)^2, (+y)^2\big) = \big((+x)^2, (-y)^2\big) = \big((-x)^2, (-y)^2\big)$.

In order to set up weak-bijective functions, one can potentially consider a generic non-linear function $\mathcal{S}$ over $\mathbb{F}_p^n$ of the form $\mathcal{S}(x) := F_0(x) \| F_1(x) \| \dots \| F_{n-1}(x)$, where $\cdot\|\cdot$ denotes concatenation and where $F_0, F_1, \dots, F_{n-1} : \mathbb{F}_p^n \to \mathbb{F}_p$ are potentially distinct functions. Here, we decided to focus on the case $\mathcal{S}(x) = y_0 \| y_1 \| \dots \| y_{n-1} \in \mathbb{F}_p^n$ where each value $y_i \in \mathbb{F}_p$ is specified according to a single *local* map $F : \mathbb{F}_p^m \to \mathbb{F}_p$ for a certain $m \leq n$. More formally:

Definition 1. Let $p \geq 3$ be a prime integer. Let $1 \leq m \leq n$, and let $F : \mathbb{F}_p^m \to \mathbb{F}_p$ be a non-linear function. The Shift-Invariant (m, n)–lifting function $\mathcal{S}_F$ over $\mathbb{F}_p^n$ induced by F is defined as

$$\mathcal{S}_F(x_0, x_1, \dots, x_{n-1}) := y_0 \| y_1 \| \dots \| y_{n-1} \tag{1}$$

where

$$y_i = F(x_i, x_{i+1}, \dots, x_{i+m-1}) \tag{2}$$

for each $i \in \{0, 1, \dots, n-1\}$, where the sub-indexes are taken modulo n.

For simplicity, we usually make used of the abbreviation "SI-lifting" function $\mathcal{S}_F$. The reason of this choice is related to the recent results presented by Grassi et al. [GOPS22] at FSE/ToSC 2022. In there, authors proved that given any quadratic function $F : \mathbb{F}_p^2 \to \mathbb{F}_p$, the corresponding function $\mathcal{S}_F$ over $\mathbb{F}_p^n$ for $n \geq 3$ as defined in Def. 1 is never invertible. A similar result holds when considering quadratic functions $F : \mathbb{F}_p^3 \to \mathbb{F}_p$ and the corresponding function $\mathcal{S}_F$ over $\mathbb{F}_p^n$ for $n \geq 5$.

By re-considering such functions, in Sect. 3 we prove that the function $\mathcal{S}_F$ over $\mathbb{F}_p^n$ for $n \geq 3$ defined via $F(x_0, x_1) = x_0^2 + \alpha \cdot x_0 + \beta \cdot x_1$ with $\beta \neq 0$ (or equivalent) is a weak-bijective function. In particular:

- the probability that a collision occurs for $\mathcal{S}_F$ over $\mathbb{F}_p^n$ for $n \geq 3$ induced by $F(x_0, x_1) = x_0^2 + \alpha \cdot x_0 + \beta \cdot x_1$ with $\beta \neq 0$ (or equivalent) is upper bounded by p^{-n}.
- if $F(x_0, x_1) = x_0^2 + \alpha \cdot x_0 + \beta \cdot x_1$ with $\beta \neq 0$ (or equivalent), a collision $\mathcal{S}_F(x) = \mathcal{S}_F(y)$ occurs only in the case in which $x_i \neq y_i$ for each $i \in \{0, 1, \dots, n-1\}$ (where z_i denotes the i-th component of $z \in \mathbb{F}_p^n$);

- among all functions $F : \mathbb{F}_p^2 \to \mathbb{F}_p$ for which $\mathcal{S}_F$ can be computed via n multiplications only, $F(x_0, x_1) = x_0^2 + \alpha \cdot x_0 + \beta \cdot x_1$ with $\beta \neq 0$ (or equivalent) is the function that minimizes the probability that a collision occurs for $\mathcal{S}_F$over $\mathbb{F}_p^n$ for $n \geq 3$.

Besides that, we point out that the function $\mathcal{S}_F$ over $\mathbb{F}_p^n$ instantiated via $F(x_0, x_1) = x_0^2 + \alpha \cdot x_0 + \beta \cdot x_1$ can be computed via n $\mathbb{F}_p$-multiplications only.

1.2 Impact on MPC-/ZK-/HE-Friendly PRFs

Even if the function $\mathcal{S}_F$ over $\mathbb{F}_p^n$ instantiated via $F(x_0, x_1) = x_0^2 + \alpha \cdot x_0 + \beta \cdot x_1$ is not invertible, it is suitable for instantiating a non-invertible symmetric primitive, due to its several benefit properties just listed. For this reason, we re-consider some symmetric primitives proposed in the literature for MPC and HE applications, and we propose variants of them instantiated with the weak bijective function just proposed. As we are going to show, this modification allows to get better results in terms of performance and/or of security. To better understand the performance improvements, we recall the following:

- MPC protocols allow several parties to jointly compute a function over their inputs, without exposing these inputs. In the most common case in which MPC protocols are evaluated via linearly homomorphic secret sharing scheme, multiplications require communication between the parties, while affine operations can be evaluated locally. In such a case, the MPC cost metric is mainly related to the number of multiplications needed to evaluate the symmetric scheme;

- HE protocols allow a user to operate on encrypted data without decrypting them. With respect to MPC applications, the cost metric in FHE applications is related to the depth of the circuit to be evaluated.

In the following, we summarize the MPC–/HE–friendly schemes considered in our analysis, and the results we are going to present. We point out that we do not exclude the possibility to adapt ZK–friendly schemes proposed in the literature by making used of a similar approach.

`MiMC++`. MiMC is an iterative Even-Mansour scheme, whose round function is instantiated via the invertible power map $x \mapsto x^d$. As a simple concrete example of a scheme instantiated with a weak bijective function, in Sect. 4 we propose the PRF `MiMC++`, a version of MiMC in which the invertible power map is replaced by the square one $x \mapsto x^2$. In order to guarantee the same security level of MiMC, the size p' of the field $\mathbb{F}_{p'}$ over which `MiMC++` operates must be triple with respect to the one used in MiMC, that is, $p' \approx p^3$ (where p is the prime number that defines the field of MiMC). At the same time, replacing $x \mapsto x^d$ with $x \mapsto x^2$ allows to decrease the multiplicative complexity, e.g., of a factor 27.5% for a security level of 128 bits (where $p \approx 2^{128}$ and $p' \approx 2^{384}$).

Pluto and Hydra++. In order to guarantee the security of the PRF `MiMC++`, we are forced to work with a prime p' that is much larger than the security level. This problem does not arise when working with e.g. the cipher HadesMiMC defined over $\mathbb{F}_p^n$. In such a case, given a certain security level and a fixed prime p, it is possible to achieve security by working on the parameter n.

The main characteristic of the Hades design strategy [GLR+20] is the uneven distribution of the S-Boxes through the rounds. The external rounds are instantiated with a full S-Box layer that allows to provide security against the statistical attacks, while the internal rounds are instantiated with a partial S-Box layer, in order to increase the overall degree of the scheme minimizing the cost. In the case of HadesMiMC, the S-Box is instantiated with an invertible power map $x \mapsto x^d$, and the linear layer via the multiplication with a MDS matrix.

The body of the recent PRF Hydra is based on the Hades design strategy. The PRF Hydra is a concrete instance of the Megafono mode of operation, a modified version of the Farfalle mode of operation [BDH+17] suitable for MPC applications. As in the case of Farfalle and of Ciminion, a scheme based on the Megafono mode of operation is composed of two phases:

- an initial phase, in which the input is mixed with the secret key via a PRP;
- an expansion phase, in which the state is expanded until the desired state size is reached.

As just pointed out, the permutation that defines the initial phase is inspired by the Hades design strategy. However, instead of having an uneven distribution of the S-Boxes as in HadesMiMC, Hydra's designers proposed two different round functions, one for the internal rounds and one for the external ones, so that

- the external full rounds are instantiated via power maps as in HadesMiMC;
- the internal partial rounds are instantiated via a generalized Lai–Massey construction.

The new PRF Pluto takes inspiration of the body of Hydra, but the power maps in the external rounds are replaced by the SI–lifting function $\mathcal{S}_F$ over $\mathbb{F}_p^n$ induced by $F(x_0, x_1) = x_0^2 + x_1$. As we show in Sect. 5, this modified version achieves (much) better performance with respect to HadesMiMC in term of multiplicative complexity, especially in the case of large $n \gg 1$. In a similar way, when replacing the body of Hydra with the keyed PRF Pluto, the multiplicative complexity of the modified PRF Hydra++ is (slightly) reduced.

Masta, Pasta, and `Rubato`. Finally, we re-consider the HE-friendly PRFs Masta, Pasta, and `Rubato`. In order to minimize the depth, these symmetric primitives are based on the design strategy initially proposed for the HE-friendly PRF Rasta, that is, (i) they are instantiated via new randomly generated affine layers for each new block to encrypt (for preventing statistical attacks), and (ii) their states have huge size (for preventing linearization attacks without increasing the number of rounds, and so the depth). Their non-linear layers are instantiated via quadratic functions, including (i) the SI–lifting function $\mathcal{S}_\chi$ over $\mathbb{F}_p^n$ defined via the local quadratic chi-map $\chi : \mathbb{F}_p^3 \to \mathbb{F}_p$ introduced in [Wol85, DGV91] and (ii) the Type–III Feistel scheme [ZMI90, Nyb96] instantiated via a quadratic map. In Sect. 6, we show that it is possible to increase the security and/or the performance of such schemes by replacing such non-linear layers with the quadratic SI–lifting function $\mathcal{S}_F$ over $\mathbb{F}_p^n$ induced by $F(x_0, x_1) = x_0^2 + x_1$.

2 "Weak Bijective" Functions

In this section, we introduce the concept of weak bijective functions. Let's start with the definition of bijective functions.

Definition 2 (Bijective)**.** Let $\mathfrak{X}, \mathfrak{Y}$ be two sets. Let $\mathcal{F} : \mathfrak{X} \to \mathfrak{Y}$ be a function. The function $\mathcal{F}$ is **bijective** if (i) each element of $\mathfrak{X}$ is paired with exactly one element of $\mathfrak{Y}$, and (ii) each element of $\mathfrak{Y}$ is paired with exactly one element of $\mathfrak{X}$.

In other words, for each $y \in \mathfrak{Y}$, there exists a *unique* preimage $x \in \mathfrak{X}$ such that $\mathcal{F}(x) = y$. In the case of a bijective function, we have that (1st) $|\mathfrak{X}| = |\mathfrak{Y}|$ (where $|\mathfrak{Z}|$ denotes the cardinality of a set $\mathfrak{Z}$) and (2nd) the probability of a collision (that is, $\mathcal{F}(x) = \mathcal{F}(x')$ for $x, x' \in \mathfrak{X}$ with $x \neq x'$) is zero.

Here we relax this definition, by introducing the concept of "weak bijective" functions.

Definition 3 (Weak Bijective)**.** Let $\mathfrak{X}, \mathfrak{Y}$ be two sets. Let $\mathcal{F} : \mathfrak{X} \to \mathfrak{Y}$ be a function. The function $\mathcal{F}$ is "**weak bijective**" if (i) each element of $\mathfrak{X}$ is paired with exactly one element of $\mathfrak{Y}$, and (ii) each element of $\mathfrak{Y}$ is paired with *at most two* elements of $\mathfrak{X}$.

In other words, for each $y \in \mathfrak{Y}$, only one of the three following events occurs:

- there is *no preimage* of y, that is, $\mathcal{F}(x) \neq y$ for each $x \in \mathfrak{X}$;
- there is *exactly one preimage* $x \in \mathfrak{X}$ of y, that is, $\mathcal{F}(x) = y$ and $\mathcal{F}(\hat{x}) \neq y$ for each $\hat{x} \in \mathfrak{X} \setminus \{x\}$;
- there are *exactly two preimages* $x, x' \in \mathcal{X}$ of y with $x \neq x'$, that is, $\mathcal{F}(x) = \mathcal{F}(x') = y$ and $\mathcal{F}(\hat{x}) \neq y$ for each $\hat{x} \in \mathfrak{X} \setminus \{x, x'\}$.

As concrete examples, $x \mapsto x^2$ is weak bijective, while $(x, y) \mapsto (x^2, y^2)$ is not.

Obviously, a bijective function is also weak bijective, while the opposite is not true in general. By definition of weak bijective function, for each $x \in \mathfrak{X}$, there is *at most* one $x' \in \mathfrak{X} \setminus \{x\}$ such that $\mathcal{F}(x) = \mathcal{F}(x')$, which implies that the number of collisions is at most $\lfloor |\mathfrak{X}|/2 \rfloor$. More formally:

Lemma 1. *Let $\mathcal{F} : \mathfrak{X} \to \mathfrak{Y}$ be a weak bijective function. Then, (i) $|\mathfrak{X}| \leq 2 \cdot |\mathfrak{Y}|$, and (ii) the probability that a collision occurs is at most*

$$\left\lfloor \frac{|\mathfrak{X}|}{2} \right\rfloor \cdot \frac{2}{|\mathfrak{X}| \cdot (|\mathfrak{X}| - 1)} \leq \frac{1}{|\mathfrak{X}| - 1} \,.$$

As an immediate consequence, if the probability of the collision event is strictly bigger than $1/(|\mathfrak{X}|-1)$, then one can immediately conclude that F is not a weak bijective function. Vice-versa is not true in general.[1]

Composition of (Weak) Bijective Functions. Next, we analyze what happens when we compose two or more (weak) bijective functions defined over the same domain. As it is well known, the composition of two bijective functions is again bijective. In the case of weak bijective functions:

Lemma 2. *The composition of a bijective function and a weak bijective function (or vice-versa) is a weak bijective function. The composition of two weak bijective functions is* ***not*** *weak bijective in general.*

E.g., consider $\mathcal{G}(x) = x^2$ over $\mathbb{F}_p$. Then $\mathcal{G}(\mathcal{G}(x)) = x^4$ is not weak bijective if $p = 1$ mod 4, due to the fact that -1 is a square modulo p. Indeed, given $y = x^4$, then $x, -x, \eta \cdot x$ and $-\eta \cdot x$ are all preimages of y, where $\eta \in \mathbb{F}_p$ satisfies $\eta^2 = -1$.

This fact has an impact on the probability that a collision occurs in the case of a function defined as the composition of weak bijective functions. Not surprisingly (see e.g. [PvO95,MS15,BDD+17] for related results on non-invertible functions), the probability that a collision occurs for an *iterated weak bijective function* (that is, a function defined via the composition of weak bijective functions) is higher than the corresponding probability for a weak bijective function. For our scope, we provide an upper bound of such probability in the next proposition.

[1] E.g., consider $\mathcal{G}$ over $\mathbb{F}_p$ for a prime p defined as $\mathcal{G}(0) = \mathcal{G}(1) = \mathcal{G}(2) = 1$, and $\mathcal{G}(x) = x$ for each $x \in \mathbb{F}_p \setminus \{0, 1, 2\}$. The probability that a collision occurs is $\frac{3}{p\cdot(p-1)}$, which is smaller than $\frac{1}{p-1}$ for each $p \geq 5$. However, the function is not weak bijective since 1 has three pre-images.

Proposition 1. *Let $\mathcal{F}_0, \mathcal{F}_1, \ldots, \mathcal{F}_{r-1} : \mathfrak{X} \to \mathfrak{X}$ be $r \geq 1$ weak bijective functions over a non-empty and non-trivial set $\mathfrak{X}$. Assume $r \leq \lfloor \log_2(|\mathfrak{X}|) \rfloor$, and let $\mathcal{G}(\cdot) := \mathcal{F}_{r-1} \circ \ldots \circ \mathcal{F}_1 \circ \mathcal{F}_0(\cdot)$. Then, the probability that a collision occurs on $\mathcal{G}$ is upper bounded by*

$$Prob(\mathcal{G}(x) = \mathcal{G}(y) \cap x \neq y) \leq \frac{2^r - 1}{|\mathfrak{X}| - 1} .$$

Proof. This result follows from the fact that each output of $\mathcal{G}$ has *at most* 2^r different preimages. Indeed, for each $x \in \mathfrak{X}$ and for each $i \in \{0, 1, \ldots, r-1\}$, there exist at most two different elements $y, y' \in \mathfrak{X}$ such that $\mathcal{F}_i(y) = \mathcal{F}_i(y') = x$. Hence, if $r = 1$, then each output of $\mathcal{G}$ has at most two preimages. If $r = 2$, each output of $\mathcal{F}_1$ has at most two preimages, and each output of $\mathcal{F}_0$ has again at most two preimages, for a total of four preimages for each output of $\mathcal{G}$. Working iteratively, each output of $\mathcal{G}$ has *at most* 2^r different preimages. □

3 Non-Invertible Quadratic SI-Lifting Functions over $\mathbb{F}_p^n$ for $n \geq 3$ via $F : \mathbb{F}_p^2 \to \mathbb{F}_p$

In [GOPS22], Grassi et al. proved that, given any quadratic function $F : \mathbb{F}_p^2 \to \mathbb{F}_p$, the corresponding SI-lifting function $\mathcal{S}_F$ over $\mathbb{F}_p^n$ for $n \geq 3$ as defined in Def. 1 is never invertible. In this section, we analyze

- the probability that a collision occurs, namely, the probability that $\mathcal{S}_F(x) = \mathcal{S}_F(y)$ given $x, y \in \mathbb{F}_p^n$ such that $x \neq y$;
- the details of the inputs $x, y \in \mathbb{F}_p^n$ for which $\mathcal{S}_F(x) = \mathcal{S}_F(y)$.

As main results, we prove that:

1. given $F(x_0, x_1) = x_0^2 + x_1$ (or equivalent), the probability that a collision occurs is $\frac{(p-1)^n}{p^n \cdot (p^n - 1)} \leq p^{-n}$. In particular, we show that if a collision $\mathcal{S}_F(x) = \mathcal{S}_F(y)$ occurs, then $x_i \neq y_i$ for each $i \in \{0, 1, \ldots, n-1\}$;
2. the SI-lifting function $\mathcal{S}_F$ over $\mathbb{F}_p^n$ induced by $F(x_0, x_1) = x_0^2 + x_1$ (or equivalent) is weak bijective;
3. given any other quadratic function $F : \mathbb{F}_p^2 \to \mathbb{F}_p$ for which $\mathcal{S}_F$ can be computed via n multiplications *independently of the value of* p, the probability that a collision occurs in $\mathcal{S}_F$ is never smaller than the one corresponding for $F(x_0, x_1) = x_0^2 + x_1$.

Motivation. While the motivation regarding the analysis of the probability that a collision occurs is clear, here we explain – via a concrete example – why we are also interested on the details of the inputs $x, y \in \mathbb{F}_p^n$ for which $\mathcal{S}_F(x) = \mathcal{S}_F(y)$. Consider a sponge hash function [BDPA08] instantiated via an iterative scheme, whose round function is of the form $x \mapsto \gamma + M \times \mathcal{S}(x)$, where $\mathcal{S} : \mathbb{F}_p^n \to \mathbb{F}_p^n$ is a non-linear layer, $M \in \mathbb{F}_p^{n \times n}$ is an invertible matrix and $\gamma \in \mathbb{F}_p^n$ is a round constant. Let r, c be respectively the rate and the capacity of the sponge hash function, where $c + r = n$. Given $F : \mathbb{F}_p^m \to \mathbb{F}_p$, let's consider a non-linear layer $\mathcal{S}$ instantiated as $\mathcal{S}_F$ over $\mathbb{F}_p^n$ as defined in Def. 1. Assume that $\mathcal{S}_F$ is not invertible and assume there exist different $x, y \in \mathbb{F}_p^n$ such that (1st) $\mathcal{S}_F(x) = \mathcal{S}_F(y)$ and such that (2nd) $x_i = y_i$ for each $i \in \{n-r-1, n-r, \ldots, n-1\}$ (that is, x and y are equal in the part corresponding to the inner part of the sponge hash function). In such a case, a collision can be obviously constructed, independently of the probability that a collision occurs for $\mathcal{S}_F$. For this reason, in such a scenario it is crucial to know the details of the inputs for which the collision occurs, besides the probability of the collision event.

Depending on the details of M, a similar result can be achieved even if the linear layer M is applied before the non-linear $\mathcal{S}_F$.

Before going on, we point out that (in general) a similar problem does not arise in the case in which the primitive depends on some secret key material, due to the fact that the concrete value of the state is unknown and "masked" by the secret key.

3.1 $F(x_0, x_1) = x_1^2 + \alpha_{1,0} \cdot x_0 + \alpha_{0,1} \cdot x_1$

Let's start by analyzing the case $F(x_0, x_1) = \alpha_{0,2} \cdot x_1^2 + \alpha_{1,0} \cdot x_0 + \alpha_{0,1} \cdot x_1$ with $\alpha_{0,2}, \alpha_{1,0} \neq 0$ (the following result is equivalent for $F(x_0, x_1) = \alpha_{2,0} \cdot x_0^2 + \alpha_{1,0} \cdot x_0 + \alpha_{0,1} \cdot x_1$ with $\alpha_{2,0}, \alpha_{0,1} \neq 0$). Without loss of generality (W.l.o.g.), we assume $\alpha_{0,2} = 1$. Indeed, note that $\mathcal{S}_F$ over $\mathbb{F}_p^n$ induced by $F(x_0, x_1) = \alpha_{0,2} \cdot x_1^2 + \alpha_{1,0} \cdot x_0 + \alpha_{0,1} \cdot x_1$ is equivalent to $\alpha_{0,2} \cdot \mathcal{S}_{F'}$ induced by $F(x_0, x_1) = x_1^2 + \alpha'_{1,0} \cdot x_0 + \alpha'_{0,1} \cdot x_1$ where $\alpha'_{1,0} = \alpha_{1,0}/\alpha_{0,2}$ and $\alpha'_{0,1} = \alpha_{0,1}/\alpha_{0,2}$. (We emphasize that the same analysis/trick holds for the functions analyzed in the next subsections.)

The collision $\mathcal{S}_F(x) = \mathcal{S}_F(y)$ occurs if and only if

$$\forall i \in \{0, \ldots, n-1\}: \qquad (x_{i+1} - y_{i+1}) \cdot (x_{i+1} + y_{i+1}) = -\alpha_{1,0} \cdot (x_i - y_i) - \alpha_{0,1} \cdot (x_{i+1} - y_{i+1}) \, .$$

Via the change of variables

$$d_i := x_i - y_i \qquad \text{and} \qquad s_i := x_i + y_i, \tag{3}$$

where $x_i = (s_i + d_i)/2$ and $y_i = (s_i - d_i)/2$, the collision $\mathcal{S}_F(x) = \mathcal{S}_F(y)$ occurs if and only if

$$\begin{bmatrix} 0 & d_1 & 0 & \ldots & 0 \\ 0 & 0 & d_2 & \ldots & 0 \\ \vdots & & & \ddots & \vdots \\ 0 & 0 & 0 & \ldots & d_{n-1} \\ d_0 & 0 & 0 & \ldots & 0 \end{bmatrix} \times \begin{bmatrix} s_0 \\ s_1 \\ s_2 \\ \vdots \\ s_{n-1} \end{bmatrix} = - \begin{bmatrix} \alpha_{1,0} \cdot d_0 + \alpha_{0,1} \cdot d_1 \\ \alpha_{1,0} \cdot d_1 + \alpha_{0,1} \cdot d_2 \\ \alpha_{1,0} \cdot d_2 + \alpha_{0,1} \cdot d_3 \\ \vdots \\ \alpha_{1,0} \cdot d_{n-1} + \alpha_{0,1} \cdot d_0 \end{bmatrix} . \tag{4}$$

The determinant of the left-hand side (l.h.s.) matrix is $-(-1)^n \cdot \prod_{i=0}^{n-1} d_i$:

- if $d_i \neq 0$ for all $i \in \{0, 1, \ldots, n-1\}$, then the system admits a solution for each given $s_0, s_1, \ldots, s_{n-1}$, which corresponds to a collision;
- if $d_i = 0$ for a certain $i \in \{0, 1, \ldots, n-1\}$, e.g. $d_1 = 0$, then the condition $d_1 \cdot s_0 = -(\alpha_{1,0} \cdot d_0 + \alpha_{0,1} \cdot d_1)$ is satisfied only by $d_0 = 0$. Working iteratively, we get that if at least one d_i is zero, then the system admits a solution if and only if all d_i are zero, which corresponds to $x = y$.

It follows that

Proposition 2. *Let $p \geq 3$ be a prime and let $n \geq 2$. Let $F : \mathbb{F}_p^2 \to \mathbb{F}_p$ be defined as $F(x_0, x_1) = x_1^2 + \alpha_{1,0} \cdot x_0 + \alpha_{0,1} \cdot x_1$ where $\alpha_{1,0} \neq 0$. Let $\mathcal{S}_F$ over $\mathbb{F}_p^n$ be defined as in Def. 1. The probability of a collision $\mathcal{S}_F(x) = \mathcal{S}_F(y)$ for $x, y \in \mathbb{F}_p^n$ such that $x \neq y$ is*[2]

$$\frac{(p-1)^n}{p^n \cdot (p^n - 1)} < \frac{p^n - 1}{p^n \cdot (p^n - 1)} = p^{-n} \, .$$

Moreover, if $x, y \in \mathbb{F}_p^n$ such that $x \neq y$ and $\mathcal{S}_F(x) = \mathcal{S}_F(y)$, then $x_i \neq y_i$ ***for each*** *$i \in \{0, 1, \ldots, n-1\}$.*

[2] $p^n - 1 - (p-1)^n = \sum_{i=1}^{n-1} \binom{n}{i} p^i > 0$ for $p \geq 3$ and $n \geq 2$.

Given a difference $d \in \mathbb{F}_p^n$ such that $d_i \neq 0$ for each $i \in \{0, 1, \ldots, n-1\}$, then there exists *only one* input $x \in \mathbb{F}_p^n$ such that $\mathcal{S}_F(x+d) = \mathcal{S}_F(x)$. Moreover, given a difference $d \in \mathbb{F}_p^n$ and a sum $s \in \mathbb{F}_p^n$ that correspond to a collision, that is, $\mathcal{S}_F((s+d)/2) = \mathcal{S}_F((s-d)/2)$, we point out that $\mathcal{S}_F((s+\omega\cdot d)/2) = \mathcal{S}_F((s-\omega\cdot d)/2)$ for each $\omega \in \mathbb{F}_p$.

In particular, we have the following result:

Lemma 3. *Let $p \geq 3$ be a prime and let $n \geq 2$. Let $F : \mathbb{F}_p^2 \to \mathbb{F}_p$ be defined as $F(x_0, x_1) = x_1^2 + \alpha_{1,0}\cdot x_0 + \alpha_{0,1}\cdot x_1$ where $\alpha_{1,0} \neq 0$. Let $\mathcal{S}_F$ over $\mathbb{F}_p^n$ be defined as in Def. 1. Two different inputs $x, y \in \mathbb{F}_p^n$ satisfies $\mathcal{S}_F(x) = \mathcal{S}_F(y)$ if for each $i \in \{0, 1, \ldots, n-1\}$ they are of the form*

$$x_i = \frac{\alpha_{0,1}}{2} + \frac{d_i}{2}\cdot\left(\frac{\alpha_{1,0}}{d_{i+1}} + 1\right) \qquad \textit{and} \qquad y_i = \frac{\alpha_{0,1}}{2} + \frac{d_i}{2}\cdot\left(\frac{\alpha_{1,0}}{d_{i+1}} - 1\right) = x_i + d_i\,,$$

where $d \in \mathbb{F}_p^n$ satisfies $d_i \in \mathbb{F}_p \setminus \{0\}$ for each $i \in \{0, 1, \ldots, n-1\}$.

Proof. In order to prove the result, it is sufficient to invert the l.h.s. diagonal matrix given in (4), and to make used of the definition of d, s given in (3). □

The Function $\mathcal{S}_F$ over $\mathbb{F}_p^n$ via $F(x_0, x_1) = x_1^2 + \alpha_{1,0}\cdot x_0 + \alpha_{0,1}\cdot x_1$ is Weak Bijective

Proposition 3. *Let $p \geq 3$ be a prime and let $n \geq 2$. Let $F : \mathbb{F}_p^2 \to \mathbb{F}_p$ be defined as $F(x_0, x_1) = x_1^2 + \alpha_{1,0}\cdot x_0 + \alpha_{0,1}\cdot x_1$ where $\alpha_{1,0} \neq 0$. The SI–lifting function $\mathcal{S}_F$ over $\mathbb{F}_p^n$ defined as in Def. 1 is **weak bijective** (based on Def. 3).*

Proof. W.l.o.g., we focus on $\alpha_{1,0} = 0$ and $\alpha_{0,1} = 1$ (the following proof is equivalent for the other cases). Let's assume that the $\mathcal{S}_F$ is not weak bijective. Hence, there exist three *distinct* elements $x, \tilde{x}, \hat{x} \in \mathbb{F}_p^n$ such that $\mathcal{S}_F(x) = \mathcal{S}_F(\tilde{x}) = \mathcal{S}_F(\hat{x})$. Based on Lemma 3:

$$\begin{aligned} x_i &= \frac{d_i}{2}\cdot\left(\frac{1}{d_{i+1}} + 1\right) \qquad &\text{and} \qquad \tilde{x}_i &= \frac{d_i}{2}\cdot\left(\frac{1}{d_{i+1}} - 1\right) = x_i + d_i\,,\\ x_i &= \frac{d'_i}{2}\cdot\left(\frac{1}{d'_{i+1}} + 1\right) \qquad &\text{and} \qquad \hat{x}_i &= \frac{d'_i}{2}\cdot\left(\frac{1}{d'_{i+1}} - 1\right) = x_i + d'_i\,, \end{aligned}$$

for certain $d, d' \in \mathbb{F}_p^n$ as before. Since $\frac{d_i}{2}\cdot\left(\frac{1}{d_{i+1}} + 1\right) = x_i = \frac{d'_i}{2}\cdot\left(\frac{1}{d'_{i+1}} + 1\right)$, it follows that

$$d_i\cdot d'_{i+1}\cdot(1 - d_{i+1}) = d'_i\cdot d_{i+1}\cdot(1 - d'_{i+1}) \qquad \longrightarrow \qquad d'_i = \frac{d_i\cdot d'_{i+1}\cdot(1-d_{i+1})}{d_{i+1}\cdot(1-d'_{i+1})}\,. \tag{5}$$

Working iteratively, we get

$$\begin{aligned} d'_0 &= \frac{d_0\cdot d'_1\cdot(1-d_1)}{d_1\cdot(1-d'_1)} = d_0\cdot\left(1 + \frac{d'_1 - d_1}{d'_1\cdot(1-d_1) - d'_1 + d_1}\right)\\ &= d_0\cdot\left(1 + \frac{d'_2 - d_2}{d'_2\cdot(1-d_2)\cdot(1-d_1) - d'_2 + d_2}\right)\\ &= d_0\cdot\left(1 + \frac{d'_3 - d_3}{d'_3\cdot(1-d_3)\cdot(1-d_2)\cdot(1-d_1) - d'_3 + d_3}\right)\\ &= \ldots\\ &= d_0\cdot\left(1 + \frac{d'_{n-1} - d_{n-1}}{d'_{n-1}\cdot\prod_{i=1}^{n-1}(1-d_i) - d'_{n-1} + d_{n-1}}\right)\\ &= d_0\cdot\left(1 + \frac{d'_0 - d_0}{d'_0\cdot\prod_{i=0}^{n-1}(1-d_i) - d'_0 + d_0}\right)\,. \end{aligned}$$

By re-writing it, we have

$$(d_0')^2 \cdot (1-d_0) \cdot \prod_{i=1}^{n-1}(1-d_i) - (d_0')^2 + d_0 \cdot d_0' = d_0 \cdot d_0' \cdot (1-d_0) \cdot \prod_{i=1}^{n-1}(1-d_i)$$

$$\longrightarrow \qquad d_0' \cdot (d_0' - d_0) \cdot \left(\prod_{i=0}^{n-1}(1-d_i) - 1\right) = 0\,,$$

which implies either $d_0' = d_0$ or $(1-d_0) \cdot \prod_{i=1}^{n-1}(1-d_i) = 1$ (note that $d_0' \neq 0$).

If $d_0' = d_0$, then $d_i' = d_i$ for each $i \in \{0, 1, \ldots, n-1\}$ due to Eq. (5), that is, $\hat{x} = \tilde{x}$. Hence, $\mathcal{S}_F$ is weak bijective.

Otherwise, if $(1-d_0) \cdot \prod_{i=1}^{n-1}(1-d_i) = 1$, then $d_0' = d_0 \cdot \left(1 + \frac{d_{n-1}' - d_{n-1}}{d_{n-1}' \cdot \prod_{i=1}^{n-1}(1-d_i) - d_{n-1}' + d_{n-1}}\right)$ implies

$$d_0' = d_0 \cdot \left(\frac{d_{n-1}'}{d_{n-1}' + (1-d_0) \cdot (d_{n-1} - d_{n-1}')}\right) \qquad \longrightarrow \qquad d_0' \cdot (1-d_0) \cdot (d_{n-1} - d_{n-1}') = 0\,,$$

that is, $d_{n-1} = d_{n-1}'$ (note that $d_0' \neq 0$ and that $d_0 \neq 1$, since $(1-d_0) \cdot \prod_{i=1}^{n-1}(1-d_i) = 1$). Working as before, it follows that $d_i' = d_i$ for each $i \in \{0, 1, \ldots, n-1\}$ due to Eq. (5), that is, $\hat{x} = \tilde{x}$. Hence, $\mathcal{S}_F$ is weak bijective. □

3.2 $F(x_0, x_1) = x_0 \cdot x_1 + \alpha_{1,0} \cdot x_0 + \alpha_{0,1} \cdot x_1$

Next, we analyze $\mathcal{S}_F$ induced by $F(x_0, x_1) = \alpha_{1,1} \cdot x_0 \cdot x_1 + \alpha_{1,0} \cdot x_0 + \alpha_{0,1} \cdot x_1$. W.l.o.g., we fix $\alpha_{1,1} = 2$ (this allows for a simpler description when using the variables s_i and d_i). Given $F(x_0, x_1) = 2 \cdot x_0 \cdot x_1 + \alpha_{1,0} \cdot x_0 + \alpha_{0,1} \cdot x_1$, the system of equations that defines the collision $\mathcal{S}_F(x) = \mathcal{S}_F(y)$ via the variables $d_i := x_i - y_i$ and $s_i := x_i + y_i$ introduced in Eq. (3) is

$$\begin{bmatrix} d_1 & d_0 & 0 & 0 & \ldots & 0 \\ 0 & d_2 & d_1 & 0 & \ldots & 0 \\ 0 & 0 & d_3 & d_2 & \ldots & 0 \\ \vdots & & & \ddots & \ddots & \vdots \\ 0 & 0 & 0 & \ldots & d_{n-1} & d_{n-2} \\ d_{n-1} & 0 & 0 & \ldots & 0 & d_0 \end{bmatrix} \times \begin{bmatrix} s_0 \\ s_1 \\ s_2 \\ \vdots \\ s_{n-2} \\ s_{n-1} \end{bmatrix} = - \begin{bmatrix} \alpha_{1,0} \cdot d_0 + \alpha_{0,1} \cdot d_1 \\ \alpha_{1,0} \cdot d_1 + \alpha_{0,1} \cdot d_2 \\ \alpha_{1,0} \cdot d_2 + \alpha_{0,1} \cdot d_3 \\ \vdots \\ \alpha_{1,0} \cdot d_{n-2} + \alpha_{0,1} \cdot d_{n-1} \\ \alpha_{1,0} \cdot d_{n-1} + \alpha_{0,1} \cdot d_0 \end{bmatrix}. \tag{6}$$

The determinant of the l.h.s. matrix is

$$(1 - (-1)^n) \cdot \prod_{i=0}^{d-1} d_i\,,$$

that is, $2 \cdot \prod_{i=0}^{d-1} d_i$ for odd n, and zero for even n (independently of the values of d_i). As we are going to show:

1. the probability that a collision occurs is strictly higher than $\frac{(p-1)^n}{p^n \cdot (p^n - 1)}$, which corresponds to the probability of having a collision for $F(x_0, x_1) = x_1^2 + x_0$ (and equivalent functions) as given in Prop. 2;

2. a collision can occur also in the case in which $n-1$ input differences d_i are equal to zero.

Analysis of $x, y \in \mathbb{F}_p^n$ such that $\mathcal{S}_F(x) = \mathcal{S}_F(y)$. About this second point, consider the case $d_j \in \mathbb{F}_p \setminus \{0\}$ and $d_i = 0$ for each $i \neq j$. In this case, the system of equation reduces to

$$d_j \cdot s_{j-1} = -\alpha_{0,1} \cdot d_j \qquad \text{and} \qquad d_j \cdot s_{j+1} = -\alpha_{1,0} \cdot d_j,$$

which is satisfied by $s_{j-1} = -\alpha_{0,1}$ and $s_{j+2} = -\alpha_{1,0}$ (no condition on all others s_l for $j \notin \{0, 2\}$).

Collision Probability for n odd. First of all, note that if $d_i \neq 0$ for each $i \in \{0, 1, \ldots, n-1\}$, then a collision can occur. Indeed, the determinant is different from zero, which means that there exist $s_0, s_1, \ldots, s_{n-1}$ that satisfy the required condition for having a collision.

Let's consider the case in which $n-1$ differences d_i are equal to zero (note that there are n different cases). This case is obviously not included in the previous one, since now the determinant is equal to zero. As pointed out in the previous paragraph, a collision can occur if s_{i-1} and s_{i+1} satisfy some particular equalities, while no condition is imposed on the other s_j. As a result, the probability of having a collision is *at least* equal

$$\frac{(p-1)^n + n \cdot p^{n-2} \cdot (p-1)}{p^n \cdot (p^n - 1)} > \frac{(p-1)^n}{p^n \cdot (p^n - 1)},$$

which is strictly bigger than the probability given in Prop. 2.

Collision Probability for n even. Since the determinant of the matrix is always equal to zero, there is a linear combination among its rows. Assuming such linear combination is defined via $\lambda_0, \lambda_1, \ldots, \lambda_{n-1} \in \mathbb{F}_p$, a collision can occur if the right-hand side (r.h.s.) of (6) satisfies the same linear relation, that is, if $\sum_{i=0}^{n-1} \lambda_i \cdot (\alpha_{1,0} \cdot d_i + \alpha_{0,1} \cdot d_{i+1}) = 0$. In such a case, this implies that one difference d_i is fixed. W.l.o.g., assuming that d_1 satisfies such linear relation, the collision takes place if

$$\begin{bmatrix} d_2 & d_1 & 0 & \ldots & 0 \\ 0 & d_3 & d_2 & \ldots & 0 \\ \vdots & & \ddots & \ddots & \vdots \\ 0 & 0 & \ldots & d_{n-1} & d_{n-2} \\ 0 & 0 & \ldots & 0 & d_0 \end{bmatrix} \times \begin{bmatrix} s_1 \\ s_2 \\ \vdots \\ s_{n-2} \\ s_{n-1} \end{bmatrix} = - \begin{bmatrix} \alpha_{1,0} \cdot d_1 + \alpha_{0,1} \cdot d_2 \\ \alpha_{1,0} \cdot d_2 + \alpha_{0,1} \cdot d_3 \\ \vdots \\ \alpha_{1,0} \cdot d_{n-2} + \alpha_{0,1} \cdot d_{n-1} \\ (\alpha_{1,0} + s_0) \cdot d_{n-1} + \alpha_{0,1} \cdot d_0 \end{bmatrix},$$

where d_1 is fixed and where no condition holds on s_0. The determinant of the l.h.s. matrix is equal to $d_0 \cdot \prod_{i=2}^{n-1} d_i$. As before, a collision can occur if $d_0, d_2, d_3, \ldots, d_{n-1} \neq 0$, since in such a case the determinant of the matrix is different from zero. This is sufficient for concluding that the probability of having a collision is *at least* equal to

$$\frac{p \cdot (p-1)^{n-1}}{p^n \cdot (p^n - 1)} > \frac{(p-1)^n}{p^n \cdot (p^n - 1)},$$

which is strictly bigger than the probability given in Prop. 2.

3.3 $F(x_0, x_1) = x_0^2 + \alpha_{0,2} \cdot x_1^2 + \alpha_{1,0} \cdot x_0 + \alpha_{0,1} \cdot x_1$

Given $F(x_0, x_1) = x_0^2 + \alpha_{0,2} \cdot x_1^2 + \alpha_{1,0} \cdot x_0 + \alpha_{0,1} \cdot x_1$ where $\alpha_{0,2} \neq 0$ (w.l.o.g., we fixed $\alpha_{2,0} = 1$), the system of equations that defines the collision $\mathcal{S}_F(x) = \mathcal{S}_F(y)$ via the variables

$d_i := x_i - y_i$ and $s_i := x_i + y_i$ introduced in Eq. (3) is

$$\begin{bmatrix} d_0 & \alpha_{0,2}\cdot d_1 & 0 & 0 & \dots & 0 \\ 0 & d_1 & \alpha_{0,2}\cdot d_2 & 0 & \dots & 0 \\ 0 & 0 & d_2 & \alpha_{0,2}\cdot d_3 & \dots & 0 \\ \vdots & & & \ddots & \ddots & \vdots \\ 0 & 0 & 0 & \dots & d_{n-2} & \alpha_{0,2}\cdot d_{n-1} \\ \alpha_{0,2}\cdot d_0 & 0 & 0 & \dots & 0 & d_{n-1} \end{bmatrix} \times \begin{bmatrix} s_0 \\ s_1 \\ s_2 \\ \vdots \\ s_{n-2} \\ s_{n-1} \end{bmatrix} = \qquad (7)$$
$$- \begin{bmatrix} \alpha_{1,0}\cdot d_0 + \alpha_{0,1}\cdot d_1 & \alpha_{1,0}\cdot d_1 + \alpha_{0,1}\cdot d_2 & \dots & \alpha_{1,0}\cdot d_{n-1} + \alpha_{0,1}\cdot d_0 \end{bmatrix}^T ,$$

where $\cdot^T$ denotes the transpose matrix/vector. The determinant of the l.h.s. matrix is equal to

$$(1 - (-\alpha_{0,2})^n) \cdot \prod_{i=0}^{n-1} d_i \,.$$

Hence, in order to give a lower bound on the probability of having a collision, we study separately the two cases: (1st) $1 \neq (-\alpha_{0,2})^n$ and (2nd) $1 = (-\alpha_{0,2})^n$. Before going on, we point out that $\mathcal{S}_F$ costs n multiplications by pre-computing $x_0^2, x_1^2, \dots, x_{n-1}^2$.

Analysis of $x, y \in \mathbb{F}_p^n$ such that $\mathcal{S}_F(x) = \mathcal{S}_F(y)$. As first step, we analyze the details of x, y such that $\mathcal{S}_F(x) = \mathcal{S}_F(y)$. In this case, a collision does not necessary occur if $n-1$ input differences d_i are equal to zero. E.g., in the case $d_1 \in \mathbb{F}_p \setminus \{0\}$ and $d_i = 0$ for each $i \neq 1$, the system of equations reduces to

$$\alpha_{0,2}\cdot d_1 \cdot s_1 = -\alpha_{0,1}\cdot d_1 \qquad \text{and} \qquad d_1 \cdot s_1 = -\alpha_{1,0}\cdot d_1 \,,$$

which is satisfied by $s_1 = -\alpha_{1,0}$ and $\alpha_{0,2}\cdot\alpha_{1,0} = \alpha_{0,1}$. If this second condition is not satisfied, then a collision cannot occur, independently of the choice of s_i. At the same time, if at least two differences are non-null (e.g., $d_1, d_2 \in \mathbb{F}_p \setminus \{0\}$), then it is always possible to have a collision (even if $\alpha_{0,1} \neq \alpha_{0,2}\cdot\alpha_{1,0}$). Indeed, in such a case, the system of equations reduces to

$$\begin{aligned} \alpha_{0,2}\cdot d_1 \cdot s_1 &= -\alpha_{0,1}\cdot d_1 \,, \\ d_1 \cdot s_1 + \alpha_{0,2}\cdot d_2 \cdot s_2 &= -\alpha_{1,0}\cdot d_1 - \alpha_{0,1}\cdot d_2 \,, \\ d_2 \cdot s_2 &= -\alpha_{1,0}\cdot d_2 \,, \end{aligned}$$

which is satisfied by $s_1 = -\alpha_{0,1}/\alpha_{0,2}$, $s_2 = -\alpha_{1,0}$ and by $d_1 = d_2 \cdot \alpha_{0,2}$.

Collision Probability for $1 - (-\alpha_{0,2})^n \neq 0$. First of all, if all d_i are non-zero, then the determinant of the l.h.s. matrix is non-zero, and a collision can occur by properly choosing $s_0, s_1, \dots, s_{n-1}$. Consider the case in which only two differences d_i, d_{i+1} are non-null, and all the others are equal to zero (note that there are n different cases). As pointed out in the previous paragraph, given d_i, a collision can occur if s_i, s_{i+1}, d_{i+1} satisfy some particular relation (note that all the others s_j for $j \in \{0, 1, \dots, n-1\} \setminus \{i, i+1\}$ are free to take any possible value). As a result, the probability of having a collision is *at least* equal to

$$\frac{(p-1)^n + n \cdot (p-1)^{n-3}}{p^n \cdot (p^n - 1)} > \frac{(p-1)^n}{p^n \cdot (p^n-1)} \,,$$

which is strictly bigger than the probability given in Prop. 2.

Collision Probability for $1-(-\alpha_{0,2})^n = 0$. If $1-(-\alpha_{0,2})^n = 0$, then the determinant of the l.h.s. matrix is equal to zero, which means that its rows satisfy a linear relation. Working as in Sect. 3.2, a collision can occur if the r.h.s. of (7) satisfies the same linear relation of the rows of the l.h.s. matrix, which implies that one difference d_i is fixed. W.l.o.g., let's assume d_0 is fixed. In such a case, the collision takes place if

$$\begin{bmatrix} d_1 & \alpha_{0,2}\cdot d_2 & 0 & \dots & 0\\ 0 & d_2 & \alpha_{0,2}\cdot d_3 & \dots & 0\\ \vdots & & \ddots & \ddots & \vdots\\ 0 & 0 & \dots & d_{n-2} & \alpha_{0,2}\cdot d_{n-1}\\ 0 & 0 & \dots & 0 & d_{n-1}\end{bmatrix} \times \begin{bmatrix} s_1\\ s_2\\ \vdots\\ s_{n-2}\\ s_{n-1}\end{bmatrix} = - \begin{bmatrix} \alpha_{1,0}\cdot d_1 + \alpha_{0,1}\cdot d_2\\ \alpha_{1,0}\cdot d_2 + \alpha_{0,1}\cdot d_3\\ \vdots\\ \alpha_{1,0}\cdot d_{n-2} + \alpha_{0,1}\cdot d_{n-1}\\ \alpha_{1,0}\cdot d_{n-1} + (\alpha_{0,1}+\alpha_{0,2}\cdot s_0)\cdot d_0\end{bmatrix},$$

where no condition on $s_0 \in \mathbb{F}_p$ holds. The determinant of the l.h.s. matrix is equal to $\prod_{i=1}^{n-1} d_i$, which is different from zero if $d_1, d_2, \dots, d_{n-1} \in \mathbb{F}_p \setminus \{0\}$. This is sufficient for concluding that the probability of having a collision is *at least* equal to

$$\frac{p\cdot(p-1)^{n-1}}{p^n\cdot(p^n-1)} > \frac{(p-1)^n}{p^n\cdot(p^n-1)},$$

which is strictly bigger than the probability given in Prop. 2.

3.4 $F(x_0,x_1) = x_0\cdot x_1 + \alpha_{2,0}\cdot x_0^2 + \alpha_{1,0}\cdot x_0 + \alpha_{0,1}\cdot x_1$

Consider $F(x_0,x_1) = \alpha_{1,1}\cdot x_0\cdot x_1 + \alpha_{2,0}\cdot x_0^2 + \alpha_{1,0}\cdot x_0 + \alpha_{0,1}\cdot x_1$ where $\alpha_{1,1}, \alpha_{2,0} \neq 0$ (the following result is equivalent for $F(x_0,x_1) = \alpha_{1,1}\cdot x_0\cdot x_1 + \alpha_{0,2}\cdot x_1^2 + \alpha_{1,0}\cdot x_0 + \alpha_{0,1}\cdot x_1$ where $\alpha_{1,1}, \alpha_{0,2} \neq 0$). W.l.o.g., we fix $\alpha_{1,1} = 2$. Given $F(x_0,x_1) = 2\cdot x_0\cdot x_1 + \alpha_{2,0}\cdot x_0^2 + \alpha_{1,0}\cdot x_0 + \alpha_{0,1}\cdot x_1$ where $\alpha_{2,0} \neq 0$, the system of equations that corresponds to the collision $\mathcal{S}_F(x) = \mathcal{S}_F(y)$ via the variables $d_i := x_i - y_i$ and $s_i := x_i + y_i$ introduced in (3) is

$$\begin{bmatrix} \alpha_{2,0}\cdot d_0 + d_1 & d_0 & 0 & 0 & \dots & 0\\ 0 & \alpha_{2,0}\cdot d_1 + d_2 & d_1 & 0 & \dots & 0\\ 0 & 0 & \alpha_{2,0}\cdot d_2 + d_3 & d_2 & \dots & 0\\ \vdots & & & \ddots & \ddots & \vdots\\ 0 & 0 & 0 & \dots & \alpha_{2,0}\cdot d_{n-2} + d_{n-1} & d_{n-2}\\ d_{n-1} & 0 & 0 & \dots & 0 & \alpha_{2,0}\cdot d_{n-1} + d_0\end{bmatrix}$$
$$\times \begin{bmatrix} s_0 & s_1 & s_2 & \dots & s_{n-2} & s_{n-1}\end{bmatrix}^T =$$
$$- \begin{bmatrix} \alpha_{1,0}\cdot d_0 + \alpha_{0,1}\cdot d_1 & \alpha_{1,0}\cdot d_1 + \alpha_{0,1}\cdot d_2 & \dots & \alpha_{1,0}\cdot d_{n-1} + \alpha_{0,1}\cdot d_0\end{bmatrix}^T .$$

Before going on, note that the function F can be computed via one multiplication only, by re-writing it as $F(x_0,x_1) = x_0\cdot(\alpha_{1,1}\cdot x_1 + \alpha_{2,0}\cdot x_0) + \alpha_{1,0}\cdot x_0 + \alpha_{0,1}\cdot x_1$.

Analysis of $x, y \in \mathbb{F}_p^n$ such that $\mathcal{S}_F(x) = \mathcal{S}_F(y)$. In this case, a collision can occur even if $n-1$ differences d_i are equal to zero. E.g., if $d_1 \neq 0$ and $d_i = 0$ for each $i \neq 1$, the system of equations reduces to

$$d_1\cdot s_0 = -\alpha_{0,1}\cdot d_1 \qquad \text{and} \qquad \alpha_{2,0}\cdot d_1\cdot s_1 + d_1\cdot s_2 = -\alpha_{1,0}\cdot d_1,$$

which is satisfied if $s_0 = -\alpha_{0,1}$ and $s_2 = -\alpha_{1,0} - \alpha_{2,0}\cdot s_1$, where $s_1, s_3, s_4, \dots, s_{n-1}$ are free to take any possible value.

Collision Probability. The determinant of the l.h.s. matrix is

$$\prod_{i=0}^{n-1}(\alpha_{2,0}\cdot d_i + d_{i+1}) - (-1)^n\cdot\prod_{i=0}^{n-1} d_i .$$

By re-writing it with respect to d_0, the determinant is equal to zero if and only if

$$\alpha_{2,0} \cdot \beta \cdot d_0^2 + \left(\alpha_{2,0}^2 \cdot \beta \cdot d_{n-1} + \beta \cdot d_1 - (-1)^n \cdot \prod_{i=1}^{n-1} d_i \right) \cdot d_0 + \alpha_{2,0} \cdot d_1 \cdot \beta \cdot d_{n-1} = 0\,,$$

where $\beta := \prod_{i=1}^{n-2} (\alpha_{2,0} \cdot d_i + d_{i+1})$.

The case $\beta \neq 0$ holds if $d_{i-1} \neq -\alpha_{2,0} \cdot d_i$, i.e., $d_{n-1} \in \mathbb{F}_p$ and $d_i \in \mathbb{F}_p \setminus \{-\alpha_{2,0} \cdot d_{i+1}\}$ for each $i \in \{1, 2, \ldots, n-2\}$. If $\beta \neq 0$, then the previous equation of degree two admits at most two solutions in d_0. This means that there are at least $p \cdot (p-1)^{n-2} \cdot (p-2)$ different values of $d_0, d_1, \ldots, d_{n-1} \in \mathbb{F}_p$ for which the matrix is invertible, and so for which a collision occurs.

As pointed out in the previous paragraph, a collision can also occur if $n-1$ differences are equal to zero. E.g., If $d_i \neq 0$, this happens if s_{i-1} and s_{i+1} satisfy some particular relations given before. Note that this case is excluded from the previous case, since the determinant is equal to zero. This is sufficient for concluding that the probability of having a collision is *at least* equal to

$$\frac{p \cdot (p-1)^{n-2} \cdot (p-2) + n \cdot (p-1) \cdot p^{n-2}}{p^n \cdot (p^n - 1)} = \frac{p \cdot (p-1) \cdot ((p-1)^{n-3} \cdot (p-2) + n \cdot p^{n-3})}{p^n \cdot (p^n - 1)}\,.$$

Since $(p-1)^{n-3} \cdot (p-2) + n \cdot p^{n-3} \geq (p-1)^{n-2}$ if and only if $n \cdot p^{n-3} \geq (p-1)^{n-3}$ (which is always satisfied), then we conclude that such probability is strictly bigger than the probability given in Prop. 2.

3.5 $F(x_0, x_1) = \alpha_{2,0} \cdot x_0^2 + x_0 \cdot x_1 + \alpha_{0,2} \cdot x_1^2 + \alpha_{1,0} \cdot x_0 + \alpha_{0,1} \cdot x_1$

Given $F(x_0, x_1) = \alpha_{2,0} \cdot x_0^2 + 2 \cdot x_0 \cdot x_1 + \alpha_{0,2} \cdot x_1^2 + \alpha_{1,0} \cdot x_0 + \alpha_{0,1} \cdot x_1$ for $\alpha_{2,0}, \alpha_{0,2} \in \mathbb{F}_p \setminus \{0\}$ (w.l.o.g., we fixed $\alpha_{1,1} = 2$), the system of equations that corresponds to the collision $\mathcal{S}_F(x) = \mathcal{S}_F(y)$ via the variables $d_i := x_i - y_i$ and $s_i := x_i + y_i$ introduced in Def. 3 is

$$\begin{aligned} &\begin{bmatrix} \alpha_{2,0} \cdot d_0 + d_1 & d_0 + \alpha_{0,2} \cdot d_1 & 0 & \ldots & 0 & 0 \\ 0 & \alpha_{2,0} \cdot d_1 + d_2 & d_2 + \alpha_{0,2} \cdot d_3 & \ldots & 0 & 0 \\ \vdots & & & \ddots & \ddots & \vdots \\ 0 & 0 & 0 & \ldots & \alpha_{2,0} \cdot d_{n-2} + d_{n-1} & d_{n-2} + \alpha_{0,2} \cdot d_{n-1} \\ d_{n-1} + \alpha_{0,2} \cdot d_0 & 0 & 0 & \ldots & 0 & \alpha_{2,0} \cdot d_{n-1} + d_0 \end{bmatrix} \\ &\qquad \times \begin{bmatrix} s_0 & s_1 & s_2 & \ldots & s_{n-2} & s_{n-1} \end{bmatrix}^T = \\ &- \begin{bmatrix} \alpha_{1,0} \cdot d_0 + \alpha_{0,1} \cdot d_1 & \alpha_{1,0} \cdot d_1 + \alpha_{0,1} \cdot d_2 & \ldots & \alpha_{1,0} \cdot d_{n-1} + \alpha_{0,1} \cdot d_0 \end{bmatrix}^T . \end{aligned} \tag{8}$$

Multiplicative Complexity for Computing $\mathcal{S}_F$. First of all, we discuss the cost of computing $\mathcal{S}_F$, keeping in mind that our goal is to consider only quadratic non-linear layers over $\mathbb{F}_p^n$ that cost n multiplications. In general, computing $\mathcal{S}_F$, costs $2 \cdot n$ multiplications, since one has to compute both $x_0^2, x_1^2, \ldots, x_{n-1}^2$ and $x_0 \cdot x_1, x_1 \cdot x_2, \ldots, x_{n-1} \cdot x_0$. However, if F can be re-written as

$$\begin{aligned} F(x_0, x_1) &= (\varphi_0 \cdot x_0 + \varphi_1 \cdot x_1) \cdot (\psi_0 \cdot x_0 + \psi_1 \cdot x_1) + \alpha_{1,0} \cdot x_0 + \alpha_{0,1} \cdot x_1 = \\ &= \varphi_0 \cdot \psi_0 \cdot x_0^2 + (\varphi_0 \cdot \psi_1 + \varphi_1 \cdot \psi_0) \cdot x_0 \cdot x_1 + \varphi_1 \cdot \psi_1 \cdot x_1^2 + \alpha_{1,0} \cdot x_0 + \alpha_{0,1} \cdot x_1 \end{aligned}$$

for certain $\varphi_0, \varphi_1, \psi_0, \psi_1 \in \mathbb{F}_p \setminus \{0\}$, then the cost decreases to exactly n multiplications for $\mathcal{S}_F$. This case occurs if the following equalities are all satisfied:

$$\alpha_{2,0} = \varphi_0 \cdot \psi_0, \qquad \alpha_{1,1} = \varphi_0 \cdot \psi_1 + \varphi_1 \cdot \psi_0, \qquad \alpha_{0,2} = \varphi_1 \cdot \psi_1\,.$$

Given $\alpha_{2,0}, \alpha_{0,2} \in \mathbb{F}_p \setminus \{0\}$ and $\alpha_{1,1} = 2$, these three equalities are satisfied if

$$\alpha_{0,2} \cdot \psi_1^2 - 2\psi_0 \cdot \psi_1 + \alpha_{2,0} \cdot \psi_0 = 0 \qquad \longrightarrow \qquad \psi_1 = \frac{\psi_0 \cdot (1 \pm \sqrt{1 - \alpha_{0,2} \cdot \alpha_{2,0}})}{\alpha_{0,2}} \,.$$

The only case in which the square root exists *independently of the value of* p is $\alpha_{0,2} \cdot \alpha_{2,0} = 1$. For this reason, we limit ourselves to work with $\alpha_{0,2} \cdot \alpha_{2,0} = 1$ in the following.

Analysis of $x, y \in \mathbb{F}_p^n$ such that $\mathcal{S}_F(x) = \mathcal{S}_F(y)$. First of all, we notice that a collision can occur even if $n-1$ differences d_i are equal to zero. E.g., if $d_1 \neq 0$ and $d_i = 0$ for each $i \neq 1$, we have

$$d_1 \cdot s_0 + \alpha_{0,2} \cdot d_1 \cdot s_1 = -\alpha_{0,1} \cdot d_1 \qquad \text{and} \qquad \alpha_{2,0} \cdot d_1 \cdot s_1 + d_1 \cdot s_2 = -\alpha_{1,0} \cdot d_1$$

which is satisfied if $s_0 = -\alpha_{0,2} \cdot s_1 - \alpha_{0,1}$ and $s_2 = -\alpha_{2,0} \cdot s_1 - \alpha_{1,0}$, where $s_1, s_3, s_4, \dots, s_{n-1}$ can take any possible value in $\mathbb{F}_p$.

Collision Probability. As before, our goal is to show that the probability of having a collision is strictly bigger than $\frac{(p-1)^n}{p^n \cdot (p^n - 1)}$. By simple computation, the determinant of the matrix is equal to

$$\prod_{i=0}^{n-1} (\alpha_{2,0} \cdot d_i + d_{i+1}) - (-1)^n \cdot \prod_{i=0}^{n-1} (d_i + \alpha_{0,2} \cdot d_{i+1}) \,.$$

Following the strategy proposed in Sect. 3.4 and by re-writing the determinant with respect to d_0, it is equal to zero if and only if

$$\begin{aligned} &d_0^2 \cdot (\alpha_{2,0} \cdot \beta + \alpha_{0,2} \cdot \gamma) + d_0 \cdot (\alpha_{2,0}^2 \cdot \beta \cdot d_{n-1} + \beta \cdot d_1 + \alpha_{0,2}^2 \cdot d_1 \cdot \gamma + d_{n-1} \cdot \gamma) \\ &+ (\alpha_{2,0} \cdot \beta + \alpha_{0,2} \cdot \gamma) \cdot d_1 \cdot d_{n-1} = 0 \,, \end{aligned} \tag{9}$$

where

$$\beta := \prod_{i=1}^{n-2} (\alpha_{2,0} \cdot d_i + d_{i+1}) \qquad \text{and} \qquad \gamma := -(-1)^n \cdot \prod_{i=1}^{n-2} (d_i + \alpha_{0,2} \cdot d_{i+1}) \,.$$

By limiting ourselves to focus on $\alpha_{0,2} \cdot \alpha_{2,0} = 1$, note that

$$\beta = \prod_{i=1}^{n-2} \left(\frac{1}{\alpha_{0,2}} \cdot d_i + d_{i+1} \right) = \left(\frac{1}{\alpha_{0,2}} \right)^{n-2} \cdot \prod_{i=1}^{n-2} (d_i + \alpha_{0,2} \cdot d_{i+1}) = - \left(- \frac{1}{\alpha_{0,2}} \right)^{n-2} \cdot \gamma \,. \tag{10}$$

This implies that $\gamma = 0$ if and only if $\beta = 0$. Moreover, the coefficient $\alpha_{2,0} \cdot \beta + \alpha_{0,2} \cdot \gamma$ of d_0^2 in (9) can be re-written as

$$\alpha_{2,0} \cdot \beta + \alpha_{0,2} \cdot \gamma = \gamma \cdot \alpha_{0,2} \cdot \left(1 - \left(- \frac{1}{\alpha_{0,2}} \right)^n \right) \,,$$

which implies that

- if $(-\alpha_{0,2})^n \neq 1$, then the coefficient $\alpha_{2,0} \cdot \beta + \alpha_{0,2} \cdot \gamma$ of d_0^2 is equal to zero if and only if $\beta = \gamma = 0$;
- if $(-\alpha_{0,2})^n = 1$, then the coefficient of d_0^2 in Eq. (9) is always equal to zero.

Case: $(-\alpha_{0,2})^n \neq 1$. As we have just seen, the coefficient $\alpha_{2,0} \cdot \beta + \alpha_{0,2} \cdot \gamma$ of d_0^2 is equal to zero if and only if $\beta = \gamma = 0$. Working as in Sect. 3.4, by choosing $d_1, \ldots, d_{n-1} \in \mathbb{F}_p$ such that $d_i \neq -\alpha_{2,0} \cdot d_{i-1}$ for each $i \in \{2, 3, \ldots, n-1\}$ (where e.g. d_1 is free to take any possible value), then $\beta, \gamma \neq 0$. In such a case, there are at most two values of d_0 that satisfies Eq. (9). In other words, there are $p \cdot (p-1)^{n-2} \cdot (p-2)$ different values of $d_0, d_1, \ldots, d_{n-1}$ for which the determinant is different from zero.

As pointed out in the previous paragraph, a collision can occur even if $n-1$ differences d_i are equal to zero (note that this case is obviously excluded from the previous one, since the determinant would be zero). If d_i is not null, s_{i-1}, s_{i+1} would be fixed, while all other s_j are free to take any possible. This is sufficient for concluding that the probability of having a collision is *at least* equal to

$$\frac{p \cdot (p-1)^{n-2} \cdot (p-2) + n \cdot (p-1) \cdot p^{n-2}}{p^n \cdot (p^n - 1)} = \frac{p \cdot (p-1) \cdot ((p-1)^{n-3} \cdot (p-2) + n \cdot p^{n-3})}{p^n \cdot (p^n - 1)},$$

which is strictly bigger than $\frac{(p-1)^n}{p^n \cdot (p^n-1)}$, that is, the probability given in Prop. 2.

Case: $(-\alpha_{0,2})^n = 1$. As we already pointed out, $\alpha_{2,0} \cdot \beta + \alpha_{0,2} \cdot \gamma$ is equal to zero in this case, which implies that Eq. (9) reduces to

$$d_0 \cdot (\alpha_{2,0}^2 \cdot \beta \cdot d_{n-1} + \beta \cdot d_1 + \alpha_{0,2}^2 \cdot d_1 \cdot \gamma + d_{n-1} \cdot \gamma) = 0\,.$$

Since $\alpha_{0,2} \cdot \alpha_{2,0} = 1$ and since $\gamma = -\alpha_{2,0}^2 \cdot \beta$ due to Eq. (10), such equation is always satisfied for each d_0, d_1, β. It follows that the determinant is always equal to zero, and so that the rows of the r.h.s. vector in Eq. (8) must satisfy the same linear relation of the rows of the l.h.s. matrix. This implies that one difference d_i is fixed. W.l.o.g., we assume d_0 satisfies such linear relation. In such a case, a collision takes place if

$$\begin{bmatrix} \alpha_{2,0} \cdot d_1 + d_2 & d_1 + \alpha_{0,2} \cdot d_2 & 0 & \ldots & 0 & 0 \\ 0 & \alpha_{2,0} \cdot d_2 + d_3 & d_2 + \alpha_{0,2} \cdot d_3 & \ldots & 0 & 0 \\ \vdots & & \ddots & \ddots & & \vdots \\ 0 & 0 & 0 & \ldots & \alpha_{2,0} \cdot d_{n-2} + d_{n-1} & d_{n-2} + \alpha_{0,2} \cdot d_{n-1} \\ 0 & 0 & 0 & \ldots & 0 & \alpha_{2,0} \cdot d_{n-1} + d_0 \end{bmatrix}$$
$$\times \begin{bmatrix} s_1 & s_2 & \ldots & s_{n-2} & s_{n-1} \end{bmatrix}^T =$$
$$- \begin{bmatrix} \alpha_{1,0} \cdot d_1 + \alpha_{0,1} \cdot d_2 & \ldots & \alpha_{1,0} \cdot d_{n-2} + \alpha_{0,1} \cdot d_{n-1} & \alpha_{1,0} \cdot d_{n-1} + \alpha_{0,1} \cdot d_0 + s_0 \cdot (d_{n-1} + \alpha_{0,2} \cdot d_0) \end{bmatrix}^T,$$

where s_0 can take any possible value in $\mathbb{F}_p$. The determinant of the l.h.s. matrix is equal to $\prod_{i=1}^{n-1} (\alpha_{2,0} \cdot d_i + d_{i+1})$. Since d_0 is fixed, there are $(p-1)^{n-1}$ values of $d_1, d_2, \ldots, d_{n-1}$ for which such matrix is invertible, and so, for which a collision can occur. This is sufficient for concluding that the probability of having a collision is *at least* equal to

$$\frac{p \cdot (p-1)^{n-1}}{p^n \cdot (p^n - 1)} > \frac{(p-1)^n}{p^n \cdot (p^n - 1)},$$

which is strictly bigger than the probability given in Prop. 2.

4 The PRF `MiMC++`: Reducing the Multiplicative Complexity of MiMC via the Square Map

MiMC [AGR+16] is an iterated Even-Mansour cipher proposed at Asiacrypt 2016 for MPC [GRR+16] applications. Here, we show that it is possible to reduce its multiplicative complexity by instantiating it with the square map. We call this modified version of MiMC as `MiMC++`.

4.1 The PRF `MiMC++`

MiMC. MiMC [AGR+16] over $\mathbb{F}_p$ is an iterated Even-Mansour cipher, whose round function is defined as $F(x) = (x + \mathtt{K} + \gamma)^d$, where $\mathtt{K} \in \mathbb{F}_p$ is the secret-key, $\gamma \in \mathbb{F}_p$ is a random round constant and $d \geq 3$ is the smallest integer that satisfies $\gcd(d, p-1) = 1$ (in order to guarantee invertibility). A final key is added. For a security level of $\kappa \approx \log_2(p)$ bits with a data-limit of $2^{\kappa/2} \approx p^{1/2}$ texts available for the attack, the number of rounds is $R = 1 + \lceil(\kappa - 2 \cdot \log_2(\kappa)) \cdot \log_d(2)\rceil$.[3]

`MiMC++`. As we already mentioned in the introduction, MiMC's designers suggest to use it in a mode of operation in which the inverse is not needed. Hence, a natural question arises: if the inverse is not needed, why not implementing it with a non-invertible function? In such a case, the simplest non-linear weak bijective map over $\mathbb{F}_p$ is the quadratic one. Given $F(x) = x^2$, a collision $x^2 = y^2$ can occur if and only if $x = \pm y$, which implies that the probability of having a collision is $(p-1)^{-1} \approx p^{-1}$.

Based on this simple observation, we propose the PRF `MiMC++` defined as following. Let κ be the security level, and let p be *any* prime number such that $p \geq 2^{3\cdot\kappa}$. Given a secret key $\mathtt{K} \in \mathbb{F}_p$, we define the PRF `MiMC++` as the iterated Even-Mansour scheme whose round function is instantiated by $F'(x) = (x + \mathtt{K} + \gamma)^2$, where $\gamma \in \mathbb{F}_p$ is a random round constant. Again, a final key is added. Since the PRF `MiMC++` is not invertible, it must be used it in a mode in which the inverse is not needed, e.g., the CTR one

$$(x, N) \mapsto (x + \mathtt{MiMC++}_{\mathtt{K}}(N), N) \tag{11}$$

where $N \in \mathbb{F}_p$ is a nonce. For a security level of κ bits and assuming a data-limit of $2^{\kappa/2}$ texts available for the attack, the number of rounds for providing security is given by

$$R_{\mathtt{MiMC++}} = 3 + \lceil \kappa - 2 \cdot \log_2(\kappa) \rceil .$$

4.2 Security Analysis for `MiMC++`

Here we justify the number of rounds $R_{\mathtt{MiMC++}}$ just given for `MiMC++`. Since the security analysis is equivalent to the one given for MiMC, we limit ourselves to adapt the security analysis of MiMC to `MiMC++`. For this reason, we focus only on the main attacks, including the differential one, the interpolation one, the GCD one, and the linearization one (see [AGR+16] for more details). As in the case of MiMC, all other attacks do not outperform the ones just listed. Besides that, as in the case of MiMC, we explicitly state that we do **not** make any security claim in the related-key setting.

4.2.1 Statistical Attacks

Differential Attack. Given pairs of inputs with some fixed input differences, differential cryptanalysis [BS90,BS93] considers the probability distribution of the corresponding output differences produced by the cryptographic primitive. Let $\Delta_I, \Delta_O \in \mathbb{F}_p^n$ be respectively the input and the output differences through a function $\mathcal{F}$ over $\mathbb{F}_p^n$. The differential probability (DP) of having a certain output difference Δ_O given a particular input difference Δ_I is equal to

$$\text{Prob}_{\mathcal{F}}(\Delta_I \to \Delta_O) = \frac{|\{x \in \mathbb{F}_p^n \mid \mathcal{F}(x + \Delta_I) - \mathcal{F}(x) = \Delta_O\}|}{p^n} .$$

[3]For completeness, we point out that MiMC in [AGR+16] is proposed only for $d = 3$, assuming $p = 2 \mod 3$. Here, we simply generalized it for a generic $d \geq 3$ such that $\gcd(d, p-1) = 1$, by re-using the same argument proposed in [AGR+16] and by noting that currently no attack [ACG+19, EGL+20] applies on the full version of MiMC defined over a prime field.

For any non-zero input and output differences $\Delta_I, \Delta_O \in \mathbb{F}_p \setminus \{0\}$, the equality $(x+\Delta_I)^2 - x^2 = \Delta_O$ admits a single solution, that is, $x = (\Delta_O - \Delta_I^2)/(2\Delta_O)$. Hence, the maximum differential probability for 1-round `MiMC++` is $1/p$. As for MiMC, few rounds are sufficient for preventing differential attacks based on trails with non-zero differences.

However, since $x \mapsto x^2$ is not invertible, a collision can occur at every round. Based on the result proposed in Prop. 1, the overall probability that a collision occurs is upper bounded by

$$\frac{2^{R_{\texttt{MiMC++}}} - 1}{p-1} = \frac{2^{3+\lceil \kappa - 2\cdot \log_2(\kappa)\rceil} - 1}{p-1} \approx \frac{8 \cdot 2^\kappa}{p \cdot \kappa^2}\,.$$

Since $p \geq 2^{3\cdot\kappa}$, such probability is always smaller than

$$\frac{8\cdot 2^\kappa}{p\cdot\kappa^2} \leq \underbrace{\frac{8}{\kappa^2}}_{<1} \cdot 2^{-2\cdot\kappa} < 2^{-2\cdot\kappa}\,.$$

Since at most $2^{\kappa/2}$ texts are available for the attack, the attacker can construct at most $\binom{2^{\kappa/2}}{2} \approx 2^{\kappa-1}$ different pairs of texts, which implies that observing a collision is very unrealistic.

Other Statistical Attacks. As in the case of MiMC, few rounds of `MiMC++` are sufficient for preventing other statistical attacks as the linear one [Mat93], the truncated differential one [Knu94], the impossible differential one [BBS99], the boomerang one [Wag99], the invariant subspace attack [LAAZ11, LMR15], and so on.

We limit ourselves to recall that $\mathbb{F}_p$ does not admit any non-trivial subspace. Moreover, we point out that the set $\mathfrak{R} := \{x^2 \in \mathbb{F}_p \mid \forall x \in \mathbb{F}_p\}$ is **not** closed with respect to the addition. Indeed, let $x \geq 2$ be the *smallest* integer that is a non-quadratic residue modulo p (that is, $x \neq y^2$ for each $y \in \mathbb{F}_p$). Note that $0 = 0^2$ and $1 = (\pm 1)^2$ are always quadratic residue. By definition of x, $x-1$ is a quadratic residue, that is, there exists y such that $x - 1 = y^2$. Equivalently, $x = (\pm1)^2 + y^2$ and $x \cdot z^2 = (\pm z)^2 + (z \cdot y)^2$ for each $z \in \mathbb{F}_p \setminus \{0\}$, where obviously $x \cdot z^2$ is a non-quadratic residue modulo p. Hence, the sum of two elements in $\mathfrak{R}$ does not belong to such set in general.

4.2.2 Algebraic Attacks

Interpolation Attack. The interpolation attack [JK97] aims to construct an interpolation polynomial that describes the scheme. Such polynomial can be exploited in order to set up a distinguisher and/or a forgery/key-recovery attack on the symmetric scheme. The interpolation polynomial cannot be constructed if the number of unknown monomials is larger than the data available for the attack. The degree of `MiMC++` after $R_{\texttt{MiMC++}}$ rounds is $2^{R_{\texttt{MiMC++}}}$, and the number of monomials is upper bounded by $2^{R_{\texttt{MiMC++}}} + 1$. Since the data limit is $2^{\kappa/2}$, then the scheme is secure against the interpolation polynomial if $2^{R_{\texttt{MiMC++}}} \geq 2^{\kappa/2}$, that is, $R_{\texttt{MiMC++}} \geq \kappa/2$ (noting that the interpolation polynomial of `MiMC++` has the same density of the interpolation of MiMC over $\mathbb{F}_p$). One more round is added for preventing key-guessing.

We also add two more rounds for preventing interpolation attacks that make used of the Meet-in-the-Middle (MitM) approach. In particular, since the overall construction is not invertible, note that:

- the inverse function can *only* be set up locally;
- such local inverse function would have high (close to maximum) degrees.

In more details, let $\mathfrak{R} := \{x^2 \in \mathbb{F}_p \mid \forall x \in \mathbb{F}_p\}$ be the set containing the quadratic residues. Since $x \mapsto F(x) = x^2$ is not invertible, it is only possible to define "local" inverses $F_+^{-1} : \mathfrak{R} \to \mathfrak{X}_+$ and $F_-^{-1} : \mathfrak{R} \to \mathfrak{X}_-$ such that

1. the sets $\mathfrak{X}_-, \mathfrak{X}_+$ satisfy: $\mathfrak{X}_- \cap \mathfrak{X}_+ = \{0\}$ and $\mathfrak{X}_- \cup \mathfrak{X}_+ = \{0, 1, 2, \ldots, p-1\}$;
2. for each $x \in \mathfrak{X}_+ \setminus \{0\}$: $-x \in \mathfrak{X}_-$ and $-x \notin \mathfrak{X}_+$;
3. $F(F_+^{-1}(x)) = x$ and $F(F_-^{-1}(x)) = x$ for each $x \in \mathfrak{R}$.

Given $\mathfrak{X}_+$ and $\mathfrak{X}_-$ as before, other couples of sets can be potentially obtained by carefully swapping elements of $\mathfrak{X}_+$ and $\mathfrak{X}_-$. The algebraic representations of the functions F_+^{-1} and F_-^{-1} obviously depend on the sets $\mathfrak{X}_+$ and $\mathfrak{X}_-$. In general, we expect that the degrees of the functions F_+^{-1} and F_-^{-1} are of the same order of p, due to Fermat's little theorem (i.e., $x^{p-1} = 1$ for each $x \in \mathbb{F}_p \setminus \{0\}$). E.g., if $p = 3 \mod 4$, then $F_\pm^{1/2}$ can be defined as $x \mapsto \pm x^{\frac{p+1}{4}}$ over certain sets $\mathfrak{X}_\pm$, since $\pm x^{\frac{p+1}{4}}$ are the square roots of the quadratic residue x. It follows that two rounds are sufficient for reaching maximum degree in the backward direction, and so preventing MitM attacks.

GCD Attack. A dedicate attack proposed for MiMC is the GCD attack. Let's denote by $E(k, x)$ the encryption of x under the key k. Given two inputs/outputs pairs (p_0, c_0) and (p_1, c_1), let $F_0(k) := E(k, p_0) - c_0$ and $F_1(k) := E(k, p_1) - c_1$. It is easy to check that the secret key is a zero of $\gcd(F_0(k), F_1(k))$, which has in general low degree. The cost of computing the GCD of two polynomials of degree (at most) d is $\mathcal{O}(d \cdot \log^2(d))$. In our case, the cost would be $\mathcal{O}(2^{R_{\texttt{MiMC++}}} \cdot \log^2(2^{R_{\texttt{MiMC++}}}))$, which implies that the number of rounds $R_{\texttt{MiMC++}}$ must satisfy

$$2^{R_{\texttt{MiMC++}}} \cdot \log^2(2^{R_{\texttt{MiMC++}}}) = 2^{R_{\texttt{MiMC++}}} \cdot (R_{\texttt{MiMC++}})^2 \geq 2^\kappa$$

for preventing such attack. Following [AGR+16, Sect. 4.2], this means that

$$R_{\texttt{MiMC++}} \geq \kappa - 2 \cdot \log_2(\kappa) + 1$$

rounds are sufficient for preventing this attack. Due to the same argument given for the interpolation attack, we conjecture that two more rounds are sufficient for preventing the Meet-in-the-Middle version of the attack.

Linearization Attack. Linearization [KS99] is a well-known technique to solve multivariate polynomial systems of equations. Given a system of polynomial equations, the idea is to turn it into a system of linear equations by adding new variables that replace all the monomials of the system whose degree is strictly greater than 1. This linear system of equations can be solved using linear algebra if there are enough equations to make the linearized system over-determined, typically at least on the same order as the number of variables after linearization.

The most straightforward way to linearize algebraic expressions in t unknowns of degree limited by d is just by introducing a new variable for every monomial. As it is well known, the number of monomials in t variables of degree at most d is given by

$$\#(d, t) := \binom{t+d}{d}.$$

Due to this fact, the computation cost of such attacks is of $\mathcal{O}(\#(d, t)^\omega)$ operations (for $2 < \omega \leq 3$), besides a memory cost of $\mathcal{O}(\#(d, t)^2)$ for storing the linear equations.

In our case, the cost of the attack on $R_{\texttt{MiMC++}}$-round MIMC++ is given by $\left(\binom{1+2^{R_{\texttt{MiMC++}}}}{2^{R_{\texttt{MiMC++}}}}\right)^\omega \geq 2^{2 \cdot R_{\texttt{MiMC++}}}$, which is similar to the one of the interpolation attack.

Other Algebraic Attacks. As in the case of MiMC, the previous number of rounds guarantee security against other algebraic attacks, as the higher-order differential one [Lai94, Knu94, BCD+20] and the factorization attack (we recall that the complexity of factorizing a polynomial over $\mathbb{F}_{p^t}$ of degree d is $\mathcal{O}(d^3 \cdot t^2 + d \cdot t^3)$ – see [Gen07] for more details).

4.3 Multiplicative Complexity: MiMC vs. `MiMC++`

Regarding the performance, the number of $\mathbb{F}_p$-multiplications required for evaluating the PRF `MiMC++` is smaller than the corresponding one required for MiMC, even if `MiMC++` requires a larger prime p – approximately, triple size – with respect to one used in MiMC for the same security level. For comparing the performances of MiMC and of `MiMC++`, we first recall that evaluating $x \mapsto x^d$ costs $\lfloor \log_2(d) \rfloor + \text{hw}(d) - 1$ $\mathbb{F}_p$-multiplications, as showed e.g. in [GOPS22]. Based on this, the number of multiplications required for evaluating MiMC corresponds to

$$(1 + \lceil (\kappa - 2 \cdot \log_2(\kappa)) \cdot \log_d(2) \rceil) \cdot (\lfloor \log_2(d) \rfloor + \text{hw}(d) - 1)\,,$$

which satisfies[4]

$$\begin{aligned}
&(1 + \lceil (\kappa - 2 \cdot \log_2(\kappa)) \cdot \log_d(2) \rceil) \cdot (\lfloor \log_2(d) \rfloor + \text{hw}(d) - 1) \\
\geq &(\lfloor \log_2(d) \rfloor + \text{hw}(d) - 1) + \lceil (\kappa - 2 \cdot \log_2(\kappa)) \cdot \underbrace{\log_d(2) \cdot (\lfloor \log_2(d) \rfloor + \text{hw}(d) - 1)}_{>1} \rceil \\
\geq &(\lfloor \log_2(d) \rfloor + \text{hw}(d) - 1) + \lceil \kappa - 2 \cdot \log_2(\kappa) \rceil \\
\geq &2 + \lceil \kappa - 2 \cdot \log_2(\kappa) \rceil\,.
\end{aligned}$$

As a result, the number of multiplications required to evaluate MiMC is always bigger than or equal to $2 + \lceil \kappa - 2 \cdot \log_2(\kappa) \rceil$, which is almost the number of multiplications required for evaluating `MiMC++`, corresponding to $3 + \lceil \kappa - 2 \cdot \log_2(\kappa) \rceil$.

As a concrete example, consider $\kappa = 128$: MiMC (with $p \approx 2^{128}$) requires 79 rounds and 158 $\mathbb{F}_p$-multiplications, while the PRF `MiMC++` (with $p' \approx 2^{384}$) requires 126 rounds and 126 $\mathbb{F}_p$-multiplications, that is, approximately 27.5% less multiplications.

5 The MPC-Friendly PRFs Pluto and Hydra++

The main drawback of `MiMC++` regards the fact that one is forced to work with a state size p that is much larger than the security level κ in order to guarantee security. This problem does not arise when working with a scheme that is defined over $\mathbb{F}_p^n$, as HadesMiMC and Hydra. In such a case, it is possible to guarantee security for a proper choice of n, keeping p fixed. Inspired by Hydra's body, we propose the PRF Pluto,[5] a modified version of HadesMiMC in which the external rounds are instantiated via the weak bijective function $\mathcal{S}_F$ over $\mathbb{F}_p^n$ previously proposed, among other things. Hydra++ is simply defined as the PRF Hydra whose body is replaced with Pluto.

5.1 The PRFs Pluto and Hydra++

5.1.1 Preliminary: HadesMiMC and Hydra

The Cipher HadesMiMC. The Hades design strategy [GLR+20] allows to design SPN schemes over $\mathbb{F}_q^n$ that aim to reduce the overall multiplicative complexity. In order to guarantee security and maximize the efficiency, both rounds with full and partial S-Box layers are used:

- the external rounds with full S-Box layers (that is, n S-Boxes in each non-linear layer) at the beginning and at the end of the construction ensure security against statistical attacks, besides masking the internal rounds;

[4] Note that $\log_d(2) \cdot (\lfloor \log_2(d) \rfloor + \text{hw}(d) - 1) > 1$ if and only if $\lfloor \log_2(d) \rfloor + \text{hw}(d) > 1 + \log_2(d)$. Since $\lfloor \log_2(d) \rfloor > \log_2(d) - 1$, this last inequality is satisfied if and only if $\text{hw}(d) \geq 2$, i.e., odd $d \geq 3$.

[5] Following the strategy of calling an updated modified version of Poseidon as Neptune proposed in [GOPS22], we decided to call this new version as Pluto, which is the Roman name of the Greek god Hades.

- the internal rounds with partial S-Box layers (that is, 1 S-Box and $n-1$ identity functions) in the middle of the construction increase the overall degree of the scheme ensuring security against algebraic attacks, besides being cheaper to evaluate.

Let $p > 2^{63}$ (or equivalently, $\lceil \log_2(p) \rceil \geq 64$), and let $n \geq 2$. Let $\mathtt{K} \in \mathbb{F}_p^n$ be the secret master key. Let κ be the security level such that $2^{80} \leq 2^\kappa \leq \min\{p^2, 2^{256}\}$. Let $d \geq 3$ be the smallest integer such that $\gcd(d, p-1) = 1$. The block cipher HadesMiMC $\mathcal{H}_\mathtt{K} : \mathbb{F}_p^n \to \mathbb{F}_p^n$ is defined as[6]

$$\mathcal{H}_\mathtt{K}(x) = \underbrace{\mathcal{E}_{R_f+R'_f-1} \circ \cdots \circ \mathcal{E}_{R_f}}_{=R'_f \text{ rounds}} \circ \underbrace{\mathcal{I}_{R_P-1} \circ \cdots \circ \mathcal{I}_0}_{=R_P \text{ rounds}} \circ \underbrace{\mathcal{E}_{R_f-1} \circ \cdots \circ \mathcal{E}_0}_{=R_f \text{ rounds}}(x + \mathtt{K})$$

where

$$\mathcal{E}_i(x) = k_i + M_\mathcal{E} \times \mathcal{S}_\mathcal{E}(x) \qquad \text{and} \qquad \mathcal{I}_j(x) = k_j + M_\mathcal{I} \times S_\mathcal{I}(x)$$

for each $i \in \{0, 1, \ldots, R_f + R'_f - 1\}$ and $j \in \{0, 1, \ldots, R_\mathcal{I} - 1\}$, such that

- the external non-linear layer is defined as $\mathcal{S}_\mathcal{E}(x_0, x_1, \ldots, x_{n-1}) = x_0^d \| x_1^d \| \ldots \| x_{n-1}^d$;
- the internal non-linear layer is defined as $\mathcal{S}_\mathcal{I}(x_0, x_1, \ldots, x_{n-1}) = x_0^d \| x_1 \| \ldots \| x_{n-1}$ (that is, the power map is applied only on the first component);
- $M_\mathcal{E} = M_\mathcal{I} \in \mathbb{F}_p^{n \times n}$ is a MDS matrix that prevents the existence of invariants subspace trails for the internal rounds – we refer to [GRS21] for a detailed description on how to choose such matrices;
- $k_i, k_j \in \mathbb{F}_p^n$ are the round sub–keys derived from the master key via an affine key–schedule of the form

$$k_i = (M_\mathcal{K})^i \times \mathtt{K} + \varphi_i \,, \qquad k_j = (M'_\mathcal{K})^j \times \mathtt{K} + \varphi'_j \tag{12}$$

  for invertible matrices $M_\mathcal{K}, M'_\mathcal{K} \in \mathbb{F}_p^{n \times n}$ and for random round constants $\varphi_i, \varphi'_j \in \mathbb{F}_p^n$ – we refer to [GLR+20, Sect. 3] for all details.

Let $2^{40} \leq 2^{\kappa/2} \leq \min\{p, 2^{128}\}$ be the data limit available for the attack. The number of rounds are given by $R_f = R'_f = 3$ and $R_P = 4 + \left\lceil \frac{\kappa}{2 \cdot \log_2(d)} \right\rceil + \lceil \log_d(n) \rceil$.[7]

The Body of Hydra. The PRF Hydra is based on the Megafono mode of operation recently introduced in [GØWS22], a modified version of the Farfalle mode of operation [BDH+17] suitable for MPC applications. As already recalled in the introduction, a scheme based on the Megafono mode of operation is composed of two phases, that is, (1st) an initial phase in which the input is mixed with the secret key via a PRP, and (2nd) an expansion phase in which the state is expanded until the desired state size is reached. We refer to [GØWS22] for more details.

The initial phase – called body – is instantiated via an Even–Mansour construction of the form

$$x \mapsto \mathtt{K} + \mathcal{B}(x + \mathtt{K}) \,, \tag{13}$$

where $\mathtt{K}$ is the secret key, and $\mathcal{B}$ is an *unkeyed* permutation. *In the case in which such initial phase is indistinguishable from a PRP*, the security of the entire construction can be heavily simplified since only few attacks apply, as a consequence of the facts that

[6] In [GLR+20], authors use the nomenclature "Full" and "Partial" rounds for referring respectively to the "External" and the "Internal" rounds. This new nomenclature has been introduced in [GOPS22].

[7] In [GLR+20], the number of rounds are provided only for the cases in which either $p \approx 2^\kappa$ or $p^n \approx 2^\kappa$. The number of rounds given here is a simple generalization of the number of rounds given in the original paper [GLR+20].

(i) the attacker does not have access to the internal states of the construction, as the inputs of the expansion phase (besides not being able to choose the outputs for e.g. a chosen ciphertext attacks), and (ii) the inputs of the expansion phase (equivalently, the outputs of the initial phase) do not have any algebraic/statistical structure. In such a scenario, one of the most powerful attacks is usually set up by considering the relation of the outputs of consecutive calls of the expansion phase permutations on related *unknown* inputs [CFG+18, CG20, BBLP22].

For the goal of this paper, we focus on the body part only. In the case of Hydra, the permutation $\mathcal{B}$ that instantiates the initial phase is based on the Hades design strategy. However, the main differences between the body of Hydra and HadesMiMC are the following:

- the body of Hydra is defined over $\mathbb{F}_p^4$, that is, $n = 4$ fixed;
- the external rounds are instantiated via power maps. Instead, the non-linear layer $\mathcal{S}_\mathcal{I}$ of the internal rounds are instantiated via a degree $2^l \geq 2$ Lai–Massey scheme that can be computed via only l $\mathbb{F}_p$-multiplications:

$$\mathcal{S}_\mathcal{I}(x_0, x_1, \ldots, x_{n-1}) = x_0 + z \| x_1 + z \| \ldots \| x_{n-1} + z \qquad \text{where}$$
$$z = \left(\left(\sum_{i=0}^{n-1} \lambda_i \cdot x_i\right)^2 + \left(\sum_{i=0}^{n-1} \lambda'_i \cdot x_i\right) + \psi\right) \cdot \left(\left(\sum_{i=0}^{n-1} \lambda_i \cdot x_i\right)^2 + \left(\sum_{i=0}^{n-1} \lambda'_i \cdot x_i\right) + \psi'\right) \tag{14}$$

  where $\lambda_0, \ldots, \lambda_{n-1}, \lambda'_0, \ldots, \lambda'_{n-1}, \psi, \psi' \in \mathbb{F}_p$ satisfy $\sum_{i=0}^{n-1} \lambda_i = \sum_{i=0}^{n-1} \lambda'_i = 0$;
- $M_\mathcal{E} \in \mathbb{F}_p^{n \times n}$ is a MDS matrix, while $M_\mathcal{I} \in \mathbb{F}_p^{n \times n}$ is an invertible matrix that aims for destroying the invariant subspace trails of the Lai–Massey construction – we refer to [GRS21, GØWS22] for all details about $M_\mathcal{I}$;
- the round sub–keys are replaced by random round–constants, and a key addition with the master key takes place only at the beginning and at the end of the permutation call (as in (13)).

Using the previous notation, the number of rounds are given by $R_f = 2$, $R_{f'} = 4$, and $R_P = \lceil 1.125 \cdot \lceil \frac{\kappa}{4} + 6 - \log_2(d) \rceil \rceil$, including a security margin of 12.5% for the internal rounds.

5.1.2 The PRFs Pluto and Hydra++

The PRF Pluto. Here, we propose the PRF Pluto as a modified version of the Hydra's body. Let $p > 2^{63}$ (or equivalently, $\lceil \log_2(p) \rceil \geq 64$), and let $n \geq 4$. Let $\mathtt{K} \in \mathbb{F}_p^n$ be the secret master key. Let κ be the security level such that $2^{80} \leq 2^\kappa \leq \min\{p^2, 2^{256}\}$. The keyed PRF Pluto $\mathcal{P}_\mathtt{K} : \mathbb{F}_p^n \to \mathbb{F}_p^n$ is defined as

$$\mathcal{P}_\mathtt{K}(x) = \underbrace{\mathcal{E}_{R_\mathcal{E}+R_{\mathcal{E}'}-1} \circ \cdots \circ \mathcal{E}_{R_\mathcal{E}}}_{=R_{\mathcal{E}'} \text{ rounds}} \circ \underbrace{\mathcal{I}_{R_\mathcal{I}-1} \circ \cdots \circ \mathcal{I}_0}_{=R_\mathcal{I} \text{ rounds}} \circ \underbrace{\mathcal{E}_{R_\mathcal{E}-1} \circ \cdots \circ \mathcal{E}_0}_{=R_\mathcal{E} \text{ rounds}}(x + \mathtt{K})$$

where

$$\mathcal{E}_i(x) = k_i + M_\mathcal{E} \times \mathcal{S}_\mathcal{E}(x) \qquad \text{and} \qquad \mathcal{I}_j(x) = k_j + M_\mathcal{I} \times S_\mathcal{I}(x)$$

for $i \in \{0, 1, \ldots, R_\mathcal{E} + R_{\mathcal{E}'} - 1\}$ and $j \in \{0, 1, \ldots, R_\mathcal{I} - 1\}$ such that

- the external non-linear layer $\mathcal{S}_\mathcal{E}$ is instantiated by the SI-lifting function $\mathcal{S}_F$ over $\mathbb{F}_p^n$ induced by $F(x_0, x_1) = x_0^2 + x_1$;
- the internal non-linear layer $\mathcal{S}_\mathcal{I}$ is instantiated via the Lai–Massey function defined in (14);

- $M_{\mathcal{E}} \in \mathbb{F}_p^{n\times n}$ is a MDS matrix, while $M_{\mathcal{I}} \in \mathbb{F}_p^{n\times n}$ is an invertible matrix that aims for destroying the invariant subspace trails of the Lai–Massey construction – we refer to [GRS21, GØWS22] for all details about $M_{\mathcal{I}}$;
- k_i, k_j are the round sub–keys defined as

$$k_i = \mathtt{K} + \varphi_i\,, \qquad k_i = \mathtt{K} + \varphi'_j$$

for random round constants $\varphi_i, \varphi'_j \in \mathbb{F}_p^n$.

Obviously, this construction is not invertible anymore, but it can be used as a stream cipher for encryption purpose, exactly as in the case of `MiMC++` – see (11). The number of rounds are given by

$$R_{\mathcal{E}} = R_{\mathcal{E}'} = 4 \qquad \text{and} \qquad R_{\mathcal{I}} = \left\lceil 1.125 \cdot \left\lceil \frac{\kappa}{4} + \frac{n}{2} + \log_2(n) + 1 \right\rceil \right\rceil\,,$$

where we add an arbitrary security margin of 12.5% for the internal rounds.

The PRF Hydra++. The keyed PRF Hydra++ is defined as the PRF Hydra whose body is replaced with the keyed PRF Pluto just defined over $\mathbb{F}_p^4$.

Remark 1. We point out that the body of the PRF Hydra is instantiated via an Even–Mansour construction $x \mapsto \mathtt{K} + \mathcal{B}(x + \mathtt{K})$, where $\mathcal{B}$ is a permutation that is independent of the secret key. In the case of Hydra++, its body is instantiated with a keyed iterated PRF in which the key addition takes place in each round. We are not aware of any attack on Hydra++ that exploits such difference.

5.2 Security Analysis of Pluto

Here we justify the number of rounds just given for Pluto (and Hydra++). Since the security analysis is equivalent to the one given for HadesMiMC/Hydra's body, we limit ourselves to adapt the security analysis of HadesMiMC/Hydra's body to Pluto. (We refer to Sect. 4.2 for the description of the attacks analyzed here.) We explicitly state that we do **not** make any security claim in the related-key setting.

5.2.1 Statistical Attacks

Let $\mathcal{A}$ over $\mathbb{F}_p^n$ be an invertible affine transformation. As in the case of HadesMiMC, our goal is to show that

$$x \mapsto \underbrace{\mathcal{E}_7 \circ \cdots \circ \mathcal{E}_4}_{=4 \text{ rounds}} \circ \mathcal{A} \circ \underbrace{\mathcal{E}_3 \circ \cdots \circ \mathcal{E}_0}_{=4 \text{ rounds}}(x) \tag{15}$$

is secure against statistical attacks. The security of Pluto follows from the fact that the security of this "weaker" scheme (15) is not reduced when $\mathcal{A}$ is replaced by internal invertible non-linear rounds (note that the internal rounds of Pluto are invertible). As in the case of HadesMiMC, two of the eight external rounds of Pluto aim to frustrate partial key–guessing attacks.

Remark 2. Before going on, we clarify the choice of the key–schedule of Pluto compared to the one of HadesMiMC. Let's assume that the attacker partially guesses one sub–key k_i. In the case of HadesMiMC, due to its linear key–schedule (12), the attacker only partially knows the relation between the entries of other sub–keys k_j for $j \neq i$, but not the exact values of its entry (in general). This choice aims to frustrate attacks in which the attacker partially guesses multiple sub–keys. Due to the conditions $p > 2^{64}$, $n \geq 4$, and $\kappa \leq \min\{2 \cdot \log_2(p), 256\}$, here we claim that a simpler key–schedule (defined via random round constants addition) is sufficient for reaching the same goal.

Differential Attack. First of all, we analyze the probability that a collision occurs in PLUTO, keeping in mind that $\mathcal{S}_F$ is a weak bijective (non-invertible) function. Based on the result proposed in Prop. 1, the probability that a collision occurs is upper bounded by

$$\frac{2^8-1}{p^n-1} \le \frac{2^8-1}{p^4-1} \approx \frac{2^8}{p^4} \le \frac{2^8}{2^{2\cdot\kappa}} = 2^{-2\cdot\kappa+8}\,,$$

where the number of external rounds is 8, $n \ge 4$, and $2^\kappa \le p^2$. Since at most $2^{\kappa/2}$ texts are available for the attack, the attacker can construct at most $\binom{2^{\kappa/2}}{2} \approx 2^{\kappa-1}$ different pairs of texts, which implies that observing a collision is very unrealistic.

Next, let's consider the case of a differential characteristic without collision. Let $\Delta^I, \Delta^O \in \mathbb{F}_p^n \setminus \{0\}$ be respectively an input/output (non-null) difference. The system of equations $\mathcal{S}_F(x+\Delta^I) - \mathcal{S}_F(x) = \Delta^O$ is satisfied by $x = (s-\Delta^I)/2 \in \mathbb{F}_p^n$ where

$$\begin{bmatrix} 0 & \Delta_1^I & 0 & \dots & 0 \\ 0 & 0 & \Delta_2^I & \dots & 0 \\ \vdots & & & \ddots & \vdots \\ 0 & 0 & 0 & \dots & \Delta_{n-1}^I \\ \Delta_0^I & 0 & 0 & \dots & 0 \end{bmatrix} \times \begin{bmatrix} s_0 \\ s_1 \\ s_2 \\ \vdots \\ s_{n-1} \end{bmatrix} = - \begin{bmatrix} \alpha_{1,0}\cdot\Delta_0^I + \alpha_{0,1}\cdot\Delta_1^I - \Delta_0^O \\ \alpha_{1,0}\cdot\Delta_1^I + \alpha_{0,1}\cdot\Delta_2^I - \Delta_1^O \\ \alpha_{1,0}\cdot\Delta_2^I + \alpha_{0,1}\cdot\Delta_3 - \Delta_2^O \\ \vdots \\ \alpha_{1,0}\cdot\Delta_{n-1}^I + \alpha_{0,1}\cdot\Delta_0^I - \Delta_{n-1}^O \end{bmatrix}.$$

Since $\Delta^O \ne 0 \in \mathbb{F}_p^n$, for each $i \in \{0,1,\dots,n-1\}$:

- if $\Delta_i^I = 0$, then the i-th equality is satisfied if and only if $\alpha_{1,0}\cdot\Delta_{i-1}^I = \Delta_i^O$;
- if $\Delta_i^I \ne 0$, then the i-th equality is satisfied if $s_i = -(\alpha_{1,0}\cdot\Delta_{i-1}^I + \alpha_{0,1}\cdot\Delta_i^I - \Delta_i^O)/\Delta_i^I$.

Hence, the number of different solutions of $\mathcal{S}_F(x+\Delta^I) - \mathcal{S}_F(x) = \Delta^O$ is *at most* equal to p^z, where $0 \le z \le n-1$ is the number of zero $\mathbb{F}_p$-components of Δ^I.

Let's now consider two consecutive rounds, and let's introduce:

- $\mathtt{a}_0$:= number of active (i.e., non-zero) $\mathbb{F}_p$-components at the *input* of the *first* non-linear layer $\mathcal{S}_F$;
- $\mathtt{a}_1$:= number of active (i.e., non-zero) $\mathbb{F}_p$-components at the *output* of the *first* non-linear layer $\mathcal{S}_F$;
- $\mathtt{a}_2$:= number of active (i.e., non-zero) $\mathbb{F}_p$-components at the *input* of the *second* non-linear layer $\mathcal{S}_F$.

Given $\mathtt{a}_0 \ge 1$ active inputs $\mathbb{F}_p$-components, then at most $\mathtt{a}_1 \le \min\{2\cdot\mathtt{a}_0, n\}$ $\mathbb{F}_p$-components are active at the output of $\mathcal{S}_F$. In particular, note that if the $\mathtt{a}_0 \le n/2$ active input $\mathbb{F}_p$-components are not in consecutive positions, then at most $2\cdot\mathtt{a}_0$ are active in output. Since $M_\mathcal{E} \in \mathbb{F}_p^{n\times n}$ is a MDS matrix (hence, its branch number is $n+1$ – see e.g. [DR01,DR02a] for details), then at least $\mathtt{a}_2 \ge n+1-\mathtt{a}_1$ $\mathbb{F}_p$-components are active at the input of the second round. Over two consecutive rounds, the probability of a differential trail is approximately given by

$$\begin{aligned}\frac{1}{p^{2\cdot n}} \cdot \underbrace{p^{n-\mathtt{a}_0}}_{\text{1st round}} \cdot \underbrace{p^{n-\mathtt{a}_2}}_{\text{2nd round}} &= p^{-\mathtt{a}_0-\mathtt{a}_2} \le p^{-\mathtt{a}_0+\mathtt{a}_1-n-1} \\ &\le p^{-\mathtt{a}_0+\min\{2\mathtt{a}_0,n\}-n-1} = p^{\min\{\mathtt{a}_0-n-1,-1-\mathtt{a}_0\}}\,,\end{aligned}$$

since $\mathtt{a}_2 \ge n+1-\mathtt{a}_1 \ge 1$ and since $\mathtt{a}_1 \le \min\{2\cdot\mathtt{a}_0, n\}$. It follows that the probability over two rounds is upper bounded by

$$\max_{1\le\mathtt{a}_0\le n} p^{\min\{\mathtt{a}_0-n-1,-1-\mathtt{a}_0\}} = \max\left\{\max_{1\le\mathtt{a}_0\le\lfloor n/2\rfloor} p^{\mathtt{a}_0-n-1}, \max_{\lceil n/2\rceil\le\mathtt{a}_0\le n} p^{-1-\mathtt{a}_0}\right\} = p^{-\lceil n/2\rceil-1}\,.$$

Given two consecutive rounds per two times, the probability of any differential trail based on trails with non-zero differences is at most equal to

$$\left(p^{-\lceil n/2\rceil-1}\right)^2 \leq p^{-2}\cdot 2^{-\frac{2}{4}\cdot n\cdot\kappa} \leq p^{-2}\cdot 2^{-2\cdot\kappa}$$

since $2^\kappa \leq p^2$ and since $n \geq 4$, which is much smaller than the security level. Moreover, besides the external rounds $\mathcal{E}$, the internal rounds $\mathcal{I}$ guarantee security against classical differential attack as well, as pointed out in e.g. [KR21]. In particular, since no invariant subspace with no active non-linear function can cover $R_{\mathcal{I}}/2$ (or more) internal rounds (see [GØWS22] for details), then the probability of each characteristic of four external rounds and $R_{\mathcal{I}}$ internal rounds is upper bounded by

$$p^{-2}\cdot 2^{-2\cdot\kappa}\cdot\left(\frac{3}{p}\right)^{\left\lfloor\frac{R_{\mathcal{I}}}{2}\right\rfloor},$$

where $3/p$ is the maximum differential probability of $\mathcal{S}_{\mathcal{I}}$, as proved in [GØWS22, App. H]. Based on this, we conclude that Pluto instantiated with eight external rounds is secure against classical differential attacks with a data limit of $2^{\kappa/2}$ texts available for the attacker.

Other Statistical Attacks. As in the case of HadesMiMC, eight external rounds are sufficient for preventing other statistical attacks, including the linear one [Mat93], the truncated differential one [Knu94], the impossible differential [Knu98, BBS99], the boomerang attack [Wag99], the integral one [DKR97], the multiple-of-n/mixture differential [GRR17,Gra18], among others. This follows from the fact that no truncated differential with probability 1 can cover more than a single external round, due to the facts that (1st) $M_{\mathcal{E}}$ is an MDS matrix and (2nd) $\mathcal{S}_F$ is a full non-linear layer.

5.2.2 Algebraic Attacks

Interpolation Attack. As explained in Sect. 4.2.2, a primitive can be considered secure against the interpolation attack [JK97] if the number of unknown monomials that defines the scheme is larger than the data available to the attacker. As before, we assume that two of the eight external rounds of Pluto aim to frustrate partial key–guessing attacks.

Focusing on the backward direction, the function $\mathcal{S}_F$ is not invertible. Working as in Sect. 4.2.2, it is possible to define three sets $\mathfrak{X}_+, \mathfrak{X}_-, \mathfrak{Z} \subset \mathbb{F}_p^n$ so that

- given $z \in \mathfrak{Z}$, then $\mathcal{S}_F(y) \neq \mathcal{S}_F(z)$ for each $y \in \mathbb{F}_p^n \setminus \{z\}$;
- given $x, x' \in \mathbb{F}_p^n \setminus \mathfrak{Z}$ such that $\mathcal{S}_F(x) = \mathcal{S}_F(x')$ and $x \neq x'$, then (i) $x \in \mathfrak{X}_+$ and $x \notin \mathfrak{X}_-$ and (ii) $x' \in \mathfrak{X}_-$ and $x \notin \mathfrak{X}_+$.

It follows that $\mathfrak{X}_+ \cup \mathfrak{X}_- \cup \mathfrak{Z} = \mathbb{F}_p^n$ and that $\mathfrak{X}_+ \cap \mathfrak{X}_- = \mathfrak{X}_+ \cap \mathfrak{Z} = \mathfrak{X}_- \cap \mathfrak{Z} = \emptyset$. Moreover, since a collision can occur if and only if $x_i \neq x'_i$ for each $i \in \{0, 1, \dots, n-1\}$, we have that

$$|\mathfrak{X}_+| = |\mathfrak{X}_-| = \frac{(p-1)^n}{2} \approx \frac{p^{n-1}\cdot(p-n)}{2} \qquad \text{and} \qquad |\mathfrak{Z}| = p^n - (p-1)^n \approx n\cdot p^{n-1}.$$

As in the case of $x \mapsto x^2$, while $\mathfrak{Z}$ is uniquely defined, there are several equivalent representations of $\mathfrak{X}_+$ and $\mathfrak{X}_-$ (by carefully swapping two elements x and x' as before). It follows that only *local* inverses can be defined (e.g., from $\mathbb{F}_p^n$ into $\mathfrak{Z} \cup \mathfrak{X}_\pm$), where (1st) their algebraic expressions depend on such representations and (2nd) they are in general of high degree. For these reasons, we conjecture that the last three external rounds $\mathcal{E}$ are sufficient for stopping backward interpolation attacks.

Focusing on the forward direction,[8] the degree growths as $2^2 \cdot 4^{R_\mathcal{I} - n/2 - \log_2(n) - 2}$, where we discount (i) one external round and $\log_2(n) + 1$ internal rounds in order to destroy possible relations existing between the coefficients of the monomials (and so, to ensure full diffusion) and (ii) $n/2 + 1$ extra internal rounds, since an attacker can cover at most $n/2$ internal rounds via an invariant subspace without activating any function of $\mathcal{I}$ (see [GØWS22] for more details).[9] As a result:

$$2^2 \cdot 4^{R_\mathcal{I} - n/2 - \log_2(n) - 2} \geq 2^{\kappa/2} \qquad \rightarrow \qquad R_\mathcal{I} \geq \frac{\kappa}{4} + \frac{n}{2} + \log_2(n) + 1 \,.$$

Other Algebraic Attacks. As in the case of HadesMiMC and based on the same argument proposed in Sect. 4.2.2, the security against the MitM interpolation attack implies security against higher-order differential attack [BCD+20], the linearization attack [KS99] and the Gröbner basis one [Buc76]. Without going into the details, the cost of a Gröbner basis attack depends on several factors, including (i) the number of non-linear equations that composed the system to solve, (ii) the number of independent variables, and (iii) the degree of each equations to solve. Moreover, the cost of a Gröbner basis attack depends on the considered representative of the system of equations. Due to the analogous Gröbner basis attack on HadesMiMC and Hydra given in [GLR+20, Sect. 4] and in [GØWS22, Sect. 7], two main strategies are possible for setting up a Gröbner basis attack:

- working with a system of equations that involved only the inputs/outputs of the entire permutation;
- considering a system of equations defined at round level.

In the first case, the number of variables if fixed, and the attacker can collect more equations than the number of possible monomials. In such a case, Gröbner basis attack reduces to a linearization attack, which does not outperform the interpolation attack just described (see the analysis proposed for `MiMC++` in Sect. 4.2.2). In the second case, the number of variables is proportional to the number of rounds. Due to the analogous result proposed for Hydra and HadesMiMC, we can conclude that the cost of such strategy is higher than the security level. Other approaches do not seem to be competitive as the ones just discussed.

5.3 Multiplicative Complexity: HadesMiMC/Hydra vs. Pluto/Hydra++

Let n be the size of the text to be encrypted. By simple computation, the number of multiplications required to evaluate HadesMiMC is

$$\underbrace{\mathbf{6 \cdot (hw(d) + \lfloor \log_2(d) \rfloor - 1)}}_{\geq 12} \cdot \mathbf{n} + \underbrace{(\text{hw}(d) + \lfloor \log_2(d) \rfloor - 1) \cdot \left(4 + \left\lceil \frac{\kappa}{2 \cdot \log_2(d)} \right\rceil + \lceil \log_d(n) \rceil \right)}_{\approx \text{ constant w.r.t. } n}$$

while the number of multiplications required to evaluate Pluto is (approximately)

$$\mathbf{9.125} \cdot \mathbf{n} + \underbrace{2 \cdot \left\lceil 1.125 \cdot \left\lceil \frac{\kappa}{4} + \log_2(n) + 1 \right\rceil \right\rceil}_{\approx \text{ constant w.r.t. } n}$$

We can notice that:

- in Pluto, the factor that multiplies n is fixed and (approximately) equal to 9, since its external rounds are instantiated with a quadratic function *independently* of p;

[8] Remember that $\deg(\mathcal{E}_i) = \deg(\mathcal{S}_F) = 2$ and that $\deg(\mathcal{I}_j) = \deg(\mathcal{S}_\mathcal{I}) = 4$.
[9] Note that we already discount three external rounds in order to prevent backward and MitM attacks.

- in HadesMiMC, such factor depends on the value of d, and it is never smaller than 12.

As a result, we reached the goal of reducing the number of multiplications of HadesMiMC without decreasing its security.

Hydra vs. Hydra++. A similar conclusion holds when comparing the body of Hydra versus the body of Hydra++, for which we remember that $n = 4$ is fixed. E.g., for the most common case $p \approx 2^{128}$, $\kappa = 128$ and $d = 3$, Hydra++'s body requires 116 $\mathbb{F}_p$-multiplications versus 132 $\mathbb{F}_p$-multiplications for the Hydra's body (that is, 13.8% more). The gap growths for bigger values of d.

6 HE-friendly Schemes: Implications on Masta, Pasta, and `Rubato`

Masta [HKC+20], Pasta [DGH+21], and `Rubato` [HKL+22] are recent PRFs over $\mathbb{F}_p^n$ proposed for Homomorphic Encryption. In the following, we recall these schemes, and we show that it is possible to achieve better performance and/or security by modifying them with the results proposed in this paper. Since all these schemes are inspired by Rasta [DEG+18], we first recall it for pointing out the main common design strategy of all these HE–friendly PRFs.

Preliminary: Rasta. Rasta is a family of HE-friendly stream ciphers over $\mathbb{F}_2^n$ for odd n proposed at Crypto 2018. Given an input $x \in \mathbb{F}_2^n$ to encrypt, a public nonce $N \in \mathbb{F}_2^n$ and a public block index counter $i \in \mathbb{N}$, the ciphertext is generated as

$$(x, N) \mapsto (x + \mathtt{K} + \mathcal{P}_{N,i}(\mathtt{K}), N)$$

for a secret key $\mathtt{K} \in \mathbb{F}_2^n$. The public permutation $\mathcal{P}_{N,i} : \mathbb{F}_2^n \to \mathbb{F}_2^n$ consists of several rounds $r \geq 1$ of affine layers and non-linear layers of the form

$$\mathcal{P}_{N,i}(\cdot) = \mathcal{A}_{r,N,i} \circ \mathcal{S}_\chi \circ \ldots \circ \mathcal{A}_{1,N,i} \circ \mathcal{S}_\chi \circ \mathcal{A}_{0,N,i}(\cdot)\,, \tag{16}$$

where

- $\mathcal{S}_\chi$ over $\mathbb{F}_2^n$ is the SI-lifting function induced by the local map $\chi : \mathbb{F}_2^3 \to \mathbb{F}_2$ defined as $\chi(x_0, x_1, x_2) = x_0 + x_2 + x_1 \cdot x_2$. We recall that $\mathcal{S}_\chi$ is invertible for odd $n \geq 3$;
- for each $j \in \{0, 1, \ldots, r\}$, $\mathcal{A}_{j,N,i} = M_{j,N,i} \times x + c_{j,N,i}$ is an affine function over $\mathbb{F}_2^n$, where $M_{j,N,i} \in \mathbb{F}_2^{n\times n}$ is an invertible matrix and $c_{j,N,i} \in \mathbb{F}_2^n$.

In order to minimize the depth, the design strategy adapted for Rasta is quite different from the one usually adopted by "traditional" symmetric primitives. In general, given the size n and the security level κ, the number of rounds R of a symmetric primitive is chosen in order to guarantee security (e.g., so that no known attack published in the literature can break the scheme, besides a possible security margin). Exactly the opposite occurs for Rasta. In order to minimize the depth (note that each round has depth one, since $\mathcal{S}_\chi$ is a quadratic function and $A_{j,N,i}$ is an affine function), given the number of rounds R and the security level κ, the size n is chosen in order to guarantee security, that is, in order to frustrate any possible attack on the scheme. This usually results in huge state size compared to "traditional" symmetric primitives.

Besides that, another crucial feature of Rasta regards the affine layers $A_{j,N,i}$, which are *not* fixed. At each new encryption, new random invertible affine layers $A_{0,N,i}, \ldots, A_{r,N,i}$ are generated via a public XOF that takes in input the nonce N and the counter i. This fact has a crucial impact on the security against statistical attacks, as linear or differential attacks. In the case of a "traditional" encryption scheme, given a set of inputs and corresponding outputs encrypted via the same algorithm, the attacker performs statistical analysis on the output distribution in order to break the scheme. However, such strategy does not work in the case of Rasta, since *each input is encrypted via a different encryption scheme*.

As a result, the main attack vector against Rasta results the linearization one. In a linearization approach, the attacker replaces all monomials of degrees greater than one by new variables, and finally tries to solve the resulting system of linear equations. The cost of the attack is proportional to the number of (non-linear) monomials that define the analyzed function, as already given in Sect. 4.2.2.

For completeness, we point out that other attacks on Rasta have been recently proposed in the literature [DMRS20, LSMI21, LSMI22]. However, since such attacks exploit the details of the non-linear $\mathcal{S}_\chi$ over $\mathbb{F}_2^n$ and they (currently) do not apply to the prime field case, we do not discuss them in this context.

About Masta. Masta can be seen as a direct translation of Rasta to $\mathbb{F}_p^n$ for a prime integer $p \geq 3$. Both in Rasta and in Masta, the non-linear layer is defined via the SI-lifting function $\mathcal{S}_\chi$ over the entire state $\mathbb{F}_q^n$ (where $q = 2$ for Rasta and $q = p$ for Masta) instantiated via the chi function $\chi : \mathbb{F}_q^3 \to \mathbb{F}_q$ defined as before. The main difference between Masta and Rasta regards the way in which the invertible matrices $M_{j,N,i}$ that define the affine layers are generated.

Our result proposed in Sect. 3.2 implies that *$\mathcal{S}_\chi$ over $\mathbb{F}_p^n$ for $p \geq 3$ and $n \geq 3$ is never invertible*. This can be easily proven by adapting the proof just given for such function. In the analyzed case, the equality given in (6) and corresponding to the collision $\mathcal{S}_\chi(x) = \mathcal{S}_\chi(y)$ re-written via the variables $d, s \in \mathbb{F}_p^n$ becomes

$$\begin{bmatrix} 0 & d_2 & d_1 & 0 & \dots & 0 \\ 0 & 0 & d_3 & d_2 & \dots & 0 \\ \vdots & & & \ddots & \ddots & \vdots \\ 0 & 0 & 0 & \dots & d_{n-1} & d_{n-2} \\ d_{n-1} & 0 & 0 & \dots & 0 & d_0 \\ d_1 & d_0 & 0 & \dots & 0 & 0 \end{bmatrix} \times \begin{bmatrix} s_0 \\ s_1 \\ s_2 \\ \vdots \\ s_{n-2} \\ s_{n-1} \end{bmatrix} = - \begin{bmatrix} d_0 + d_2 \\ d_1 + d_3 \\ d_2 + d_4 \\ \vdots \\ d_{n-2} + d_0 \\ d_{n-1} + d_1 \end{bmatrix}.$$

Note that the l.h.s. matrix in this equality corresponds to the l.h.s. matrix in (6) after a re-arrangement of the rows. Since the collision event $\mathcal{S}_\chi(x) = \mathcal{S}_\chi(y)$ only depends on the details of such matrix, the result follows immediately.

Hence, by replacing $\mathcal{S}_\chi$ in Masta with the SI-lifting function $\mathcal{S}_F$ over $\mathbb{F}_p^n$ induced by $F(x_0, x_1) = x_0^2 + x_1$, we would have the following advantages:

- the costs of $\mathcal{S}_\chi$ and of $\mathcal{S}_F$ in terms of multiplications is equal;
- the resistant of Masta against linearization attacks does not change, since the number of quadratic monomials of $\mathcal{S}_F$ and of $\mathcal{S}_\chi$ are equal;
- the probability that a collision occurs is smaller.

Regarding this last point, at the current state of the art, no attack on Masta based on the fact that $\mathcal{S}_\chi$ is not invertible has been proposed in the literature. As in the case of Rasta, this is related to the fact that statistical attacks are frustrated by the change of the affine layer at every encryption. At the same time, the proposed change allows to reduce the collision probability without any other counter–effect.

About Pasta. With respect to Masta, Pasta is a variant of Rasta over $\mathbb{F}_p^n$ instantiated with invertible non-linear layers only, and in which the feed–forward is replaced by a final truncation. That is, given an input $x \in \mathbb{F}_p^n$ to encrypt, a public nonce $N \in \mathbb{F}_p^n$ and a public block index counter $i \in \mathbb{N}$, the ciphertext is generated as

$$(x, N) \mapsto (x + \mathcal{T}_{2n,n} \circ \mathcal{P}_{N,i}(\mathtt{K}), N)$$

for a secret key $\mathtt{K} \in \mathbb{F}_p^{2n}$, a public permutation $\mathcal{P}_{N,i} : \mathbb{F}_p^{2n} \to \mathbb{F}_p^{2n}$, and a truncation function $\mathcal{T}_{2n,n} : \mathbb{F}_p^{2n} \to \mathbb{F}_p^n$. As in the case of Rasta, the public permutation $\mathcal{P}_{N,i}$ consists of several rounds $R \geq 1$ of affine layers and non-linear layers of the form (16), where the non-linear layers in the first $R - 1$ rounds is instantiated via two parallel (independent) Type–III Feistel schemes over $\mathbb{F}_p^n$ of the form

$$(x_0, \dots, x_{n-2}, x_{n-1}) \mapsto (x_0 + (x_1)^2, \dots, x_{n-2} + (x_{n-1})^2, x_{n-1}) \, , \tag{17}$$

while the last round is instantiated via power maps $(x_0, x_1, \dots, x_{2n-1}) \mapsto (x_0^d, x_1^d, \dots, x_{2n-1}^d)$ for an integer $d \geq 3$ such that $\gcd(d, p-1) = 1$.

One of the selling points of Pasta compared to Masta regards the fact that no internal collision can occur due to the non-invertibility of the non-linear layer, see [DGH+21, Sect. 5.2]: "*the χ-function is in general* [actually, never] *no permutation when working over $\mathbb{F}_p^t$, which is why we consider some alternatives*". (Remember that a collision at the output can always occur, since Pasta as well as Masta and Rasta are not invertible due to the feed–forward/truncation construction). However, as we already pointed out, it seems hard that an internal collision can be translated into an attack on the entire scheme. Indeed, assume that a collision occurs at the $\tilde{R}$-th round for $\tilde{R} < R$, that is,

$$\mathcal{S}_\chi \circ \mathcal{A}_{\tilde{R},N,i} \circ \ldots \circ \mathcal{A}_{1,N,i} \circ \mathcal{S}_\chi \circ \mathcal{A}_{0,N,i}(x) = \mathcal{S}_\chi \circ \mathcal{A}_{\tilde{R},N',i'} \circ \ldots \circ \mathcal{A}_{1,N',i'} \circ \mathcal{S}_\chi \circ \mathcal{A}_{0,N',i'}(x)$$

for $N \neq N'$ and $i \neq i'$. Since $\mathcal{A}_{j,N,i}$ is (generally) different from $\mathcal{A}_{j,N',i'}$ for each $j \in \{\tilde{R}+1, \dots, R\}$, such collision does not survive the next rounds. Hence, by replacing the Type–III Feistel scheme as in (17) with the SI-lifting function $\mathcal{S}_F$ over $\mathbb{F}_2^n$ induced by $F(x_0, x_1) = x_0^2 + x_1$, we would have the following advantages:

- the depth of the entire construction (and so the overall cost) does not change;
- the obtained construction would be (slightly) more resistant against the linearization attack, since the number of quadratic monomials in $\mathcal{S}_F$ is slightly bigger than in the Type–III Feistel scheme as in (17).

In particular, it is crucial to keep in mind that a collision occurs with probability approximately p^{-n}, which is much smaller than the security level due to the huge size of Pasta (and of Rasta-like design schemes in general). For all these reasons, we claim that the advantages just proposed do not imply a smaller security level.

About `Rubato`. `Rubato` is a family of *noisy* stream ciphers over $\mathbb{F}_p^n$ based on the Rasta design strategy, targeting the transciphering framework for approximate homomorphic encryption. The main difference with Pasta regards the way in which the encryption is performed. First of all, given an input $x \in \mathbb{F}_p^n$ to encrypt and a public nonce $N \in \mathbb{F}_p^n$, the ciphertext is generated as

$$(x, N) \mapsto (x + \mathcal{E}_{\mathtt{K}}(N), N)$$

for a secret key $\mathtt{K} \in \mathbb{F}_p^n$, and a cipher $\mathcal{E}_{\mathtt{K}} : \mathbb{F}_p^n \to \mathbb{F}_p^n$. With respect to the public permutation $\mathcal{P}_{N,i}$ used in Rasta, Masta, and Pasta:

- a round–key addition takes place at each round of $\mathcal{E}_{\mathtt{K}}$;

- the affine layer of $\mathcal{E}_{\mathtt{K}}$ are *fixed*, that is, they do not change at each encryption;
- the round-keys are generated via an affine maps that change at every encryption (as before, such affine maps are generated via a public XOF that takes in input the nonce N).

That is, the l-round sub–key $k_l \in \mathbb{F}_p^n$ for $l \in \{0, 1, \ldots, r\}$ of the i encryption for $i \geq 0$ is defined as $k_l = \mathcal{A}'_{l,N,i}(\mathtt{K})$ for an invertible affine layer $\mathcal{A}'_{l,N,i}$ over $\mathbb{F}_p^n$. We refer to [HKL+22] for more details (in particular, we point out that here we omitted the noise addition for simplicity reason only, since it does not play any role in the following argument).

Having said that, since `Rubato` is instantiated with the same non-linear layers of Pasta, that is, the quadratic Type–III Feistel scheme given in (17), the observations just proposed for Pasta translate directly to `Rubato` as well.

Acknowledgments. Author thanks Morten Øygarden for his valuable comments that helped to improve the quality of the paper. Lorenzo Grassi is supported by the European Research Council under the ERC advanced grant agreement under grant ERC-2017-ADG Nr. 788980 ESCADA.

References

[AAB+20] Abdelrahaman Aly, Tomer Ashur, Eli Ben-Sasson, Siemen Dhooghe, and Alan Szepieniec. Design of Symmetric-Key Primitives for Advanced Cryptographic Protocols. *IACR Trans. Symmetric Cryptol.*, 2020(3):1–45, 2020.

[ACG+19] Martin R. Albrecht, Carlos Cid, Lorenzo Grassi, Dmitry Khovratovich, Reinhard Lüftenegger, Christian Rechberger, and Markus Schofnegger. Algebraic Cryptanalysis of STARK-Friendly Designs: Application to MARVELlous and MiMC. In *ASIACRYPT 2019*, volume 11923 of *LNCS*, pages 371–397, 2019.

[AGP+19] Martin R. Albrecht, Lorenzo Grassi, Léo Perrin, Sebastian Ramacher, Christian Rechberger, Dragos Rotaru, Arnab Roy, and Markus Schofnegger. Feistel Structures for MPC, and More. In *ESORICS 2019*, volume 11736 of *LNCS*, pages 151–171, 2019.

[AGR+16] Martin R. Albrecht, Lorenzo Grassi, Christian Rechberger, Arnab Roy, and Tyge Tiessen. MiMC: Efficient Encryption and Cryptographic Hashing with Minimal Multiplicative Complexity. In *ASIACRYPT 2016*, volume 10031 of *LNCS*, pages 191–219, 2016.

[AMT22] Tomer Ashur, Mohammad Mahzoun, and Dilara Toprakhisar. Chaghri – an FHE-friendly Block Cipher. Cryptology ePrint Archive, Paper 2022/592, 2022. https://eprint.iacr.org/2022/592.

[BBC+22] Clémence Bouvier, Pierre Briaud, Pyrros Chaidos, Léo Perrin, and Vesselin Velichkov. Anemoi: Exploiting the Link between Arithmetization-Orientation and CCZ-Equivalence. Cryptology ePrint Archive, Paper 2022/840, 2022. https://eprint.iacr.org/2022/840.

[BBLP22] Augustin Bariant, Clémence Bouvier, Gaëtan Leurent, and Léo Perrin. Algebraic Attacks against Some Arithmetization-Oriented Primitives. *IACR Transactions on Symmetric Cryptology*, 2022(3):73–101, 2022.

[BBS99] Eli Biham, Alex Biryukov, and Adi Shamir. Cryptanalysis of Skipjack Reduced to 31 Rounds Using Impossible Differentials. In *EUROCRYPT 1999*, volume 1592 of *LNCS*, pages 12–23, 1999.

[BCD+20] Tim Beyne, Anne Canteaut, Itai Dinur, Maria Eichlseder, Gregor Leander, Gaëtan Leurent, María Naya-Plasencia, Léo Perrin, Yu Sasaki, Yosuke Todo, and Friedrich Wiemer. Out of Oddity - New Cryptanalytic Techniques Against Symmetric Primitives Optimized for Integrity Proof Systems. In *CRYPTO 2020*, volume 12172 of *LNCS*, pages 299–328, 2020.

[BDD+17] Ritam Bhaumik, Nilanjan Datta, Avijit Dutta, Nicky Mouha, and Mridul Nandi. The Iterated Random Function Problem. In *Advances in Cryptology - ASIACRYPT 2017*, volume 10625 of *LNCS*, pages 667–697, 2017.

[BDH+17] Guido Bertoni, Joan Daemen, Seth Hoffert, Michaël Peeters, Gilles Van Assche, and Ronny Van Keer. Farfalle: parallel permutation-based cryptography. *IACR Trans. Symmetric Cryptol.*, 2017(4):1–38, 2017.

[BDPA08] Guido Bertoni, Joan Daemen, Michaël Peeters, and Gilles Van Assche. On the Indifferentiability of the Sponge Construction. In *EUROCRYPT 2008*, volume 4965 of *LNCS*, pages 181–197, 2008.

[BDPA13] Guido Bertoni, Joan Daemen, Michaël Peeters, and Gilles Van Assche. Keccak. In *EUROCRYPT 2013*, volume 7881 of *LNCS*, pages 313–314, 2013.

[BS90] Eli Biham and Adi Shamir. Differential Cryptanalysis of DES-like Cryptosystems. In *CRYPTO 1990*, volume 537 of *LNCS*, pages 2–21, 1990.

[BS93] Eli Biham and Adi Shamir. *Differential Cryptanalysis of the Data Encryption Standard*. Springer, 1993.

[Buc76] Bruno Buchberger. A theoretical basis for the reduction of polynomials to canonical forms. *SIGSAM Bull.*, 10(3):19–29, 1976.

[CFG+18] Colin Chaigneau, Thomas Fuhr, Henri Gilbert, Jian Guo, Jérémy Jean, Jean-René Reinhard, and Ling Song. Key-Recovery Attacks on Full Kravatte. *IACR Trans. Symmetric Cryptol.*, 2018(1):5–28, 2018.

[CG20] Tingting Cui and Lorenzo Grassi. Algebraic Key-Recovery Attacks on Reduced-Round Xoofff. In *Selected Areas in Cryptography - SAC 2020*, volume 12804 of *LNCS*, pages 171–197, 2020.

[DEG+18] Christoph Dobraunig, Maria Eichlseder, Lorenzo Grassi, Virginie Lallemand, Gregor Leander, Eik List, Florian Mendel, and Christian Rechberger. Rasta: A Cipher with Low ANDdepth and Few ANDs per Bit. In *CRYPTO 2018*, volume 10991 of *LNCS*, pages 662–692, 2018.

[DGGK21] Christoph Dobraunig, Lorenzo Grassi, Anna Guinet, and Daniël Kuijsters. Ciminion: Symmetric Encryption Based on Toffoli-Gates over Large Finite Fields. In *EUROCRYPT 2021*, volume 12697 of *LNCS*, pages 3–34, 2021.

[DGH+21] Christoph Dobraunig, Lorenzo Grassi, Lukas Helminger, Christian Rechberger, Markus Schofnegger, and Roman Walch. Pasta: A Case for Hybrid Homomorphic Encryption. Cryptology ePrint Archive, Report 2021/731, 2021. https://ia.cr/2021/731.

[DGV91] J. Daemen, R. Govaerts, and J. Vandewalle. Efficient pseudorandom sequence generation by cellular automata. In *Proceedings of the 12th Symposium on Information Theory in the Benelux, Werkgemeenschap voor Informatie- en Communicatietheorie*, 1991.

[DKR97] Joan Daemen, Lars R. Knudsen, and Vincent Rijmen. The block cipher Square. In *FSE 1997*, volume 1267 of *LNCS*, pages 149–165, 1997.

[DMRS20] Christoph Dobraunig, Farokhlagha Moazami, Christian Rechberger, and Hadi Soleimany. Framework for faster key search using related-key higher-order differential properties: applications to Agrasta. *IET Inf. Secur.*, 14(2):202–209, 2020.

[DR01] Joan Daemen and Vincent Rijmen. The Wide Trail Design Strategy. In *IMA International Conference 2001*, volume 2260 of *LNCS*, pages 222–238, 2001.

[DR02a] Joan Daemen and Vincent Rijmen. Security of a Wide Trail Design. In *Progress in Cryptology - INDOCRYPT 2002*, volume 2551 of *LNCS*, pages 1–11, 2002.

[DR02b] Joan Daemen and Vincent Rijmen. *The Design of Rijndael: AES - The Advanced Encryption Standard.* Information Security and Cryptography. Springer, 2002.

[EGL+20] Maria Eichlseder, Lorenzo Grassi, Reinhard Lüftenegger, Morten Øygarden, Christian Rechberger, Markus Schofnegger, and Qingju Wang. An Algebraic Attack on Ciphers with Low-Degree Round Functions: Application to Full MiMC. In *ASIACRYPT 2020*, volume 12491 of *LNCS*, pages 477–506, 2020.

[Gen07] Giulio Genovese. Improving the algorithms of berlekamp and niederreiter for factoring polynomials over finite fields. *J. Symb. Comput.*, 42(1-2):159–177, 2007.

[GHR+22] Lorenzo Grassi, Yonglin Hao, Christian Rechberger, Markus Schofnegger, Roman Walch, and Qingju Wang. A New Feistel Approach Meets *Fluid*-SPN: Griffin for Zero-Knowledge Applications. Cryptology ePrint Archive, Report 2022/403, 2022. https://ia.cr/2022/403.

[GKL+21] Lorenzo Grassi, Dmitry Khovratovich, Reinhard Lüftenegger, Christian Rechberger, Markus Schofnegger, and Roman Walch. Reinforced Concrete: Fast Hash Function for Zero Knowledge Proofs and Verifiable Computation. Cryptology ePrint Archive, Report 2021/1038, 2021. https://ia.cr/2021/1038 – accepted at ACM CCS 2022.

[GKR+21] Lorenzo Grassi, Dmitry Khovratovich, Christian Rechberger, Arnab Roy, and Markus Schofnegger. Poseidon: A New Hash Function for Zero-Knowledge Proof Systems. In *USENIX Security 2021*. USENIX Association, 2021.

[GLR+20] Lorenzo Grassi, Reinhard Lüftenegger, Christian Rechberger, Dragos Rotaru, and Markus Schofnegger. On a Generalization of Substitution-Permutation Networks: The HADES Design Strategy. In *EUROCRYPT 2020*, volume 12106 of *LNCS*, pages 674–704, 2020.

[GOPS22] Lorenzo Grassi, Silvia Onofri, Marco Pedicini, and Luca Sozzi. Invertible Quadratic Non-Linear Layers for MPC-/FHE-/ZK-Friendly Schemes over $\mathbb{F}_p^n$: Application to Poseidon. *IACR Trans. Symmetric Cryptol.*, 2022(3):20–72, 2022.

[GØWS22] Lorenzo Grassi, Morten Øygarden, Roman Walch, and Markus Schofnegger. From Farfalle to Megafono via Ciminion: The PRF Hydra for MPC Applications. Cryptology ePrint Archive, Report 2022/342, 2022. https://ia.cr/2022/342.

[Gra18] Lorenzo Grassi. Mixture Differential Cryptanalysis: a New Approach to Distinguishers and Attacks on round-reduced AES. *IACR Trans. Symmetric Cryptol.*, 2018(2):133–160, 2018.

[GRR+16] Lorenzo Grassi, Christian Rechberger, Dragos Rotaru, Peter Scholl, and Nigel P. Smart. MPC-Friendly Symmetric Key Primitives. In *CCS*, pages 430–443. ACM, 2016.

[GRR17] Lorenzo Grassi, Christian Rechberger, and Sondre Rønjom. A New Structural-Differential Property of 5-Round AES. In *EUROCRYPT 2017*, volume 10211 of *LNCS*, pages 289–317, 2017.

[GRS21] Lorenzo Grassi, Christian Rechberger, and Markus Schofnegger. Proving Resistance Against Infinitely Long Subspace Trails: How to Choose the Linear Layer. *IACR Trans. Symmetric Cryptol.*, 2021(2):314–352, 2021.

[HKC+20] Jincheol Ha, Seongkwang Kim, Wonseok Choi, Jooyoung Lee, Dukjae Moon, Hyojin Yoon, and Jihoon Cho. Masta: An HE-Friendly Cipher Using Modular Arithmetic. *IEEE Access*, 8:194741–194751, 2020.

[HKL+22] Jincheol Ha, Seongkwang Kim, ByeongHak Lee, Jooyoung Lee, and Mincheol Son. Rubato: Noisy Ciphers for Approximate Homomorphic Encryption. In *Advances in Cryptology - EUROCRYPT 2022*, volume 13275 of *LNCS*, pages 581–610, 2022.

[JK97] Thomas Jakobsen and Lars R. Knudsen. The Interpolation Attack on Block Ciphers. In *FSE 1997*, volume 1267 of *LNCS*, pages 28–40, 1997.

[Knu94] Lars R. Knudsen. Truncated and Higher Order Differentials. In *FSE 1994*, volume 1008 of *LNCS*, pages 196–211, 1994.

[Knu98] Lars R. Knudsen. DEAL - A 128-bit Block Cipher. *Technical Report, Department of Informatics, Bergen, Norway*, 1998.

[KR21] Nathan Keller and Asaf Rosemarin. Mind the Middle Layer: The HADES Design Strategy Revisited. In *EUROCRYPT 2021*, volume 12697 of *LNCS*, pages 35–63, 2021.

[KS99] Aviad Kipnis and Adi Shamir. Cryptanalysis of the HFE Public Key Cryptosystem by Relinearization. In *Advances in Cryptology - CRYPTO 1999*, volume 1666 of *LNCS*, pages 19–30, 1999.

[LAAZ11] Gregor Leander, Mohamed Ahmed Abdelraheem, Hoda AlKhzaimi, and Erik Zenner. A Cryptanalysis of PRINTcipher: The Invariant Subspace Attack. In *CRYPTO 2011*, volume 6841 of *LNCS*, pages 206–221, 2011.

[Lai94] X. Lai. Higher order derivatives and differential cryptanalysis. Communications and Cryptography: Two Sides of One Tapestry, 1994.

[LMR15] Gregor Leander, Brice Minaud, and Sondre Rønjom. A Generic Approach to Invariant Subspace Attacks: Cryptanalysis of Robin, iSCREAM and Zorro. In *EUROCRYPT 2015*, volume 9056 of *LNCS*, pages 254–283, 2015.

[LSMI21] Fukang Liu, Santanu Sarkar, Willi Meier, and Takanori Isobe. Algebraic Attacks on Rasta and Dasta Using Low-Degree Equations. In *Advances in Cryptology - ASIACRYPT 2021*, volume 13090 of *LNCS*, pages 214–240, 2021.

[LSMI22] Fukang Liu, Santanu Sarkar, Willi Meier, and Takanori Isobe. The Inverse of χ and Its Applications to Rasta-like Ciphers. Cryptology ePrint Archive, Paper 2022/399, 2022. https://eprint.iacr.org/2022/399.

[Mat93] Mitsuru Matsui. Linear Cryptanalysis Method for DES Cipher. In *Advances in Cryptology - EUROCRYPT 1993*, volume 765 of *LNCS*, pages 386–397, 1993.

[MS15] Brice Minaud and Yannick Seurin. The Iterated Random Permutation Problem with Applications to Cascade Encryption. In *Advances in Cryptology - CRYPTO 2015*, volume 9215 of *LNCS*, pages 351–367, 2015.

[Nyb96] Kaisa Nyberg. Generalized Feistel Networks. In *Advances in Cryptology - ASIACRYPT 1996*, volume 1163 of *LNCS*, pages 91–104. Springer, 1996.

[PvO95] Bart Preneel and Paul C. van Oorschot. MDx-MAC and Building Fast MACs from Hash Functions. In *Advances in Cryptology - CRYPTO 1995*, volume 963 of *LNCS*, pages 1–14, 1995.

[Sze21] Alan Szepieniec. On the Use of the Legendre Symbol in Symmetric Cipher Design. Cryptology ePrint Archive, Report 2021/984, 2021. https://ia.cr/2021/984.

[Wag99] David A. Wagner. The boomerang attack. In *Fast Software Encryption – FSE 1999*, volume 1636 of *LNCS*, pages 156–170, 1999.

[Wol85] Stephen Wolfram. Cryptography with Cellular Automata. In *CRYPTO 1985*, volume 218 of *LNCS*, pages 429–432, 1985.

[ZMI90] Yuliang Zheng, Tsutomu Matsumoto, and Hideki Imai. On the Construction of Block Ciphers Provably Secure and Not Relying on Any Unproved Hypotheses. In *Advances in Cryptology – CRYPTO 1989*, volume 435 of *LNCS*, pages 461–480, 1990.