



An Algebraic Attack on Ciphers with Low-Degree Round Functions: Application to Full MiMC

Maria Eichlseder¹(✉), Lorenzo Grassi^{1,2}(✉), Reinhard Lüftenegger¹(✉), Morten Øyegarden³(✉), Christian Rechberger¹(✉), Markus Schofnegger¹(✉), and Qingju Wang⁴(✉)

¹ IAIK, Graz University of Technology, Graz, Austria

{[maria.eichlseder](mailto:maria.eichlseder@iaik.tugraz.at), [reinhard.luftenegger](mailto:reinhard.luftenegger@iaik.tugraz.at), [christian.rechberger](mailto:christian.rechberger@iaik.tugraz.at), [markus.schofnegger](mailto:markus.schofnegger@iaik.tugraz.at)}@iaik.tugraz.at, lgrassi@science.ru.nl

² Digital Security Group, Radboud University, Nijmegen, The Netherlands

³ Simula UiB, Bergen, Norway

morten.oyegarden@simula.no

⁴ SnT, University of Luxembourg, Luxembourg City, Luxembourg
qingju.wang@uni.lu

Abstract. Algebraically simple PRFs, ciphers, or cryptographic hash functions are becoming increasingly popular, for example due to their attractive properties for MPC and new proof systems (SNARKs, STARKs, among many others).

In this paper, we focus on the algebraically simple construction MiMC, which became an attractive cryptanalytic target due to its simplicity, but also due to its use as a baseline in a competition for more recent algorithms exploring this design space.

For the first time, we are able to describe key-recovery attacks on all full-round versions of MiMC over \mathbb{F}_{2^n} , requiring half the code book. In the chosen-ciphertext scenario, recovering the key from this data for the n -bit full version of MiMC takes the equivalent of less than $2^{n-\log_2(n)+1}$ calls to MiMC and negligible amounts of memory.

The attack procedure is a generalization of higher-order differential cryptanalysis, and it is based on two main ingredients. First, we present a higher-order distinguisher which exploits the fact that the algebraic degree of MiMC grows significantly slower than originally believed. Secondly, we describe an approach to turn this distinguisher into a key-recovery attack without guessing the full subkey. Finally, we show that approximately $\lceil \log_3(2 \cdot R) \rceil$ more rounds (where $R = \lceil n \cdot \log_3(2) \rceil$ is the current number of rounds of MiMC- n/n) can be necessary and sufficient to restore the security against the key-recovery attack presented here.

The attack has been practically verified on toy versions of MiMC. Note that our attack does not affect the security of MiMC over prime fields.

Keywords: Algebraic attack · MiMC · Higher-order differential

1 Introduction

The design of symmetric cryptographic constructions exhibiting a clear and ideally low-degree algebraic structure is motivated by many recent use cases, for example the increasing popularity of new proof systems such as STARKs [8], SNARKs (e.g., Pinocchio [44]), Bulletproofs [19], and other concepts like secure multi-party computation (MPC). To provide good performance in these new applications, ciphers and hash functions are designed in order to minimize specific characteristics (e.g., the total number of multiplications, the depth, or other parameters related to the nonlinear operations). In contrast to traditional cipher design, the size of the field over which these constructions are defined has only a small impact on the final cost. In order to achieve this new performance goal, some crucial differences arise between these new designs and traditional ones. For example, we can consider the substitution (S-box) layer, that is, the operation providing nonlinearity in the permutation: In these new schemes, the S-boxes composing this layer are relatively large compared to the ones used in classical schemes (e.g., they operate over 64 or 128 bits instead of 4 or 8 bits) and/or they can usually be described by a simple low-degree nonlinear function (e.g., $x \mapsto x^d$ for some d). Examples of these schemes include LowMC [4], MiMC [3], JARVIS/FRIDAY [6], GMiMC [2], HadesMiMC [31], *Vision/Rescue* [5], and STARKAD/POSEIDON [30].

The structure of these schemes has a significant impact on the attacks that can be mounted. While statistical attacks (including linear [42] and differential [11] ones) are among the most powerful techniques against traditional schemes, algebraic attacks turned out to be especially effective against these new primitives. In other words, these constructions are naturally more vulnerable to algebraic attacks than those which do not exhibit a clear and simple algebraic structure. For example, this has been shown in [1], in which algebraic strategies covering the full-round versions of the attacked primitives are described. Although the approaches can be quite different, most of them exploit the low degree of the construction.

In this paper, we focus on MiMC [3]. The MiMC design constructs a cryptographic permutation by iterated cubing, interleaved with additions of random constants to break any symmetries. A secret key is added after every such round to obtain a block cipher. The design of MiMC is very flexible and can work with binary strings as well as integers modulo some prime number. Security analysis by the designers rules out various statistical attacks, and the final number of rounds is derived from an analysis of attack vectors that exploit the simple algebraic structure. We remark that the designers chose the number of rounds with a minimal security margin for efficiency. For a more detailed specification and a summary of previous analysis, we refer to Sect. 2.3.

Since its publication in 2016, MiMC has become the preferred choice for many use cases that benefit from a low multiplication count or algebraic simplicity

Table 1. Various attacks on MiMC. In this representation, n denotes the block size (and key size). The unit for the attack complexity is usually the cost of a single encryption (number of multiplications over \mathbb{F}_{2^n} necessary for a single encryption). The SK and KR attacks can be implemented using chosen plaintexts CP and/or chosen ciphertexts CC. The memory complexity is negligible for all approaches listed.

Type	n	Rounds	Time	Data	Source
KR*	129	38	$2^{65.5}$	$2^{60.2}$ CP	[41]
SK	129	80	2^{128} XOR	2^{128} CP/CC	Sect. 4.1
SK	n	$\lceil \log_3(2^{n-1} - 1) \rceil - 1$	2^{n-1} XOR	2^{n-1} CP/CC	Sect. 4.1
KK	129	160 ($\approx 2 \times$ full)	–	2^{128}	Sect. 4.3
KK	n	$2 \cdot \lceil \log_3(2^{n-1} - 1) \rceil - 2$	–	2^{n-1}	Sect. 4.3
KR	129	82 (full)	$2^{122.64}$	2^{128} CC	Sect. 5
KR	255	161 (full)	$2^{246.67}$	2^{254} CC	Sect. 5
KR	n	$\lceil n \cdot \log_3(2) \rceil$ (full)	$\leq 2^{n - \log_2(n) + 1}$	2^{n-1} CC	Sect. 5

KR \equiv Key-Recovery, KR* \equiv attack on a variant of MiMC proposed in a low-memory scenario, SK \equiv Secret-Key Distinguisher, KK \equiv Known-Key Distinguisher

[32, 45]. It also serves as a baseline for various follow-up designs evaluated in the context of the public “STARK-Friendly Hash Challenge” competition¹.

1.1 Our Contribution

As the main results in this paper, we present

- (1) a new upper bound for the algebraic degree growth in key-alternating ciphers with low-degree round functions,
- (2) a secret-key higher-order distinguisher on almost full MiMC over \mathbb{F}_{2^n} ,
- (3) a known-key zero-sum distinguisher on almost double the rounds of MiMC,
- (4) the first key-recovery attack on *full-round* MiMC over \mathbb{F}_{2^n} .

We also show that the technique we use for MiMC is sufficiently generic to apply to any permutation fulfilling specific properties, which we will define in detail. Our attacks and distinguishers on MiMC, as well as other attacks in the literature, are listed in Table 1.

Secret-Key Higher-Order Distinguishers. After recalling some preliminary facts about higher-order differentials, in Sect. 3 we analyze the growth of the algebraic degree for key-alternating ciphers whose round function can be described as a low-degree polynomial over \mathbb{F}_{2^n} .

For an SPN cipher over a field \mathbb{F} where each round has algebraic degree δ , the algebraic degree of the cipher is expected to grow essentially exponentially in δ . Several analyses made in the literature [17, 18, 20] confirm this growth for most

¹ <https://starkware.co/hash-challenge/>.

ciphers, except when the algebraic degree of the function is close to its maximum. As a result, the number of rounds necessary for security against higher-order differential attacks generally grows logarithmically in the size of \mathbb{F} . Different behaviour has been observed for certain non-SPN designs, such as some designs with partial nonlinear layers where the algebraic degree grows exponentially in some (not necessarily integer) value smaller than δ [26].

In Sect. 3, we show that if the round function can be described as an invertible low-degree polynomial function in \mathbb{F}_{2^n} , then the algebraic degree grows linearly with the number of rounds, and not exponentially as generally expected. More precisely, let d denote the exponent of the power function $x \mapsto x^d$ used to define the S-boxes. Then, we show that in the case of key-alternating ciphers over \mathbb{F}_{2^n} , the algebraic degree $\delta(r)$ as a function in the number of rounds r is

$$\delta(r) \in \mathcal{O}(\log_2(d^r)) = \mathcal{O}(r).$$

As an immediate consequence, our observation implies that roughly $n \cdot \log_d(2)$ rounds are necessary to provide security against higher-order differential attacks, much more than the expected $\approx \log_\delta(n - 1)$ rounds.

Distinguishers on MiMC over \mathbb{F}_{2^n} . Our new bounds on the number of rounds necessary to provide security against higher-order differential cryptanalysis have a major impact on key-alternating ciphers with large S-boxes. A concrete example for this class of ciphers is MiMC [3], a key-alternating cipher defined over \mathbb{F}_{2^n} (for odd $n \in \mathbb{N}$), where the round function is simply defined as the cube map $x \mapsto x^3$. Since any cubic function over \mathbb{F}_{2^n} has algebraic degree 2, one may expect that approximately $\log_2(n)$ rounds are necessary to prevent higher-order differential attacks. Our new bound implies that a much larger number of rounds is required to provide security, namely approximately $n \cdot \log_3(2)$.

As a concrete example, in Sect. 4 we show that MiMC- n/n has a security margin of only 1 or 2 rounds against (secret-key) higher-order distinguishers (depending on n), which is much smaller than expected by the designers. Moreover, we can set up a known-key distinguisher for approximately double the number of rounds of MiMC, by showing that the same number of rounds is necessary to reach the maximum degree in the decryption direction. Our findings have been practically verified on toy versions.

We remark that the designers presented other non-random properties (including GCD and interpolation attacks) that can cover a similar number of rounds. The number of rounds proposed by the designers were chosen in order to provide security against key-recovery attacks based on these properties. As we are going to show, the number of rounds is not sufficient against our new attack based on a higher-order differential property.

Results using the Division Property. For completeness, in Sect. 4.5 we search for higher-order distinguishers for MiMC- n/n with the division property [46] proposed by Todo at Eurocrypt 2015, a powerful tool for finding the best integral distinguishers for block ciphers. By modeling the most recently proposed variant of the bit-based division property, which is called *three-subset bit-based division*

property without unknown subset in [34], we are able to reproduce exactly the same higher-order distinguishers for cases with small n -bit S-boxes, where $n \in \{5, 7, 9\}$. However, as far as we know, it is an open problem to model the three-subset bit-based division property for a larger S-box of size bigger than 9 in practical time. Therefore, we conclude that the division property is unlikely to help us for the ciphers we focus on.

Key-Recovery Attack on MiMC- n/n and on Generic Ciphers. A trivial way to extend an r -round distinguisher to an $(r + 1)$ -round key-recovery attack is based on guessing the last round key, partially decrypting/encrypting, and finally exploiting the distinguisher to filter wrong key guesses. Unfortunately, this strategy does not work for MiMC, since guessing the full last round key required to invert the large S-box is equivalent to exhaustive key search. Another key-recovery approach that has been combined with integral distinguishers is based on interpolating the Boolean polynomials that define the final rounds. However, this strategy requires evaluating the distinguisher several times to collect enough equations, which is not feasible for our distinguisher due to its large data complexity.

In Sect. 5, we show how to solve this problem. Instead of guessing the last round key, we set up an equation over \mathbb{F}_{2^n} with the master key as a variable. To obtain this equation, we symbolically express the zero sum at the input to the last round as a polynomial function of the key, whose coefficients depend on the queried ciphertexts. We show how the resulting polynomial equation can be solved efficiently to recover the key. As a result, in the chosen-ciphertext case only, recovering the key from this data for the *full* n -bit version of MiMC takes the equivalent of less than $2^{n-\log_2(n)+1}$ calls to MiMC, 2^{n-1} chosen ciphertexts, and negligible amounts of memory. Moreover, we show that approximately $\lceil \log_3(2 \cdot R) \rceil$ more rounds (where $R = \lceil n \cdot \log_3(2) \rceil$ is the current number of rounds of MiMC- n/n) can be necessary and sufficient to restore the security against the key-recovery attack presented here. This would, for example, imply that we need to add 5 more rounds for the most used version MiMC-129/129 (which currently has 82 rounds).

A Generic Strategy. Our strategy is an instance of a broader class of algebraic key-recovery approaches based on solving equations in the key variables. As such, it shares some ideas with other algebraic approaches like optimized interpolation attacks. However, while most algebraic key-recovery approaches of the last years construct and solve systems of many Boolean linear equations, we use a single univariate equation of higher degree that can be solved with polynomial factoring algorithms such as Berlekamp's algorithm. In Sect. 6, we outline a more detailed and generic procedure for such an attack. It is interesting to note that a comparatively old technique which basically disappeared for the cryptanalysis of AES-like ciphers turns out to be very competitive for schemes with large S-boxes.

2 Preliminaries

In this section, we recall the most important results about polynomial representations of Boolean functions and summarize the currently best known results

regarding the growth of the algebraic degree in the context of SP networks. We also provide the specification of MiMC and give an overview of previous cryptanalytic results.

We emphasize that in general it is only possible to give a *lower* bound regarding the number of rounds which we can attack using higher-order differential techniques, in the following denoted as “necessary number of rounds to provide security”. While upper-bounding the algebraic degree is more important from an adversary’s point of view, lower bounds on the degree are much more relevant when arguing about security against algebraic attacks (such as e.g. [24, 38, 40, 49]) from a designer’s viewpoint. However, at the current state of the art and to the best of our knowledge, it seems hard to find such a lower bound for a given cipher without investigating concrete instances experimentally – which, of course, limits the scope of any analysis.

2.1 Polynomial Representations over Binary Extension Fields

We denote addition (and subtraction) in binary extension fields by the symbol \oplus . For $n \in \mathbb{N}$, every function $F : \mathbb{F}_{2^n} \rightarrow \mathbb{F}_{2^n}$ can be uniquely represented by an n -tuple (F_1, F_2, \dots, F_n) of polynomials over \mathbb{F}_2 in n variables with a maximum degree of 1 in each variable. In this representation, F_i is of the form

$$F_i(X_1, \dots, X_n) = \bigoplus_{u=(u_1, \dots, u_n) \in \{0,1\}^n} \varphi_i(u) \cdot X_1^{u_1} \cdots X_n^{u_n}, \tag{1}$$

where the coefficients $\varphi_i(u)$ can be computed by the *Moebius transform*.

As is common, we denote functions $F : \mathbb{F}_2^n \rightarrow \mathbb{F}_2$ as *Boolean functions* and functions of the form $F : \mathbb{F}_2^n \rightarrow \mathbb{F}_2^m$, for $n, m \in \mathbb{N}$, as *vectorial Boolean functions*.

Definition 1. *The algebraic normal form (ANF) of a Boolean function $F : \mathbb{F}_2^n \rightarrow \mathbb{F}_2$, as given in Eq. (1), is the unique representation as a polynomial over \mathbb{F}_2 in n variables and with a maximum univariate degree of 1. The algebraic degree $\delta(F)$ of F – or δ for simplicity – is the degree of the above representation of F as a multivariate polynomial over \mathbb{F}_2 . If $G : \mathbb{F}_2^n \rightarrow \mathbb{F}_2^m$ is a vectorial Boolean function and (G_1, \dots, G_m) is its representation as an n -tuple of multivariate polynomials over \mathbb{F}_2 , then its algebraic degree $\delta(G)$ is defined as $\delta(G) := \max_{1 \leq i \leq m} \delta(G_i)$.*

The link between the algebraic degree and the univariate degree of a vectorial Boolean function is well-known, and is for example established in [22]: the algebraic degree of $F : \mathbb{F}_{2^n} \rightarrow \mathbb{F}_{2^n}$ can be computed from its univariate polynomial representation, and is equal to the maximum hamming weight of the 2-ary expansion of its exponents.

Lemma 1. *Let $F : \mathbb{F}_{2^n} \rightarrow \mathbb{F}_{2^n}$ be a function and let $F(X) = \sum_{i=0}^{2^n-1} \varphi_i \cdot X^i$ denote the corresponding univariate polynomial description over \mathbb{F}_{2^n} . The algebraic degree $\delta(F)$ of F as a vectorial Boolean function is the maximum hamming weight² of its exponents, i.e., it is $\delta(F) = \max_{0 \leq i \leq 2^n-1} \{\text{hw}(i) \mid \varphi_i \neq 0\}$.*

² Given $x = \sum_{i=0}^x x_i \cdot 2^i$ for $x_i \in \{0, 1\}$, the hamming weight of x is $\text{hw}(x) = \sum_{i=0}^x x_i$.

2.2 Higher-Order Differential Cryptanalysis

Higher-order differential attacks [38,40] form a prominent class of attacks exploiting the low algebraic degree of a nonlinear transformation such as a classical block cipher. If this degree is sufficiently low, an attack using multiple input texts and their corresponding output texts can be mounted. In more detail, if the algebraic degree of a Boolean function f is δ , then, when applying f to all elements of an affine vector space $\mathcal{V} \oplus c$ of dimension greater than δ and taking the sum of these values, the result is 0, i.e., $\bigoplus_{v \in \mathcal{V} \oplus c} f(v) = 0$.

Security Against Higher-Order Differential Attacks – State of the Art.

To prevent higher-order differential attacks against iterated block ciphers, one would usually want the maximum algebraic degree to be reached (well) within the suggested number of rounds. To achieve this goal, and to assess the security margins, it is crucial to estimate how the algebraic degree grows with the number of rounds.

The algebraic degree of composing two functions, $F, G : \mathbb{F}_2^n \rightarrow \mathbb{F}_2^n$, can be generically bounded by

$$\deg(F \circ G) \leq \deg(F) \cdot \deg(G), \tag{2}$$

and hence an upper bound is found by iterative use of this on the round function. The resulting bound does, however, fail to reflect the real growth of the algebraic degree for many cryptosystems, and the problem of estimating the growth has been widely studied in the literature. After the initial work of Canteaut and Videau [20], a tighter upper bound was presented by Boura, Canteaut, and De Cannière [18] at FSE'11. There, the authors show how to deduce a new bound for the algebraic degree of iterated permutations for a special category of SP networks over $(\mathbb{F}_2^n)^t$, which includes functions that have a number $t \geq 1$ of balanced S-boxes as their nonlinear layer. Specifically, the authors show that the algebraic degree of the considered SP network grows almost exponentially, except when it is close to its maximum.

Proposition 1 ([18]). *Let F be a function from \mathbb{F}_2^N to \mathbb{F}_2^N corresponding to the concatenation of t smaller S-boxes S_1, \dots, S_t defined over \mathbb{F}_2^n . Then, for any function G from \mathbb{F}_2^N to \mathbb{F}_2^N , we have*

$$\deg(G \circ F(\cdot)) \leq \min \left\{ \deg(F) \cdot \deg(G), N - \frac{N - \deg(G)}{\gamma} \right\}, \text{ where} \tag{3}$$

$$\gamma = \max_{i=1, \dots, n-1} \frac{n-i}{n-\delta_i} \leq n-1, \tag{4}$$

and where δ_i is the maximum degree of the product of any i coordinates of any of the smaller S-boxes.

Thus, the number of rounds necessary to prevent higher-order differential attacks is in general bigger than the one obtained using the trivial bound in Eq. (2).




Fig. 1. The MiMC encryption function with r rounds.

2.3 Specification and Previous Analysis of MiMC

MiMC [3] is a key-alternating n -bit block cipher, where in each round the same n -bit key is added to the state. The nonlinear component of the construction is the evaluation of the cube function $f(x) = x^3$ over \mathbb{F}_{2^n} . Additionally, a different round constant is added in each round to break symmetries, where the first round constant is 0. The total number of rounds is then

$$r = \lceil n \cdot \log_3(2) \rceil,$$

and we refer to Fig. 1 for a graphical representation of the encryption function.

MiMC is defined to work over prime fields and binary fields. In this paper, we focus on the binary field versions of MiMC³, for which the block size n has to be odd in order for the S-box to be a permutation.

MiMC: Related Attacks in the Literature. The designers recommend MiMC with $\lceil n \cdot \log_3(2) \rceil$ rounds [3]. They derive this number of rounds by considering a variety of different key-recovery attacks on MiMC. According to their analysis, the most powerful attacks are interpolation [36] and GCD attacks. About higher-order differential attacks, the authors claim that “the large number of rounds ensures that the algebraic degree of MiMC in its native field will be maximum or almost maximum. This naturally thwarts higher-order differential attacks [...]”.

The first attack on MiMC- n/n [41], presented at SAC 2019, targets a reduced-round version of MiMC proposed by the designers for a scenario in which the attacker has only limited memory, but it does not affect the security claims of full-round MiMC. The Feistel version of MiMC was attacked shortly after, by using generic properties of the used Feistel construction instead of exploiting properties of the primitive itself [16]. Finally, a specific attack on MiMC using Gröbner bases was considered in [1]. The authors state that by introducing a new intermediate variable in each round, the resulting multivariate system of equations is already a Gröbner basis and thus the first step of a Gröbner basis attack is for free. However, recovering univariate polynomials from this representation and then applying techniques like the GCD attack will result in a prohibitively large computational complexity, since the recovered polynomials will be of degree $\approx 3^r$ after r rounds. Hence, the authors conclude that MiMC cannot be attacked directly by using known Gröbner basis techniques.

³ Since the only subspaces of \mathbb{F}_p , where p is a prime number, are $\{0\}$ and \mathbb{F}_p itself, our attack does not affect the security of MiMC over prime fields.

3 Higher-Order Differentials of Key-Alternating Ciphers

Our bound on the growth of the algebraic degree does not depend on the cubing of the round function in MiMC, so we introduce the following generalization of the result on MiMC from Sect. 2.3.

3.1 Setting

Let $E_k^r : \mathbb{F}_{2^n} \rightarrow \mathbb{F}_{2^n}$ be a key-alternating cipher defined by

$$E_k^r(x) := k_r \oplus R(\dots R(k_1 \oplus R(k_0 \oplus x)) \dots) \tag{5}$$

over $r \geq 1$ rounds, where $k_0, k_1, \dots, k_r \in \mathbb{F}_{2^n}$ are derived from a master key $k \in \mathbb{F}_{2^n}$ using a key schedule. Each round function $R : \mathbb{F}_{2^n} \rightarrow \mathbb{F}_{2^n}$ is defined as some invertible univariate polynomial function

$$R(x) := \rho_0 \oplus \bigoplus_{i=1}^d \rho_i \cdot x^i \tag{6}$$

of univariate degree $d \geq 3$, where $\rho_i \in \mathbb{F}_{2^n}$ and $\rho_d \neq 0$. We will, without loss of generality, assume $d \leq d_{\text{inv}}$, where d_{inv} denotes the degree of the compositional inverse of R (otherwise, an attacker would target the decryption function instead). Furthermore, we assume that the round function has *low* univariate degree, i.e., low compared to the size of \mathbb{F}_{2^n} . In other words, we work with $d \ll 2^n - 1$.

3.2 Growth of the Degree

In this section, we show that the algebraic degree δ of a key-alternating cipher E_k^r grows much slower than commonly presented in the literature. More precisely, in some cases it can grow linearly in the number of rounds and not exponentially.

Proposition 2. *Let E_k^r be a an r -round key-alternating block cipher with a round function R of degree d , as defined in Eq. (5). If $r \leq \mathcal{R}_{\text{lin}} - 1$, where*

$$\mathcal{R}_{\text{lin}} = \lceil \log_d(2^{n-1} - 1) \rceil \approx (n - 1) \cdot \log_d(2), \tag{7}$$

then the algebraic degree δ of E_k^r is at most $n - 2$. Then, a (secret-key) higher-order distinguisher using at most 2^{n-1} data can be applied to E_k^r .⁴

Proof. Due to the relation between the word-level degree and the algebraic degree, E_k^r reaches its maximum algebraic degree of $n - 1$ if at least one monomial with the exponent $2^n - 2^j - 1$ (for $0 \leq j < n$) appears in the polynomial representation. Indeed, note that all these monomials have an algebraic degree of $n - 1$. Since the smallest exponent of this form is $2^n - 2^{n-1} - 1 = 2^{n-1} - 1$,

⁴ We denote our bound by \mathcal{R}_{lin} to indicate the almost linear growth of the algebraic degree for this specific class of constructions.

and since the degree of E_k^r after r rounds is at most d^r , we require $d^r \geq 2^{n-1} - 1$ to make $x^{2^{n-1}-1}$ appear, or equivalently,

$$r \geq \lceil \log_d(2^{n-1} - 1) \rceil.$$

Hence, the degree is not maximal for $r < \lceil \log_d(2^{n-1} - 1) \rceil$ and a higher-order distinguisher using at most 2^{n-1} data can be applied. \square

The Difficulty of Lower-Bounding the Growth of the Degree. We point out that it is always possible to set up a (secret-key) higher-order distinguisher if the number of rounds is smaller than \mathcal{R}_{lin} . However, a number of rounds greater than or equal to \mathcal{R}_{lin} does not necessarily provide security.

One of the main problems in order to derive a sufficient condition for the number of rounds that provides security is the difficulty of analyzing the non-vanishing coefficients in the polynomial representation of E_k^r . Note, in general it is not easy to give a condition guaranteeing that a particular monomial appears, since many factors (including the secret key, the constant addition, and the details of the S-box) influence the result.

Without going into the details, we consider the influence of the S-box in some concrete examples. Working with $R(x) = x^d$ for a certain $3 \leq d \leq 2^n - 2$ (where $d \neq 2^{d'}$ for $d' \in \mathbb{N}$), we focus for simplicity only on two extreme cases $d = 2^{d'} \pm 1$. By exploiting Lucas's Theorem⁵:

- If $d = 2^{d'} + 1$ for some $d' \in \mathbb{N}$, then the output of a single round is sparse:

$$(x \oplus y)^{2^{d'}+1} = x^{2^{d'}+1} \oplus x^{2^{d'}} \cdot y \oplus y^{2^{d'}} \cdot x \oplus y^{2^{d'}+1}$$

(note that it contains only 4 terms instead of $d + 1 = 2^{d'} + 2$).

- If $d = 2^{d'} - 1$ for some $d' \in \mathbb{N}$, then the output of a single round is full, since

$$(x \oplus y)^{2^{d'}-1} = \bigoplus_{i=0}^{2^{d'}-1} x^i \cdot y^{2^{d'}-1-i}.$$

Even if a single round is not sparse, the output of several combined rounds is not guaranteed to be full (even if it is in general dense). As a concrete example, while the output of $(x \oplus k_0)^3 \oplus k_1$ is full, the same is not true for

$$\begin{aligned} ((x \oplus k_0)^3 \oplus k_1)^3 \oplus k_2 &= x^9 \oplus x^8 \cdot k_0 \oplus x^6 \cdot k_1 \oplus x^4 \cdot k_0^2 \cdot k_1 \oplus x^3 \cdot k_1^2 \\ &\oplus x^2 \cdot (k_0 \cdot k_1^2 \oplus k_0^2 \cdot k_1^2 \oplus k_0^4 \cdot k_1) \oplus x \cdot k_0^8 \oplus c(k_0, k_1, k_2), \end{aligned} \tag{8}$$

where both x^5 and x^7 are missing, and where $c(k_0, k_1, k_2)$ is a function that depends only on the keys. This simple example emphasizes the difficulty of analyzing the sparsity of the polynomial that defines E_k .

⁵ By Lucas's Theorem, $\binom{n}{m} \equiv \prod_{i=0}^k \binom{n_i}{m_i} \pmod{2}$, it follows that where $n = \sum_{i=0}^k n_i \cdot 2^i$ and $m = \sum_{i=0}^k m_i \cdot 2^i$ is the 2-ary expansion of n and m , respectively.

3.3 Comparison with Other Bounds

We now compare the new number of rounds necessary to provide security against secret-key higher-order distinguishers with other possible bounds. An alternative strategy is to apply generic bounds focusing on the algebraic degree of the round function, as recalled in Proposition 1. Recall that \mathcal{R}_{lin} is the number of rounds from Proposition 2, and we will denote the number of round based on generic bounds by \mathcal{R}_{gen} . The comparison will make use of $\delta_{\text{lin}}(r)$, the upper bound on the algebraic degree after r rounds following Proposition 2. The upper bound from Eq. (3) will be denoted by $\delta_{\text{gen}}(r)$. Note that $\delta_{\text{gen}}(r)$ can, for example, take advantage of a slower growth in the algebraic degree, as in Eq. (8) by considering two rounds instead of one. Despite this, the overall trend of $\delta_{\text{gen}}(r)$ will still be exponential. On the other hand, if the round function can be described by a polynomial of low univariate degree d over \mathbb{F}_{2^n} , we expect a linear behaviour in $\delta_{\text{lin}}(r)$:

$$\delta_{\text{lin}}(r) \leq \lfloor \log_2(d^r + 1) \rfloor \approx r \cdot \log_2(d).$$

As a result, the round numbers \mathcal{R}_{lin} and \mathcal{R}_{gen} necessary to provide security grow respectively linearly and logarithmically in the size n of the field, namely

$$\mathcal{R}_{\text{lin}} \in \mathcal{O}(n) \quad \text{and} \quad \mathcal{R}_{\text{gen}} \in \mathcal{O}(\log_\delta(n)).$$

A concrete comparison of $\delta_{\text{lin}}(r)$ and $\delta_{\text{gen}}(r)$ for MiMC-129/129 is given in Fig. 2. In this setting we have $\delta_{\text{lin}}(r) = \lfloor \log_2(3^r + 1) \rfloor$, and $\delta_{\text{gen}}(r)$ has been derived using the observation that two rounds of MiMC have algebraic degree two (see [28, App. A] for more details). In particular, we find $\mathcal{R}_{\text{gen}} = 13$ and $\mathcal{R}_{\text{lin}} = 81$.

Remark. We emphasize that every (invertible) S-box/round function in \mathbb{F}_2^n can be rewritten as a polynomial over \mathbb{F}_{2^n} . The crucial point here is that given a “random” S-box/round function over \mathbb{F}_2^n , the corresponding polynomial over \mathbb{F}_{2^n} has in general a high univariate degree (e.g., $d \approx 2^n - \varepsilon$ for some small ε). In such a case, even if our argument still holds, the final result becomes meaningless, since $\log_d(2^n - 1) \approx \log_{2^n - \varepsilon}(2^n - 1) \approx 1$ is basically constant (i.e., it does not grow linearly with n). Hence, our results turn out to be relevant only for S-boxes/round functions for which the corresponding polynomial over \mathbb{F}_{2^n} has “small” degree (namely, small compared to the field size, i.e., $d \ll 2^n$).

4 Distinguishers for Reduced-Round and Full MiMC

Exploiting the previous result, we now discuss the possibility to set up higher-order differential distinguishers and attacks on MiMC [3]. We show that

- (1) MiMC has a security margin of only 1 or 2 round(s) against (secret-key) higher-order distinguishers, depending on n , and that
- (2) a zero-sum known-key distinguisher can be set up for approximately double the number of rounds of MiMC.

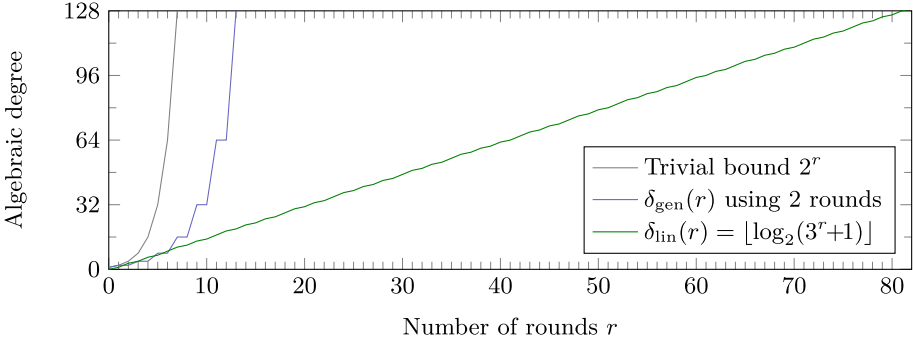


Fig. 2. Different upper bounds of the growth of the algebraic degree for MiMC-129/129. The trivial bound is 2^r . A tighter bound, $\delta_{\text{gen}}(r)$, exploits the observation that 2 rounds only have degree 2 (see Eq. (8)). Our new bound, $\delta_{\text{lin}}(r)$, is linear in the number of rounds.

4.1 Secret-Key Higher-Order Distinguisher for MiMC

The results just presented allow to set up a nontrivial (secret-key) higher-order distinguisher on $\lceil \log_3(2^{n-1} - 1) \rceil - 1$ rounds of MiMC, where $\lceil \log_3(2^{n-1} - 1) \rceil - 1 < \lceil n \cdot \log_3(2) \rceil$ for all n . Consequently, the security margin is reduced to

$$1 \leq \lceil n \cdot \log_3(2) \rceil - (\lceil \log_3(2^{n-1} - 1) \rceil - 1) \leq 2$$

rounds. To give some concrete examples, MiMC has 1 round of security margin for $n \in \{33, 63, 255\}$, and 2 rounds of security margin for $n \in \{31, 65, 127, 129\}$.

4.2 Practical Results

In this section we compare the results from Proposition 2 with practical results from scaled-down versions of MiMC. The tests⁶ have been performed in the following way: Instead of computing the ANF of a keyed permutation (which is expensive even for small field sizes), we evaluate the higher-order differential zero-sum property (as given in Sect. 2.2) for a specific input vector space. Namely, for random keys, random constants, and an input subspace of dimension $n - 1$, we look for the minimum number of rounds r for which the corresponding sum of the ciphertexts is different from zero. Such a number corresponds to the number of rounds necessary to prevent higher-order distinguishers. In order to avoid the influence of weak keys or round constants, we repeated the tests multiple times (with new random keys and round constants). The practical number of rounds we give in each row is *the smallest number of rounds among all tested keys and round constants necessary to prevent higher-order distinguishers*. This means that a potentially higher number of rounds can be attacked by choosing the keys and round constants in a particular way.

⁶ The source code for the attacks and the tests is available on <https://github.com/IAIK/mimc-analysis>.

Table 2. Theoretical and practical round numbers *necessary* to prevent higher-order distinguishers for MiMC over \mathbb{F}_{2^n} .

Param.	Theoretical		Practical
	\mathcal{R}_{lin}	\mathcal{R}_{gen}	\mathcal{R}
7	4	5	5
9	6	5	6
11	7	7	7
13	8	7	9
15	9	7	10
17	11	7	11
33	21	9	21
65	41	11	-
129	81	13	-

The results, denoted \mathcal{R} , are given in Table 2. We also present \mathcal{R}_{lin} (from Proposition 2) and \mathcal{R}_{gen} (see [28, App. A]) for comparison. We emphasize that the theoretical values predicted by \mathcal{R}_{lin} match the practical results in about half of the cases, and are off by at most one.

4.3 Known-Key Zero-Sum Distinguisher for MiMC

A known-key distinguisher is a scenario introduced in [39] where the attacker knows the key, and it is important in all settings in which no secret material is present. To succeed, the attacker has to discover some property of the attacked cipher that holds with a probability higher than for an ideal cipher, or is believed to be hard to exhibit generically. The goal of a known-key zero-sum distinguisher is to find a set of plaintexts and ciphertexts whose sums are equal to zero. To do this, the idea is to exploit the inside-out approach. By choosing a subspace of texts \mathcal{V} , one simply defines the plaintexts as the r_{dec} -round decryption of \mathcal{V} and the ciphertexts as the r_{enc} -round encryption of \mathcal{V} . Such a distinguisher can then cover $r_{\text{enc}} + r_{\text{dec}}$ rounds. Examples of this approach are given in the literature for KECCAK [7, 10, 18], Luffa [7, 18], or PHOTON [50].

In the case of MiMC, the idea is to choose \mathcal{V} as a subspace of \mathbb{F}_{2^n} of dimension $n - 1$. The maximum number of encryption rounds r_{enc} for which it is possible to guarantee a zero sum has been given in the previous paragraph. Based on Sect. 4.2, we can set up a known-key distinguisher on (more than) full MiMC- n/n . For our distinguisher on MiMC, we first recall the following result from [17].

Proposition 3 (Corollary 3 of [17]). *Let F be a permutation of \mathbb{F}_2^n . Then, $\deg(F^{-1}) = n - 1$ if and only if $\deg(F) = n - 1$.*

Corollary 1. *Let r_{enc} be the number of rounds necessary for MiMC over \mathbb{F}_{2^n} to reach its maximum algebraic degree in the encryption direction. The same*

number of rounds is necessary for reaching the maximum algebraic degree in the decryption direction, i.e., $r_{dec} = r_{enc} = \lceil \log_3(2^{n-1} - 1) \rceil$.

It follows that, given a subspace $\mathcal{V} \subseteq \mathbb{F}_{2^n}$ of dimension $n - 1$, the sums of the corresponding texts after $r_{dec} - 1$ decryption rounds and $r_{enc} - 1$ encryption rounds are always equal to zero, i.e.,

$$\underbrace{\bigoplus_{w \in \mathcal{V} \oplus v} R^{-(r_{dec}-1)}(w)}_{\text{Zero sum}} = 0 \xleftarrow{R^{-(r_{dec}-1)}} \mathcal{V} \oplus v \xrightarrow{R^{r_{enc}-1}} 0 = \underbrace{\bigoplus_{w \in \mathcal{V} \oplus v} R^{r_{enc}-1}(w)}_{\text{Zero sum}}$$

for each $v \in \mathbb{F}_{2^n}$. Hence, a known-key zero-sum distinguisher can be set up for

$$\begin{aligned} 2 \cdot (\lceil \log_3(2^{n-1} - 1) \rceil - 1) &\approx 2(n - 1) \cdot \log_3(2) - 2 = \\ &= \underbrace{n \cdot \log_3(2)}_{= \text{full MiMC}} + [(n - 2) \cdot \log_3(2) - 2] \end{aligned}$$

rounds of MiMC- n/n , which is much more than *full* MiMC- n/n .

4.4 Impact of the Known-Key Distinguisher on Full MiMC

Sponge Function. In [3], the authors propose a hash function by instantiating a sponge construction with MiMC^π , a fixed-key version of MiMC. The sponge hash function is indifferentiable from a random oracle up to $2^{c/2}$ calls to the internal permutation P (where c is the capacity) if P is modeled as a randomly chosen permutation [9]. Thus, even if it is not strictly necessary, it is desirable that MiMC is resistant against known-key distinguishers.

For completeness, we mention that even if there is a way to distinguish a permutation from a random one, it seems difficult to exploit a zero-sum distinguisher of the internal permutation of a sponge construction in order to attack the hash function. To give a concrete example, consider the case of KECCAK: As a consequence of the zero-sum distinguisher found on 18-round KECCAK- $f[1600]$, the number of rounds has been increased from 18 to 24 in the second round of the SHA-3 competition in order to avoid “non-ideal” properties (see [10, 18] for more details). However, the best known attack on the KECCAK hash function can only be set up when using 6-/7-round KECCAK- f [33].

In any case, we remark that such distinguishers based on zero sums cannot be set up for an arbitrary number of rounds, and they do indeed exploit the internal properties of a primitive using the inside-out approach found in this paper and in other literature. Hence, they cannot be considered meaningless.

Other Approaches. Even though the original MiMC paper only specifies a sponge-based hash function using MiMC, there are various applications and/or specific considerations that would make a block-cipher-based approach more advantageous (like, for example, being forced to use a block size which is too

small for a sponge-based approach). Another way to turn a block cipher into a hash function is to use a compression function like the Davies–Meyer one together with something like the Merkle–Damgård construction. Similar to the case of sponge constructions, the security of such an algorithm is proven in the ideal cipher model [12]. This choice is, however, not supported by the MiMC designers, who use our results to support their advice against using a block-cipher-based approach (even though such implementations can still be found⁷). It follows that, since the attacker has control of the key in such scenarios, it is desirable for MiMC to be resistant against known- and chosen-key distinguishers, even if it does not seem to be strictly necessary.

4.5 Results Using the Division Property

Finally, in [28, App. C] we present our practical results obtained using “Mixed Integer Linear Programming (MILP)”, which models the propagation of the (conventional) bit-based division property.

The (conventional) bit-based division property [48] was proposed to investigate integral characteristics of block ciphers at a bit level. With this approach, the integral property of each bit is studied independently. Naturally, this strategy allows to capture more information of the propagation than the word-level version, and thus integral characteristics for more rounds can be found with this new technique. For example, the integral distinguishers of SIMON32 have been improved from 10 rounds [46] (the current best result at word level) to 14 rounds [52] (obtained by the experimental method cited before).

Instead of separating the parity into the two cases “0” and “unknown” as for the (conventional) bit-based division property, three-subset bit-based division property [48] was introduced to enhance the accuracy of the conventional one, where the parity is separated into three sets, i.e., “0”, “1”, and “unknown”. It shows that the three-subset bit-based division property can indeed be more accurate than the two-subset bit-based division property for some ciphers [35, 53]. However, it becomes harder to efficiently model the three-subset division property propagation even for ciphers with simple structures. Recently, [34] pointed out that the three-subset division property has a couple of known problems when applied to cube attacks, and proposed a modified three-subset bit-based division property without the “unknown” set to overcome these problems. By modeling this modified version of the three-subset bit-based division property for our cases with small n -bit S-boxes, where $n \in \{5, 7, 9\}$, we confirm the practical results given in Table 2.

However, as far as we know, it is still an open problem to model the (modified) three-subset bit-based division property for a larger S-box of size bigger than 9. The S-boxes we focus on in this paper can be described as a (low-degree) polynomial function in \mathbb{F}_{2^n} , where n is much larger than 9. Therefore, *the division property, which is commonly believed as the most efficient tool to find the best integral distinguishers, might not help us as much for the ciphers we focus on.*

⁷ <https://github.com/HarryR/ethsnarks/blob/master/src/gadgets/mimc.hpp>.

5 Key-Recovery Attack on MiMC

Since the security margin of MiMC with respect to a (secret-key) higher-order distinguisher is of only 1 or 2 round(s) depending on n , it is potentially possible to extend a distinguisher to a key-recovery attack. Given a subspace \mathcal{V} of plaintexts whose sum is equal to zero after r rounds, we can consider $r + 1$ rounds, partially guess the last subkey and decrypt, and filter wrong key guesses that do not satisfy the zero sum:

$$\mathcal{V} \oplus v \xrightarrow{R^r(\cdot)} \underbrace{\bigoplus_{w \in \mathcal{V} \oplus v} R^r(w) = 0}_{\text{Higher-order distinguisher}} \xleftarrow[\text{Key guessing}]{R^{-1}(\cdot)} \underbrace{\{R^{r+1}(w) \mid w \in \mathcal{V} \oplus v\}}_{\text{Ciphertexts}}.$$

However, since the subkeys of MiMC are equal to the master key plus constants, and due to the single full-state S-box, even a (partial) decryption of a single round requires guessing the full key. As a result, a key-recovery attack on full MiMC based on this strategy seems infeasible.

In this section, we present an alternative strategy that allows to break full-round MiMC. Since a trivial key-guessing approach is inefficient, our idea is to construct a polynomial of low degree, which we can then try to solve.

5.1 Strategy of the Attack

From Proposition 2 and Proposition 3, a zero sum can be set up for at least $\lceil (n - 1) \log_3(2) \rceil - 1 = \lceil n \log_3(2) \rceil - \varepsilon$ rounds in the encryption and decryption direction with a vector space $\mathcal{V} \oplus v$ of dimension $n - 1$, where $\varepsilon \in \{1, 2\}$. Recalling that $\lceil n \cdot \log_3(2) \rceil$ is the number of rounds of full MiMC, we define r_{ZS} , r_{KR} as

$$r_{\text{ZS}} = \lceil (n - 1) \log_3(2) \rceil - 1 \quad \text{and} \quad r_{\text{KR}} = 1 + (\lceil n \log_3(2) \rceil - \lceil (n - 1) \log_3(2) \rceil),$$

where r_{ZS} is the number of rounds that we can cover with a zero sum, $r_{\text{KR}} = \lceil n \cdot \log_3(2) \rceil - r_{\text{ZS}} \in \{1, 2\}$.

Let $f^r(x, K)$ be the function corresponding to r rounds of $\text{MiMC}_k(\cdot)$ (and $f^{-r}(x, K)$ be r rounds of decryption, $\text{MiMC}_k^{-1}(\cdot)$), where x is the input text and K is a symbolic variable that represents the secret key k . We intend to use these functions to create a polynomial from which we can deduce k . More precisely, for a fixed vector space $\mathcal{V} \oplus v$, we consider the equations

$$\underbrace{\bigoplus_{x \in \text{MiMC}_k^{-1}(\mathcal{V} \oplus v)} f^{r_{\text{KR}}}(x, K) = 0}_{= F(K)} \quad \text{and} \quad \underbrace{\bigoplus_{x \in \text{MiMC}_k(\mathcal{V} \oplus v)} f^{-r_{\text{KR}}}(x, K) = 0}_{= G(K)}. \quad (9)$$

After having received all x values from an oracle, the attacker can construct one of the polynomials $F(K) = 0$ or $G(K) = 0$. The secret key k can now be determined by finding the roots of either of these polynomials.

In the case of MiMC, the degree of a single encryption round is 3, while the degree of a single decryption round is $(2^{n+1} - 1)/3$ (which is significantly larger than 3 for large n). Due to the slow degree growth in the encryption direction of MiMC, we will focus on finding the roots of $F(K)$ given in Eq. (9).

Finding the Roots of Univariate Polynomials. Let $F(X) \in \mathbb{F}_{2^n}[X]/\langle X^{2^n} + X \rangle$ be a univariate polynomial of degree D . Furthermore, let $M(D)$ denote a number such that multiplying two polynomials of degree $\leq D$ over \mathbb{F}_{2^n} requires $\mathcal{O}(M(D))$ operations in \mathbb{F}_{2^n} . For instance, a straightforward method would yield $M(D) = D^2$, whereas $M(D) = D \cdot \log(D) \cdot \log \log(D)$ holds for methods based on fast Fourier transforms [21]. The *Berlekamp algorithm* for determining the roots of F is then expected to require $\mathcal{C} \in \mathcal{O}(M(D) \log(D) \log(2^n D))$ operations in \mathbb{F}_{2^n} (see [29, Chapter 14.5]).

5.2 Details of the Attack

Assume $\mathcal{V} \oplus v$ is a coset of a subspace \mathcal{V} of dimension $n - 1$. We define

$$\mathcal{W} = \text{MiMC}_k^{-1}(\mathcal{V} \oplus v) \equiv \{\text{MiMC}_k^{-1}(x) \in \mathbb{F}_{2^n} \mid x \in \mathcal{V} \oplus v\}$$

under a fixed secret key k . Here, we present the details of the attack for the cases $r_{\text{KR}} = 1$ and $r_{\text{KR}} = 2$, and we analyze the computational cost. We introduce the following notation:

$$\forall d \in \mathbb{N} : \quad \mathcal{P}_d := \bigoplus_{x \in \mathcal{W}} x^d, \tag{10}$$

and whenever possible we will make use of the fact that squaring is a linear operation over \mathbb{F}_{2^n} . More specifically, computing \mathcal{P}_{2d} only requires a single squaring operation once \mathcal{P}_d is calculated:

$$\mathcal{P}_{2d} := \bigoplus_{x \in \mathcal{W}} x^{2d} = \left(\bigoplus_{x \in \mathcal{W}} x^d \right)^2 = \mathcal{P}_d^2. \tag{11}$$

This allows to reduce the total number of XOR operations.

Case: $r_{\text{KR}} = 1$. Since a single round of MiMC is described by $(x \oplus k)^3 = k^3 \oplus k^2 \cdot x \oplus k \cdot x^2 \oplus x^3$, the function $F(K)$ is given by

$$F(K) = K^2 \cdot \mathcal{P}_1 \oplus K \cdot \mathcal{P}_2 \oplus \mathcal{P}_3.$$

A complete pseudo code of the attack can be found in Algorithm 1, which makes it easy to see that the cost of the attack is well approximated by

- $|\mathcal{V}| = 2^{n-1}$ multiplications,
- $|\mathcal{V}| = 2^{n-1} + 1$ squarings,
- $2 \cdot |\mathcal{V}| + 1 = 2^n + 1$ n -bit XOR operations,
- cost of finding the roots of a univariate polynomial of degree 2.

Algorithm 1: Attack on MiMC – Case: $r_{\text{KR}} = 1$.

Input: Vector subspace \mathcal{V} of ciphertexts of dimension $\dim(\mathcal{V}) = n - 1$.

Output: Secret key k .

```

1  $\mathcal{P}_1, \mathcal{P}_2, \mathcal{P}_3 \leftarrow 0$ .
2 for  $x \in \mathcal{V} \oplus v$  do
3    $p \leftarrow \text{MiMC}_k^{-1}(x)$  from the decryption oracle.
4    $\mathcal{P}_1 \leftarrow \mathcal{P}_1 \oplus p$ .
5    $q \leftarrow p^2$ .
6    $\mathcal{P}_3 \leftarrow \mathcal{P}_3 \oplus q \cdot p$ .
7    $\mathcal{P}_2 \leftarrow (\mathcal{P}_1)^2$ .
8    $F(K) = \mathcal{P}_1 \cdot K^2 \oplus \mathcal{P}_2 \cdot K \oplus \mathcal{P}_3$ .
9   Find a solution  $k$  of  $F(K) = 0$  – see Section 5.1 (filter multiple solutions by
   brute force).
10 return  $k$ .
```

Case: $r_{\text{KR}} = 2$. The attack for the case $r_{\text{KR}} = 2$ is similar. From Eq. (8) (using $k_0 = k, k_1 = k \oplus c_1$ and $k_2 = 0$), the function $F(K)$ is described by

$$\begin{aligned}
 F(K) = & K^8 \cdot \mathcal{P}_1 \oplus K^5 \cdot \mathcal{P}_2 \oplus K^4 \cdot (\mathcal{P}_2 \cdot c_1 \oplus \mathcal{P}_1) \oplus K^3 \cdot (\mathcal{P}_4 \oplus \mathcal{P}_2) \\
 & \oplus K^2 \cdot (\mathcal{P}_4 \cdot c_1 \oplus \mathcal{P}_3 \oplus \mathcal{P}_1 \cdot c_1^2) \oplus K \cdot (\mathcal{P}_8 \oplus \mathcal{P}_6 \oplus \mathcal{P}_2 \cdot c_1^2) \oplus (\mathcal{P}_9 \oplus \mathcal{P}_6 \cdot c_1 \oplus \mathcal{P}_3 \cdot c_1^2),
 \end{aligned}$$

where c_1 is the round constant of the first round. As also noted in Sect. 3.2, while \mathcal{P}_9 is the largest \mathcal{P}_d in this expression, both \mathcal{P}_5 and \mathcal{P}_7 are missing, and hence do not need to be computed. A complete pseudo code of the attack can be found in Algorithm 2. Again, it is easy to see that the cost of the attack is well approximated by

- $2 \cdot |\mathcal{V}| + 6 = 2^n + 6$ multiplications,
- $2 \cdot |\mathcal{V}| + 4 = 2^n + 4$ squarings,
- $3 \cdot |\mathcal{V}| + 8 = 3 \cdot 2^{n-1} + 8$ n -bit XOR operations,
- cost of finding the roots of a univariate polynomial of degree 8.

5.3 Complexity Estimation

As we have just seen, our attack requires half of the code book (namely, 2^{n-1} chosen ciphertexts). Here we show that our attacks are better than exhaustive search (from the computational point of view). In order to do this, we measure the time complexities in equivalent encryption operations.

A single encryption round in MiMC requires one addition, one squaring operation, and one multiplication in the extension field. Since the cost of a single n -bit XOR operation is much smaller than the cost of a multiplication over \mathbb{F}_{2^n} , and since the number of XOR operations is similar to the number of multiplications, in the following we do not consider XOR operations. After this simplification, we find that the time complexity of $r_{\text{KR}} = 1$ is dominated by 2^{n-1} squaring and multiplication operations or, equivalently, 2^{n-1} encryption rounds. A similar line of reasoning reveals that $r_{\text{KR}} = 2$ is comparable to 2^n encryption rounds.

Algorithm 2: Attack on MiMC – Case: $r_{\text{KR}} = 2$.**Input:** Vector subspace \mathcal{V} of ciphertexts of dimension $\dim(\mathcal{V}) = n - 1$.**Output:** Secret key k .

```

1  $\mathcal{P}_1, \mathcal{P}_2, \mathcal{P}_3, \dots, \mathcal{P}_9 \leftarrow 0$ .
2 for  $x \in \mathcal{V} \oplus v$  do
3    $p \leftarrow \text{MiMC}_k^{-1}(x)$  from the decryption oracle.
4    $\mathcal{P}_1 \leftarrow \mathcal{P}_1 \oplus p$ .
5    $q_2 \leftarrow p^2$ .
6    $q_3 \leftarrow q_2 \cdot p$ .
7    $\mathcal{P}_3 \leftarrow \mathcal{P}_3 \oplus q_3$ .
8    $q_6 \leftarrow q_3^2$ .
9    $\mathcal{P}_9 \leftarrow \mathcal{P}_9 \oplus q_6 \cdot q_3$ .
10  $\mathcal{P}_2 \leftarrow (\mathcal{P}_1)^2$ .
11  $\mathcal{P}_4 \leftarrow (\mathcal{P}_2)^2$ .
12  $\mathcal{P}_6 \leftarrow (\mathcal{P}_3)^2$ .
13  $\mathcal{P}_8 \leftarrow (\mathcal{P}_4)^2$ .
14  $F(K) = K^8 \cdot \mathcal{P}_1 \oplus K^5 \cdot \mathcal{P}_2 \oplus K^4 \cdot (\mathcal{P}_2 \cdot c_1 \oplus \mathcal{P}_1) \oplus K^3 \cdot (\mathcal{P}_4 \oplus \mathcal{P}_2) \oplus K^2 \cdot (\mathcal{P}_4 \cdot$ 
    $c_1 \oplus \mathcal{P}_3 \oplus \mathcal{P}_1 \cdot c_1^2) \oplus K \cdot (\mathcal{P}_8 \oplus \mathcal{P}_6 \oplus \mathcal{P}_2 \cdot c_1^2) \oplus (\mathcal{P}_9 \oplus \mathcal{P}_6 \cdot c_1 \oplus \mathcal{P}_3 \cdot c_1^2)$ .
15 Find a solution  $k$  of  $F(K) = 0$  (filter multiple solutions by brute force).
16 return  $k$ .
```

Since the cost of solving a single low-degree equation is negligible, and one unit of encryption contains $\lceil n \cdot \log_3(2) \rceil$ rounds, it follows that the cost of our attacks is about

$$\frac{r_{\text{KR}} \cdot 2^{n-1}}{\lceil n \cdot \log_3(2) \rceil}$$

encryptions for $r_{\text{KR}} \in \{1, 2\}$. That is, the computational cost of the key-recovery part of our attacks is upper-bounded by $2^{n - \log_2(n) + 1}$, and hence the total cost is smaller than that of a brute-force attack (namely, 2^n encryptions) for each $n \geq 3$.

5.4 Practical Verification

We implemented Algorithm 1 and Algorithm 2 in the computer algebra system Magma, and verified both algorithms for all odd integers $n \in [5, 35]$. We note that Algorithm 1 ($r_{\text{KR}} = 1$) yields the correct answer for all the tested $5 \leq n \leq 35$, even if $\lceil n \log_3(2) \rceil \neq \lceil (n-1) \log_3(2) \rceil$. Namely, in practice it is possible to cover one more round with a zero sum than what we theoretically expect. In other words, $\lceil (n-1) \log_3(2) \rceil$ rounds of the decryption function of MiMC fail to obtain the maximum algebraic degree for these parameters, which is reached after $\lceil (n-1) \log_3(2) \rceil + 1$ rounds (see [28, App. B] for more details on the degree growth of MiMC^{-1}). Since we are not able to prove this behavior for larger values of n , we leave it as an open question whether Algorithm 1 can be applied to MiMC for odd integers $n > 35$.

Considerations on Data and Computational Costs of This Attack. A possible drawback of our attack is the cost. Since we are not able to provide an estimation of the growth of the degree in the decryption direction, we can only exploit the fact that a certain number of rounds are necessary in order to achieve maximum degree. It follows that the attacker is forced to use half of the code book in order to set up the attack, which also has an impact on the computational cost.

Even if our attack is not practical, we believe it provides valuable theoretical insight. It is also in line with several other attacks found in the literature, which are set up under a similar assumption on the data and/or computational cost. To give some concrete examples, consider the case of zero-correlation attacks [14], which exploit linear approximations that hold with probability $\frac{1}{2}$. The crucial limitation for basic zero-correlation linear cryptanalysis is that it requires half of the code book. Only follow-up works have been able to reduce this data requirement, including the more powerful distinguisher called multiple zero-correlation (MPZC) linear distinguisher proposed in [15], which exploits the fact that there are numerous zero-correlation linear approximations in susceptible ciphers. While needing (close to) the full code book is an inherent property of zero-correlation attacks, the reason for the high data complexity in our case is purely due to the specification of MiMC and the attacked number of rounds, and not due to an inherent property of our attack.

Splice-and-cut meet-in-the-middle attacks and biclique attacks are other examples of attacks that often come with time complexities relatively close to exhaustive search. Indeed, an extension of the biclique approach first described in [13] has a brute-force phase for a number of rounds as part of the attack. It can in principle work for any number of rounds and is hence best described as a particular optimization of brute-force key guessing. However, later variants then showed examples where the gain over brute force was in the order of millions [37]. Still, we note that the complexity of biclique attacks scales differently than our attack, whose runtime cost depends strongly on the details of the target cipher MiMC.

Finally, we point out that any attack that is better than brute force is relevant, even if it requires unrealistic amounts of data or storage. Indeed, the main goal of cryptanalysis is finding a “certificated weakness”, that is, an evidence that the cipher does not perform as advertised. In other words, in academic cryptography, a weakness or a break in a scheme is usually defined quite conservatively: It may require impractical amounts of time, memory, or data.

The Number of Rounds Needed for Security. It may be of interest to estimate the number of rounds needed for MiMC to be resistant against this attack. To this end, we bound the operations needed to compute all monomials of odd degree, up to a maximum degree D .

Lemma 2. *Let $1 \leq D \leq 2^n - 1$ and $x \in \mathbb{F}_{2^n}$. The overall number of operations needed to compute all odd powers x^i for $i \in [3, D]$ is given by 1 squaring and $\lfloor \frac{D-1}{2} \rfloor$ multiplications.*

Proof. From x , calculate and store $q := x^2$. The odd powers of x can now be successively computed as $x^{i+2} = x^i \cdot q$ for all odd integers i in the interval $[1, D - 2]$. This yields a total of 1 squaring and $\lfloor \frac{D-1}{2} \rfloor$ multiplications. \square

Assume for simplicity that $\lceil n \cdot \log_3(2) \rceil - 1$ rounds can be covered by a zero sum, and that the cost of solving the final polynomial equation is negligible. As before, we expect the time complexity to be dominated by the number of operations needed to construct the polynomial $F(K)$. Since the degree of this polynomial is upper-bounded by $3^{r_{KR}}$, by Lemma 2 at most $\lceil (3^{r_{KR}} - 1)/2 \rceil \cdot 2^{n-1}$ multiplications are required to compute all monomials with odd exponents in $F(K)$ (where all monomials with even exponents are computed via Eq. (11)).

Since one encryption of MiMC costs $\lceil n \cdot \log_3(2) \rceil$ multiplications, the number of extra rounds ρ for MiMC must satisfy

$$(3^{\rho+1} - 1) \cdot 2^{n-2} \geq 2^n \cdot (\lceil n \cdot \log_3(2) \rceil + \rho)$$

in order to provide security against the attack just presented. This would, for example, require at least $\rho = 5$ extra rounds for $n = 129$ (more generally, if R is the number of rounds of MiMC- n/n , then $\rho \approx \lceil \log_3(2 \cdot R) \rceil$ more rounds are sufficient to restore the security⁸). We remark that *this rough estimation is not intended to replace the number of rounds proposed by the designers.*

6 An Algebraic Attack on Ciphers with Low-Degree Round Functions

Here we generalize the key-recovery attack on MiMC described in Sect. 5 and discuss a generic attack strategy for any block cipher working over $(\mathbb{F}_{2^n})^t$, where $n, t \in \mathbb{N}$, $t \geq 2$ and $n \geq 3$.

6.1 Setting

We consider an r -round block cipher $E_k^r : (\mathbb{F}_{2^n})^t \rightarrow (\mathbb{F}_{2^n})^t$ with

$$E_k^r(x) = (R_r \circ R_{r-1} \circ \dots \circ R_1)(x \oplus k),$$

and where $R, R_i : (\mathbb{F}_{2^n})^t \rightarrow (\mathbb{F}_{2^n})^t$ are defined by $R_i(x) = R(x) \oplus k^{(i)}$. Here, R denominates the (nonlinear) round function. Since E_k^r consists of t components, we can write

$$E_k^r(x) = (E_{k,1}^r(x), \dots, E_{k,t}^r(x)),$$

where $E_{k,i}^r : (\mathbb{F}_{2^n})^t \rightarrow \mathbb{F}_{2^n}$. We denote the compositional inverse of E_k^r by E_k^{-r} . We assume that

⁸ In more details, $\rho \geq \log_3(4 \cdot (R + \rho) + 1) - 1$. The previous estimation is obtained by assuming $\rho \leq R/2$.

- (1) the i -th round key $k^{(i)} \in (\mathbb{F}_{2^n})^t$ is derived from the master key $k = (k_1, \dots, k_t) \in (\mathbb{F}_{2^n})^t$ by some *low-degree* (e.g., linear) key schedule,
- (2) the round function R can be described by a polynomial

$$R(x = (x_1, \dots, x_t)) = \bigoplus_{\substack{j=(j_1, \dots, j_t) \in \{0, 1, \dots, 2^n - 1\}^t \\ j_1 + \dots + j_t \leq d}} \alpha_j \cdot x_1^{j_1} \cdots x_t^{j_t}$$

of *low-degree* d with coefficients $\alpha_j \in (\mathbb{F}_{2^n})^t$.

Our attack requires the symbolic evaluation of the encryption function $E_k^{r'}$ for a small number of rounds r' to be relatively easy, which motivates the requirements of a low-degree round function R and a low-degree key schedule. This ensures that the polynomial representation of $E_k^{r'}$ can be computed efficiently. In both cases, *low-degree* means *low compared to the size of the field* \mathbb{F}_{2^n} , i.e., $d \ll 2^n - 1$. A cipher in the literature that satisfies above assumptions and does indeed use low-degree round functions is, e.g., HadesMiMC [31].

6.2 Strategy of the Attack

The idea of our generic attack is to recover the secret master key k of a cipher E_k^r by exploiting a given higher-order distinguisher over the subset $\mathcal{X} \subseteq (\mathbb{F}_{2^n})^t$ covering $1 \leq r_{ZS} < r$ rounds in the encryption or the decryption direction. For the sake of simplicity, we follow the approach of the attack on MiMC in Sect. 5 and assume that the higher-order distinguisher covers r_{ZS} rounds in the decryption direction.

In our attack, we symbolically evaluate $E_k^{r_{KR}}(y)$ with respect to the remaining $r_{KR} := r - r_{ZS}$ rounds in the encryption direction and obtain polynomials ($1 \leq i \leq t$)

$$E_{(K_1, \dots, K_t), i}^{r_{KR}}(Y) \in \mathbb{F}_{2^n}[K_1, \dots, K_t, Y_1, \dots, Y_t]$$

over \mathbb{F}_{2^n} with the master key words K_j and plaintext variables $(Y_1, \dots, Y_t) =: Y$ as indeterminates – in short, one polynomial for each of the t components of $E_k^{r_{KR}}(y)$. In general, we work with $r_{KR} \ll r_{ZS}$, since the symbolic evaluation of $E_k^{r_{KR}}(y)$ is expensive.

Having a zero sum after r_{ZS} rounds in the decryption direction with respect to the subset $\mathcal{X} \subseteq (\mathbb{F}_{2^n})^t$ means that

$$\bigoplus_{x \in \mathcal{X}} E_k^{-r_{ZS}}(x) = 0.$$

The main observation behind our attack is the following: We exploit the relation⁹

$$0 = \bigoplus_{x \in \mathcal{X}} E_k^{-r_{ZS}}(x) = \bigoplus_{x \in \mathcal{X}} (E_k^{r_{KR}} \circ E_k^{-r})(x) = \bigoplus_{y \in E_k^{-r}(\mathcal{X})} E_k^{r_{KR}}(y) \quad (12)$$

⁹ Note that in this representation, $E_k^r = E_k^{r_{ZS}} \circ E_k^{r_{KR}}$ and $E_k^{-r_{ZS}} = E_k^{r_{KR}} \circ E_k^{-r}$.

Algorithm 3: Attack on a generic cipher E_k^r over $(\mathbb{F}_{2^n})^t$.

Input: Number of rounds r of the cipher E_k^r , number of rounds r_{ZS} in the decryption direction and a subset $\mathcal{X} \subseteq (\mathbb{F}_{2^n})^t$ satisfying the zero sum $\bigoplus_{x \in \mathcal{X}} E_k^{-r_{\text{ZS}}}(x) = 0$.

Output: Secret key $k = (k_1, \dots, k_t)$.

```

1  $r_{\text{KR}} \leftarrow r - r_{\text{ZS}}$ .
2 for each  $1 \leq i \leq t$  do
3   Compute the symbolic evaluation
    $f_i = f_i(Y_1, \dots, Y_t, K_1, \dots, K_t) = E_{(K_1, \dots, K_t), i}^{r_{\text{KR}}}(Y_1, \dots, Y_t)$  of word  $i$  in the
   encryption direction for  $r_{\text{KR}}$  rounds.
4   for each monomial  $Y_1^{i_1} \dots Y_t^{i_t} \cdot K_1^{j_1} \dots K_t^{j_t}$  in  $f_i$  with  $i_1 + \dots + i_t \geq 1$  do
5      $\mathcal{P}_{i_1, \dots, i_t} \leftarrow 0$ .
6     for each  $x \in \mathcal{X}$  do
7        $y = (y_1, \dots, y_t) \leftarrow E_k^{-r}(x)$ , via the decryption oracle.
8        $\mathcal{P}_{i_1, \dots, i_t} \leftarrow \mathcal{P}_{i_1, \dots, i_t} \bigoplus y_1^{i_1} \dots y_t^{i_t}$ .
9       Replace  $Y_1^{i_1} \dots Y_t^{i_t} \cdot K_1^{j_1} \dots K_t^{j_t}$  with  $\mathcal{P}_{i_1, \dots, i_t} \cdot K_1^{j_1} \dots K_t^{j_t}$ .
10     $F_i(K_1, \dots, K_t) \leftarrow f_i(K_1, \dots, K_t)$ .
11 Find a solution  $k = (k_1, \dots, k_t)$  of  $F_1(k_1, \dots, k_t) = \dots = F_t(k_1, \dots, k_t) = 0$ .
12 return  $k = (k_1, \dots, k_t)$ .
```

to set up the following equations ($1 \leq i \leq t$) over \mathbb{F}_{2^n} in the variables k_1, \dots, k_t :

$$F_i(k_1, \dots, k_t) := \bigoplus_{y \in E_k^{-r}(\mathcal{X})} E_{(k_1, \dots, k_t), i}^{r_{\text{KR}}}(y) = 0. \tag{13}$$

Again, $E_{(k_1, \dots, k_t), i}^{r_{\text{KR}}}(y)$ denotes the symbolic evaluation of the i -th word after r_{KR} rounds in the encryption direction with the master key words as variables k_1, \dots, k_t and evaluated at $y \in \mathbb{F}_{2^n}$. Once we have set up the equation system arising from Eq. (13), we apply Gröbner basis techniques to solve this system over \mathbb{F}_{2^n} for the key variables k_1, \dots, k_t .

In Algorithm 3 we summarize the approach of our generic attack and present a pseudo code of the attack procedure. For completeness, a rough complexity estimation of the attack is derived in [28, App. E].

6.3 Comparison with Related Work

Interpolation Attacks. Originally introduced as a standalone attack, interpolation attacks [36] are algebraic attacks that express the (potentially round-reduced) cipher as a polynomial equation with unknown, key-dependent coefficients, and recover these coefficients from known inputs and outputs. More recently, this approach has been combined as a key-recovery approach together with integral distinguishers.

Attack on CAST. In an attack [43] on the CAST cipher the authors use a higher-order differential distinguisher to set up an equation system and finally solve

this systems for the key variables. In contrast to our attack, the authors of [43] work with linear equation systems over \mathbb{F}_2 . While this is sufficient for CAST, working at bit level is in general more expensive than working on word level when analyzing ciphers that are natively defined at word level.

Optimized Interpolation Attacks. One type of optimized interpolation attacks was described in [23], where the authors find attacks on reduced-round versions of LowMC which are more efficient than previous attacks based on key guessing [25]. A similar attack was also used to break the full-round version of the FRIT permutation in an Even–Mansour setting [26]. The overall strategy of this interpolation attack is to find a distinguisher (for example a constant sum in the encryption direction in the case of LowMC) with which one attacks the construction by finding the unknown monomials of the sums of the symbolic representations in the inverse direction. By determining these (key-dependent) monomials, the full key can eventually be found. Since the approach in [23] shares some similarities with our proposal, we describe the differences between these two strategies in detail.

The main difference regarding the two strategies concerns the way in which the system of equations $F_i(K) = 0$ is constructed and consequently solved:

- In [23], the idea is to construct the function using a “standard” interpolation technique. Specifically, the attacker does not care about the specification of the monomials of F , which are simply considered as unknowns. Hence, the idea is to recover (interpolate) the unknown coefficients of $F_K(C)$, and then use various ad-hoc techniques (which are not part of the framework described in this section) in order to recover the actual secret key.
- In our case, we heavily exploit the simple algebraic structure of the round function in order to construct the system of equations $F_i(K) = 0$. In other words, the system of equations is constructed by using a symbolic evaluation and not by interpolation techniques.

We emphasize that the possibility to set up one of the two attacks does not imply the possibility to set up the other one. For example, it seems hard to use the attack presented in [23] against full-round MiMC, while we show that our strategy can break it. Indeed, since we already need 2^{n-1} data for the distinguishing property (i.e., half of the code book), we do not see how to apply the approach from [23] to MiMC without further increasing the data complexity due to data needed for the interpolation step.

Attack on PyjamaMask. Only recently, a similar attack on PyjamaMask, competing in the ongoing NIST call for lightweight authenticated encryption, has been presented [27]. The authors propose an attack on the full block cipher PyjamaMask-96 by combining higher-order differentials with an in-depth ad-hoc analysis of the system of equations obtained for 2.5 rounds of PyjamaMask-96. As is the case for CAST, the attack is set up at bit level.

Cube Attacks. Although our attack and cube attacks [24] exploit low degrees in the polynomial description of a cipher, they are quite different from a conceptual point of view and can be regarded as two different cryptanalytic methods. To justify this conclusion, we briefly present the idea behind cube attacks and contrast them with our attack ideas.

Given a cipher with input variables x_0, \dots, x_{n-1} as the public variables (IV bits, plaintext bits, tweak bits, etc.), and x_n, \dots, x_{n+m-1} as the secret variables (key bits), the output of the cipher can be regarded as a polynomial $f = f(x)$ in $x = (x_0, \dots, x_{n+m-1})$. For every set $I \subset \{0, \dots, n-1\}$, f can be uniquely decomposed into

$$f = t_I \cdot f_{S(I)} + q,$$

where $t_I := \prod_{i \in I} x_i$ denotes the product of all variables indexed by elements in I , the polynomial $f_{S(I)}$ does not contain any variables from t_I , and where q misses at least one variable from t_I . The polynomial $f_{S(I)}$ is also called the *superpoly* with respect to I . For any subset $I \subseteq \{0, \dots, n-1\}$ of size $|I|$, the authors of [24] call the set C_I of $2^{|I|}$ vectors, where all the $|I|$ variables indexed by I range over all possible combinations of elements in \mathbb{F}_2 and the remaining $n+m-|I|$ variables remain undetermined, a $|I|$ -dimensional *Boolean cube*. Then the sum of f over all values in the cube C_I yields the equation of polynomials

$$\bigoplus_{v \in C_I} f(v) = f_{S(I)}.$$

Cube attacks consist of two steps. First, attackers recover the superpoly in the offline phase. In this phase, the attacker might need to try sufficiently many cubes and assignments for the remaining public variables such that the superpoly $f_{S(I)}$ is a balanced function of the secret variables. Moreover, determining the actual coefficients of $f_{S(I)}$ requires the additional assumption that the attacker is allowed to tweak both public and secret variables. Then, with this usable superpoly, during the online phase, the attacker leaves the secret variables undetermined and queries the encryption oracle with every value $c \in C_I$ and gets $f(c) \in \mathbb{F}$. Eventually, the attacker computes

$$f_I := \bigoplus_{c \in C_I} f(c).$$

The secret key information can be recovered by solving the corresponding equation system $f_I = f_{S(I)}$.

Compared with our attack, cube attacks involve an initial step of finding balanced superpolies that contain independent secret variables. Apart from that, cube attacks do *not* exploit the algebraic structure of a cipher, since they rely on the assumption of tweakable black box polynomials. In this sense, our attack is different, since it makes heavy use of the algebraic structure of a cipher when symbolically evaluating a certain number of rounds. Furthermore, cube attacks use the assumption that both key and plaintext variables are tweakable, while we rely on the assumption that some rounds of the cipher can be efficiently evaluated symbolically (which is why we work with low-degree round functions).

7 Concluding Remarks and Future Work

Reducing the Cost of the Attack. As shown in [28, App. E], two steps – namely, (1st) the construction of the system of equations $F_i(k_1, \dots, k_t) = 0$ for $1 \leq i \leq t$ and (2nd) solving such a system – mainly constitute the cost of the attack. In general, it could make sense to balance the costs of the two steps in order to either minimize the total cost of the attack or maximize the number of rounds that can be broken.

In more detail, consider the case in which the cost of the attack is well approximated by the cost of *constructing* the system of equations $F_i(K) = 0$. Since this cost grows with the size of the subspace \mathcal{V} , one strategy could be to consider a smaller subset \mathcal{X} .¹⁰ Obviously, this implies in general the possibility to cover fewer rounds r_{ZS} using a higher-order distinguisher, which means that more rounds r_{KR} must be covered in general. However, the overall cost of the attack may benefit from this strategy. On the other hand, the case in which the attack cost is well approximated by the cost of *solving* the system of equations $F_i(K) = 0$ requires the opposite strategy.

Moreover, we point out that the attacks can be improved by exploiting the details of the cipher. To give a concrete example, consider the case of MiMC given in Algorithm 1: The attack and its computational complexity benefit from the fact that $F(K)$ does not depend on \mathcal{P}_5 or \mathcal{P}_7 . As another example, consider the case of an SPN cipher where the round function is defined as

$$R(x = (x_1, \dots, x_t)) = M \times (S(x_1), S(x_2), \dots, S(x_t)),$$

where $M \in (\mathbb{F}_{2^n})^{t \times t}$ and $S : \mathbb{F}_{2^n} \rightarrow \mathbb{F}_{2^n}$ (here, ‘ \times ’ denotes matrix-vector multiplication). The cost of the attack can potentially be reduced by taking into account the fact that all monomials in the polynomial representation R depend only on a single variable x_i .

Further Generalization: Ciphers over \mathbb{F}_p . Finally, the attack strategy can be generalized to include ciphers over $(\mathbb{F}_p)^t$ for a prime p . This is of particular importance since many of the new applications named in the introduction (e.g., STARKs and MPC) natively work over \mathbb{F}_p , which means that many of the recently proposed primitives are natively constructed over \mathbb{F}_p . We remark that the strategy of the attack does not depend on the details of the field \mathbb{F} . Hence, the only thing that seems to preclude this possibility seems to be a lack of knowledge regarding efficient distinguishers over $(\mathbb{F}_p)^t$. Indeed, while it is well-known how to find a higher-order distinguisher over Boolean fields (e.g., by exploiting division property tools present in the literature [47, 51, 53]), the same is not yet true for prime fields.

¹⁰ We note that we cannot adopt this strategy for MiMC since we are not able to predict the growth of the degree of MiMC^{-1} . With such an estimation, the strategy proposed here can potentially reduce the cost of the attack.

Acknowledgements. The authors thank the anonymous reviewers for their valuable comments and suggestions. Qingju Wang is funded by the University of Luxembourg Internal Research Project (IRP) FDISC. This project has received funding from the European Union’s Horizon 2020 research and innovation programme (H2020 ICT 825225: Safe-DEED) and the European Research Council (ERC 681402: SOPHIA).

References

1. Albrecht, M.R., Cid, C., Grassi, L., Khovratovich, D., Lüftenegger, R., Rechberger, C., Schafneger, M.: Algebraic cryptanalysis of STARK-friendly designs: application to MARVELLOUS and MiMC. In: Galbraith, S.D., Moriai, S. (eds.) ASIACRYPT 2019. LNCS, vol. 11923, pp. 371–397. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-34618-8_13
2. Albrecht, M.R., Grassi, L., Perrin, L., Ramacher, S., Rechberger, C., Rotaru, D., Roy, A., Schafneger, M.: Feistel structures for MPC, and more. In: Sako, K., Schneider, S., Ryan, P.Y.A. (eds.) ESORICS 2019. LNCS, vol. 11736, pp. 151–171. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-29962-0_8
3. Albrecht, M., Grassi, L., Rechberger, C., Roy, A., Tiessen, T.: MiMC: efficient encryption and cryptographic hashing with minimal multiplicative complexity. In: Cheon, J.H., Takagi, T. (eds.) ASIACRYPT 2016. LNCS, vol. 10031, pp. 191–219. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-53887-6_7
4. Albrecht, M.R., Rechberger, C., Schneider, T., Tiessen, T., Zohner, M.: Ciphers for MPC and FHE. In: Oswald, E., Fischlin, M. (eds.) EUROCRYPT 2015. LNCS, vol. 9056, pp. 430–454. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-46800-5_17
5. Aly, A., Ashur, T., Ben-Sasson, E., Dhooghe, S., Szepieniec, A.: Design of Symmetric-Key Primitives for Advanced Cryptographic Protocols. IACR Cryptology ePrint Archive, Report 2019/426 (2019)
6. Ashur, T., Dhooghe, S.: MARVELLOUS: a STARK-Friendly Family of Cryptographic Primitives. IACR Cryptology ePrint Archive, Report 2018/1098 (2018)
7. Aumasson, J.P., Meier, W.: Zero-sum distinguishers for reduced Keccak-f and for the core functions of Luffa and Hamsi (2009), presented at the Rump Session of CHES 2009, <https://131002.net/data/papers/AM09.pdf>
8. Ben-Sasson, E., Bentov, I., Horesh, Y., Riabzev, M.: Scalable, transparent, and post-quantum secure computational integrity. IACR Cryptology ePrint Archive **2018**, 46 (2018)
9. Bertoni, G., Daemen, J., Peeters, M., Van Assche, G.: On the indistinguishability of the sponge construction. In: Smart, N. (ed.) EUROCRYPT 2008. LNCS, vol. 4965, pp. 181–197. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-78967-3_11
10. Bertoni, G., Daemen, J., Peeters, M., Van Assche, G.: Note on zero-sum distinguishers of Keccak-f. <http://keccak.noekeon.org/NoteZeroSum.pdf>
11. Biham, E., Shamir, A.: Differential cryptanalysis of DES-like cryptosystems. In: Menezes, A.J., Vanstone, S.A. (eds.) CRYPTO 1990. LNCS, vol. 537, pp. 2–21. Springer, Heidelberg (1991). https://doi.org/10.1007/3-540-38424-3_1
12. Black, J., Rogaway, P., Shrimpton, T.: Black-box analysis of the block-cipher-based hash-function constructions from PGV. In: Yung, M. (ed.) CRYPTO 2002. LNCS, vol. 2442, pp. 320–335. Springer, Heidelberg (2002). https://doi.org/10.1007/3-540-45708-9_21

13. Bogdanov, A., Khovratovich, D., Rechberger, C.: Biclique cryptanalysis of the full AES. In: Lee, D.H., Wang, X. (eds.) ASIACRYPT 2011. LNCS, vol. 7073, pp. 344–371. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-25385-0_19
14. Bogdanov, A., Rijmen, V.: Linear hulls with correlation zero and linear cryptanalysis of block ciphers. *Des. Codes Cryptogr.* **70**(3), 369–383 (2014). <https://doi.org/10.1007/s10623-012-9697-z>
15. Bogdanov, A., Wang, M.: Zero correlation linear cryptanalysis with reduced data complexity. In: Canteaut, A. (ed.) FSE 2012. LNCS, vol. 7549, pp. 29–48. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-34047-5_3
16. Bonnetain, X.: Collisions on feistel-MiMC and univariate GMiMC. *IACR Cryptology ePrint Archive* **2019**, 951 (2019)
17. Boura, C., Canteaut, A.: On the influence of the algebraic degree of F^{-1} on the algebraic degree of $G \circ F$. *IEEE Trans. Inf. Theor.* **59**(1), 691–702 (2013)
18. Boura, C., Canteaut, A., De Cannière, C.: Higher-order differential properties of KECCAK and *Luffa*. In: Joux, A. (ed.) FSE 2011. LNCS, vol. 6733, pp. 252–269. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-21702-9_15
19. Büinz, B., Bootle, J., Boneh, D., Poelstra, A., Wuille, P., Maxwell, G.: Bulletproofs: short proofs for confidential transactions and more. In: *IEEE Symposium on Security and Privacy*, pp. 315–334. IEEE Computer Society (2018)
20. Canteaut, A., Videau, M.: Degree of composition of highly nonlinear functions and applications to higher order differential cryptanalysis. In: Knudsen, L.R. (ed.) EUROCRYPT 2002. LNCS, vol. 2332, pp. 518–533. Springer, Heidelberg (2002). https://doi.org/10.1007/3-540-46035-7_34
21. Cantor, D.G., Kaltofen, E.: On fast multiplication of polynomials over arbitrary algebras. *Acta Inf.* **28**(7), 693–701 (1991)
22. Carlet, C., Charpin, P., Zinoviev, V.A.: Codes, bent functions and permutations suitable for DES-like cryptosystems. *DCC* **15**(2), 125–156 (1998)
23. Dinur, I., Liu, Y., Meier, W., Wang, Q.: Optimized interpolation attacks on LowMC. In: Iwata, T., Cheon, J.H. (eds.) ASIACRYPT 2015. LNCS, vol. 9453, pp. 535–560. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-48800-3_22
24. Dinur, I., Shamir, A.: Cube attacks on tweakable black box polynomials. In: Joux, A. (ed.) EUROCRYPT 2009. LNCS, vol. 5479, pp. 278–299. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-01001-9_16
25. Dobraunig, C., Eichlseder, M., Mendel, F.: Higher-order cryptanalysis of LowMC. In: Kwon, S., Yun, A. (eds.) ICISC 2015. LNCS, vol. 9558, pp. 87–101. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-30840-1_6
26. Dobraunig, C., Eichlseder, M., Mendel, F., Schofnegger, M.: Algebraic cryptanalysis of variants of FRIT. In: Paterson, K.G., Stebila, D. (eds.) SAC 2019. LNCS, vol. 11959, pp. 149–170. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-38471-5_7
27. Dobraunig, C., Rotella, Y., Schoone, J.: Algebraic and higher-order differential cryptanalysis of Pyjamask-96. *IACR Trans. Symmetric Cryptology* **2020**(1), 289–312 (2020)
28. Eichlseder, M., Grassi, L., Lüftenegger, R., Øygarden, M., Rechberger, C., Schofnegger, M., Wang, Q.: An algebraic attack on ciphers with low-degree round functions: application to full MiMC. *IACR Cryptol. ePrint Arch.* **2020**, 182 (2020)
29. von zur Gathen, J., Gerhard, J.: *Modern Computer Algebra*, (3ed) Cambridge University Press, New York (2013)

30. Grassi, L., Kales, D., Khovratovich, D., Roy, A., Rechberger, C., Schafneger, M.: Starkad and Poseidon: New Hash Functions for Zero Knowledge Proof Systems. Cryptology ePrint Archive, Report 2019/458 (2019)
31. Grassi, L., Lüftenecker, R., Rechberger, C., Rotaru, D., Schafneger, M.: On a generalization of substitution-permutation networks: the HADES design strategy. In: Canteaut, A., Ishai, Y. (eds.) EUROCRYPT 2020. LNCS, vol. 12106, pp. 674–704. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-45724-2_23
32. Grassi, L., Rechberger, C., Rotaru, D., Scholl, P., Smart, N.P.: Mpc-friendly symmetric key primitives. In: ACM Conference on Computer and Communications Security, pp. 430–443. ACM (2016)
33. Guo, J., Liao, G., Liu, G., Liu, M., Qiao, K., Song, L.: Practical collision attacks against round-reduced SHA-3. *J. Cryptology* **33**(1), 228–270 (2020)
34. Hao, Y., Leander, G., Meier, W., Todo, Y., Wang, Q.: Modeling for three-subset division property without unknown subset. In: Canteaut, A., Ishai, Y. (eds.) EUROCRYPT 2020. LNCS, vol. 12105, pp. 466–495. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-45721-1_17
35. Hu, K., Wang, M.: Automatic search for a variant of division property using three subsets. In: Matsui, M. (ed.) CT-RSA 2019. LNCS, vol. 11405, pp. 412–432. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-12612-4_21
36. Jakobsen, T., Knudsen, L.R.: The interpolation attack on block ciphers. *FSE. LNCS* **1267**, 28–40 (1997)
37. Khovratovich, D., Leurent, G., Rechberger, C.: Narrow-bicliques: cryptanalysis of full Idea. In: Pointcheval, D., Johansson, T. (eds.) EUROCRYPT 2012. LNCS, vol. 7237, pp. 392–410. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-29011-4_24
38. Knudsen, L.R.: Truncated and higher order differentials. In: Preneel, B. (ed.) FSE 1994. LNCS, vol. 1008, pp. 196–211. Springer, Heidelberg (1995). https://doi.org/10.1007/3-540-60590-8_16
39. Knudsen, L.R., Rijmen, V.: Known-key distinguishers for some block ciphers. In: Kurosawa, K. (ed.) ASIACRYPT 2007. LNCS, vol. 4833, pp. 315–324. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-76900-2_19
40. Lai, X.: Higher Order Derivatives and Differential Cryptanalysis. In: Blahut, R.E., Costello, D.J., Maurer, U., Mittelholzer, T. (eds.) Communications and Cryptography. The Springer International Series in Engineering and Computer Science (Communications and Information Theory), vol. 276, pp. 227–233. Springer, Boston (1994) https://doi.org/10.1007/978-1-4615-2694-0_23
41. Li, C., Preneel, B.: Improved interpolation attacks on cryptographic primitives of low algebraic degree. In: Paterson, K.G., Stebila, D. (eds.) SAC 2019. LNCS, vol. 11959, pp. 171–193. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-38471-5_8
42. Matsui, M.: Linear cryptanalysis method for DES cipher. In: Helleseht, T. (ed.) EUROCRYPT 1993. LNCS, vol. 765, pp. 386–397. Springer, Heidelberg (1994). https://doi.org/10.1007/3-540-48285-7_33
43. Moriai, S., Shimoyama, T., Kaneko, T.: Higher order differential attack of a CAST cipher. In: Vaudenay, S. (ed.) FSE 1998. LNCS, vol. 1372, pp. 17–31. Springer, Heidelberg (1998). https://doi.org/10.1007/3-540-69710-1_2
44. Parno, B., Howell, J., Gentry, C., Raykova, M.: Pinocchio: Nearly practical verifiable computation. In: IEEE Symposium on Security and Privacy, pp. 238–252. IEEE Computer Society (2013)
45. Rotaru, D., Smart, N.P., Stam, M.: Modes of operation suitable for computing on encrypted data. *IACR Trans. Symmetric Cryptol.* **2017**(3), 294–324 (2017)

46. Todo, Y.: Structural evaluation by generalized integral property. In: Oswald, E., Fischlin, M. (eds.) EUROCRYPT 2015. LNCS, vol. 9056, pp. 287–314. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-46800-5_12
47. Todo, Y., Isobe, T., Hao, Y., Meier, W.: Cube attacks on non-blackbox polynomials based on division property. In: Katz, J., Shacham, H. (eds.) CRYPTO 2017. LNCS, vol. 10403, pp. 250–279. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-63697-9_9
48. Todo, Y., Morii, M.: Bit-based division property and application to SIMON family. In: Peyrin, T. (ed.) FSE 2016. LNCS, vol. 9783, pp. 357–377. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-52993-5_18
49. Vielhaber, M.: Breaking ONE.FIVIUM by AIDA an algebraic IV differential attack. IACR Cryptology ePrint Archive 2007, 413 (2007)
50. Wang, Q., Grassi, L., Rechberger, C.: Zero-sum partitions of PHOTON permutations. In: Smart, N.P. (ed.) CT-RSA 2018. LNCS, vol. 10808, pp. 279–299. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-76953-0_15
51. Wang, Q., Hao, Y., Todo, Y., Li, C., Isobe, T., Meier, W.: Improved division property based cube attacks exploiting algebraic properties of superpoly. In: Shacham, H., Boldyreva, A. (eds.) CRYPTO 2018. LNCS, vol. 10991, pp. 275–305. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-96884-1_10
52. Wang, Q., Liu, Z., Varici, K., Sasaki, Yu., Rijmen, V., Todo, Y.: Cryptanalysis of reduced-round SIMON32 and SIMON48. In: Meier, W., Mukhopadhyay, D. (eds.) INDOCRYPT 2014. LNCS, vol. 8885, pp. 143–160. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-13039-2_9
53. Wang, S., Hu, B., Guan, J., Zhang, K., Shi, T.: MILP-aided method of searching division property using three subsets and applications. In: Galbraith, S.D., Moriai, S. (eds.) ASIACRYPT 2019. LNCS, vol. 11923, pp. 398–427. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-34618-8_14