Beyond Linear Subspace Trails: Nonlinear Subspaces for Gröbner Basis Attacks on Poseidon/Poseidon2 and Neptune

Enyan Li[1], Fukang Liu[2], and Gaoli Wang[1]

[1] Software Engineering Institute, East China Normal University, Shanghai, China
lienyan6@gmail.com glwang@sei.ecnu.edu.cn
[2] Institute of Science Tokyo, Tokyo, Japan
liufukangs@gmail.com

Abstract. Poseidon/Poseidon2 and Neptune are prominent primitives for zero-knowledge proof systems. Their arithmetic circuit cost is reduced mainly through partial S-box layers and low-degree finite field operations. Algebraic attacks are therefore a central part of their security analysis, and Gröbner basis methods are a main tool for studying such attacks. For such attacks, controlling the algebraic degree of the polynomial systems induced by partial rounds is a central issue. Previous work has shown that linear subspace trails can reduce the algebraic degree of partial rounds in constrained-input-constrained-output (CICO) problems. Therefore, subspace analysis has become an important tool for evaluating the algebraic security of Poseidon-like permutations.

The main contribution of this paper is to extend the existing linear subspace trail framework to nonlinear subspaces. More precisely, we first introduce a parametric Macaulay matrix method. This method transforms the search for algebraic constraints that reduce degree growth into the problem of solving a parametric system. It provides a general algebraic approach for constructing longer nonlinear subspace trails that suppress degree growth over more internal partial rounds. Second, for the CICO problem with Ec extra constraints, we give a concrete constraint pattern that extends a linear subspace trail into a nonlinear one. In this nonlinear construction, the first Ec subspace constraints generate an ideal, and further compatible subspace constraints can be added along the chain without enlarging this ideal. As a result, the nonlinear subspace trail can cover up to $2Ec$ internal rounds, whereas the previous linear subspace trail can cover up to Ec rounds. We further show that the balancing matrix required by this construction is generically nonsingular. Furthermore, we propose subspace modeling variants without variable substitution. These variants impose linear or nonlinear constraints directly on high-degree intermediate states.

For the Poseidon/Poseidon2 and Neptune instances proposed by Grassi et al. in ToSC 2025, our experiments show that, under the same complexity bound and the same Gröbner basis cost model, the nonlinear subspace model can analyze approximately twice as many internal partial rounds as the linear subspace model considered in ToSC 2025. For

several concrete instances, our method reaches or even exceeds the recommended number of internal rounds given by the designers in sponge mode or compression mode.

Keywords: Poseidon/Poseidon2, Neptune, Algebraic cryptanalysis, Subspace trails, Gröbner basis

1 Introduction

1.1 Background and Motivation

With the rapid development of zero-knowledge proofs, verifiable computation, and privacy-preserving protocols, arithmetization-oriented hash functions have become important building blocks in these settings. Traditional hash functions are usually designed for bit operations and efficient execution on general purpose processors. However, in succinct noninteractive arguments of knowledge (SNARKs), scalable transparent arguments of knowledge (STARKs), and related proof systems, the cost of a hash function is often dominated by finite field multiplications and nonlinear constraints. Therefore, hash functions that can be implemented over large finite fields with low algebraic degree and low multiplication complexity are useful for reducing the cost of arithmetic circuits in proof systems. In particular, the Ethereum community has considered potential applications of Poseidon and has promoted more systematic cryptanalytic studies of Poseidon-like designs [25].

Poseidon is one of the most representative designs in this line of work. It has received substantial attention and has been used in zero-knowledge protocols, decentralized computation and storage, and privacy-preserving tools [21]. Later, Poseidon2 further improved the performance by redesigning the linear layer [24]. Neptune follows the same Poseidon-like design philosophy, but introduces new nonlinear layers for the external rounds in order to further reduce the multiplication complexity [23].

The efficiency of Poseidon-like permutations largely comes from an asymmetric round structure inherited from HADES [20]. Instead of applying the nonlinear layer uniformly throughout the permutation, the construction concentrates full S-box layers near the two ends and uses cheaper partial rounds in the middle. During these internal rounds, nonlinear operations are applied to only a small part of the state, while the linear layer propagates their influence to the remaining coordinates. This organization substantially reduces the arithmetic cost of the permutation and is therefore well suited to proof-system applications.

For traditional symmetric primitives such as AES, Keccak, and SHA-3, statistical attacks have long played a central role. In contrast, for arithmetization-oriented primitives such as Poseidon, Poseidon2, and Neptune, the round functions have simple polynomial descriptions over finite fields. As a result, algebraic attacks must be carefully considered when evaluating their security margins. In recent years, Gröbner basis attacks, resultant attacks, univariate solving methods, and related algebraic attacks against arithmetization-oriented hash

functions and permutations have developed rapidly. They have become popular tools for understanding the security of these designs [5,7,25,27,34].

The partial-round structure also creates useful algebraic modeling opportunities for an attacker. By imposing suitable affine relations on an intermediate state, the inputs to the active S-boxes can be restricted so that several successive internal rounds contribute no additional degree, or substantially less degree, to the attack system [22,25,29]. Exploiting such intermediate-state restrictions is therefore an important technique in current algebraic analyses of Poseidon/Poseidon2 and Neptune.

1.2 Contributions

The main contributions of this paper are summarized as follows.

1. We introduce parametric Macaulay matrices as a general tool for constructing subspace trails. The established subspace constraints generate an ideal, and the central question is how the extra constraints should be constructed so that more internal round constraints lie in this ideal without enlarging it. We transform this question into a parameter solving problem. The parameter solutions yield subspace constraint ideals that extend the length of the corresponding subspace trails. Consequently, the algebraic model can pass through more internal partial rounds without increasing its degree. Thus, parametric Macaulay matrices provide a new viewpoint on the CICO problem for ZK-friendly hash functions. They shift the focus from solving the attack polynomial system to solving an auxiliary parameter system.
2. We give a concrete nonlinear constraint pattern in Sect. 3.3, together with an algebraic analysis of the resulting subspace trails. Suppose that each internal partial round activates s S-boxes. Following [29], let Ec denote the number of extra constraints that the attacker can impose while keeping the resulting ideal system solvable and zero-dimensional. Then our construction extends a linear subspace trail of length about $\lfloor Ec/s \rfloor$ to a nonlinear subspace trail of length $2\lfloor Ec/s \rfloor$. We further prove that the balancing matrix arising from this construction is generically nonsingular. In particular, when $s = 1$, the length is extended from Ec to $2Ec$. The complexity estimates in Sect. 4.2 show that, under the same complexity bound and the same Gröbner basis solving model, the nonlinear subspace model can analyze approximately twice as many internal rounds of Poseidon/Poseidon2 and Neptune as the traditional linear subspace model. For the Poseidon/Poseidon2 instance with $t = 24$, $\alpha = 7$, $c = d = 1$, and $p \approx 2^{256}$, the number of internal rounds analyzed by our model exceeds the designers' recommended round number, and is slightly below the recommendation after the security margin of internal rounds is applied.
3. We compare several attack models in a systematic way. In Sect. 4.1, we also consider subspace modeling variants without variable substitution, where linear or nonlinear constraints are imposed directly on a high-degree intermediate state. Tables 1 and 2 summarize the competitive modeling strategies for Poseidon/Poseidon2 and Neptune. The experiments in Sect. 4.2

Table 1: Overview of algebraic attack models for Poseidon2 in sponge mode. Here α is the S-box degree, $r_F = 2r_f$ is the number of full rounds, r_P is the number of partial rounds, t is the state size, c is the capacity, d is the digest size, τ is the number of partial rounds before entering the subspace, and Ec is the number of extra constraints available beyond the output constraints. Analogous results hold for Poseidon, for which the initial full S-Box layer can be skipped for free [7], and for Poseidon2/Poseidon in compression mode

Modeling strategy	# equations	Degree	References and notes
Basic attack	d	$\alpha^{r_F+r_P}$	Designers' analysis [21]
Forward+substitution+linear subspace	t	$\alpha^{r_f+\tau}$	Designers' analysis revisited [25]
	d	$\alpha^{r_f+r_P-\tau-Ec}$	
Forward+substitution+nonlinear subspace	t	$\alpha^{r_f+\tau}$	Model 3 in Sect. 4.1
	Ec	α	
	d	$\alpha^{r_f+r_P-\tau-2Ec}$	
Forward+nonsubstitution+linear subspace	Ec	α^{r_f}	Model 4 in Sect. 4.1
	d	$\alpha^{r_F+r_P-Ec}$	
Forward+nonsubstitution+nonlinear subspace	Ec	α^{r_f+1}	Model 5 in Sect. 4.1
	d	$\alpha^{r_F+r_P-2Ec}$	

show that the analyzable round number is mainly determined by the solving method and by whether linear or nonlinear subspaces are used. We also identify three compression mode parameter sets where nonlinear subspace modeling reaches or exceeds the designers' recommended number of rounds.

1.3 Related Work

subspace trail constructions in Poseidon and Neptune. The security relevance of subspaces in HADES-like permutations has been recognized since the early analyses of their partial-round structure. The central issue is the interaction between the linear layer and the sparsely activated nonlinear coordinates: unfavorable linear layers may preserve directions that keep the active S-box inputs constrained across multiple rounds [10,26]. This observation motivated design criteria aimed at ruling out trails that persist indefinitely [22]. For cryptanalysis, however, even a finite sequence of constrained partial rounds can be useful because it lowers the algebraic complexity of the resulting model. Such finite trails already appear in the original security analysis of Poseidon [21], and their concrete impact on Poseidon/Poseidon2 and Neptune was later quantified more precisely in [25].

Rigorous algebraic analysis of linear subspace trails. The linear-subspace framework used in this paper relies on the characterization developed in [29]. That work separates two limitations that had previously been discussed somewhat independently: the structural length permitted by the internal linear layer and the number of constraints that an attacker can actually spend on realizing such a trail. For an internal round activating s S-boxes, it derives the finite-trail

Table 2: Overview of algebraic attack models for NEPTUNE in sponge mode. Here α is the internal S-box degree, r_f and r'_f are the numbers of initial and final full rounds, $r_F = r_f + r'_f$, r_P is the number of partial rounds, t is the state size, c is the capacity, d is the digest size, τ is the number of partial rounds before entering the subspace, and Ec is the number of extra constraints available beyond the output constraints.

Modeling strategy	# equations	Degree	References and notes
Basic attack	d	$4^{r_F} \cdot \alpha^{r_P}$	Designers' analysis [23]
Forward+substitution+linear subspace	t	$4^{r_f} \cdot \alpha^{\tau}$	Designers' analysis revisited [25]
	d	$4^{r'_f} \cdot \alpha^{r_P-\tau-Ec}$	
Forward+substitution+nonlinear subspace	t	$4^{r_f} \cdot \alpha^{\tau}$	Model 3 in Sect. 4.1
	Ec	α	
	d	$4^{r'_f} \cdot \alpha^{r_P-\tau-2Ec}$	
Forward+nonsubstitution+linear subspace	Ec	4^{r_f}	Model 4 in Sect. 4.1
	d	$4^{r_F} \cdot \alpha^{r_P-Ec}$	
Forward+nonsubstitution+nonlinear subspace	Ec	$4^{r_f}\alpha$	Model 5 in Sect. 4.1
	d	$4^{r_F} \cdot \alpha^{r_P-2Ec}$	
SitM+linear subspace	c	4^{r_f}	Designers' analysis revisited [25]
	d	$4^{r'_f} \cdot \alpha^{r_P-Ec}$	
SitM+nonlinear subspace	c	4^{r_f}	Model 7 in Sect. 4.1
	Ec	α	
	d	$4^{r'_f} \cdot \alpha^{r_P-2Ec}$	

bound $\lceil t/s \rceil - 1$ once infinite invariant subspaces are excluded. It also refines the CICO formulation for sponge preimage attacks and obtains the usable extra-constraint budget $Ec = r - \min\{c, d\}$. Consequently, the realizable linearized path may be limited by the available CICO degrees of freedom even when a longer finite trail is structurally permitted. We take these bounds as the baseline and investigate whether nonlinear constraints can make more effective use of the same budget.

Round skipping and degree reduction. Besides the subspace trail approach, several works have studied the algebraic security of POSEIDON and its variants from the perspective of round skipping and degree reduction. Bariant et al. studied algebraic attacks on POSEIDON in the context of Ethereum Foundation challenges. They showed that some full rounds can be skipped in the algebraic model. These issues were later corrected or mitigated in the designs of POSEIDON2 and NEPTUNE [24,23,7]. Nonetheless, Merz and Rodríguez García showed that the concrete linear layers and operation modes of POSEIDON2 and POSEIDON2B can be exploited to strengthen round skipping attacks, both for CICO modeling and for direct algebraic preimage modeling [30]. More recently, Sanso and Vitto introduced degree annihilation, which reduces the effective algebraic degree by imposing constraints that cancel dominant degree terms without removing the corresponding rounds, and can be combined with classical round-skipping techniques [32]. Vitto subsequently studied the reachability of finite subspace trails through the initial full rounds of POSEIDON, introducing full round absorption

and constructing constrained input families and MDS linear layers for which the full round prefix acts as a controlled reparametrization into a finite trail [33]. These techniques are complementary to the approach developed in this paper. Our construction focuses on bypassing internal partial rounds more effectively. Combining the two directions may therefore be a promising topic for future work.

Other algebraic modeling and solving techniques. Faugère and Perret gave closed-form expressions for the degree of the univariate polynomial appearing in the lexicographic Gröbner basis [18]. Ashur, Buschman, and Mahzoun revisited Gröbner basis attacks based on round-level modeling. Their experiments suggested that the security estimates for some non-standard security level parameters may be optimistic [1]. Bak et al. further demonstrated the practical power of resultant based algebraic attacks by solving round-reduced CICO instances of Poseidon and Poseidon2 from the Ethereum Foundation bounty program [3].

Recent years have seen several more algebraic methods for arithmetization-oriented primitives. The FreeLunch and CheapLunch methods exploit favorable monomial orderings to simplify algebraic solving. FreeLunch applies this idea to certain round-level CICO-1 models and gives efficient attacks on full round instances of Anemoi, Arion, and Griffin, while CheapLunch extends it to CICO-k through weighted models and ideal-degree bounds, with applications to primitives such as Poseidon, Neptune, and XHash8 [5,2]. In both methods, the main cost of Gröbner basis solving is shifted to the phase of change of ordering or FGLM step. Similarly, recent iterated resultant methods show that resultant computation or univariate solving can be more suitable than multivariate Gröbner basis modeling when the number of variables is small [6,34]. In another direction, Sanso and Vitto applied Graeffe-based root finding to Poseidon and Poseidon2 over NTT-friendly fields [31]. Zhao and Ding subsequently developed a related Graeffe transform approach that accelerated root finding for high-degree univariate polynomials over finite fields and significantly improved the practical efficiency of interpolation attacks in the Poseidon challenges [35]. The *Six Worlds of Gröbner Basis Cryptanalysis* framework by Koschatko et al. provides a more systematic way to compare the costs of different stages in Gröbner basis attacks [27]. We use insights from these works when estimating the cost of univariate solving, resultant solving, and Gröbner basis solving.

1.4 Organization

The rest of this paper is organized as follows. Section 2 gives the necessary preliminaries. It reviews Poseidon/Poseidon2 and Neptune, the traditional CICO problem, and the complexity models for univariate solving, resultant computation, and Gröbner basis computation used later in the paper. Section 3 introduces the parametric Macaulay matrix method and develops the nonlinear subspace construction. Section 4 applies these constructions to concrete attack models and compares the complexity of different modeling strategies for Poseidon/Poseidon2 and Neptune parameters. Finally, Sect. 5 concludes the paper.

2 Preliminaries

2.1 Structures of Poseidon and Neptune

Poseidon, Poseidon2, and Neptune are arithmetization-oriented hash functions over finite fields. Poseidon was originally specified in sponge mode, and Neptune was also proposed mainly for sponge hashing [21,23]. Poseidon2 supports both sponge mode and compression mode. However, the only formally specified compression mode currently available is the small state instance with $t = 3$ [24,25].

In sponge mode, if the underlying permutation is modeled as random, then the sponge construction is indistinguishable from a random oracle up to about $p^{c/2}$ queries [8]. Thus, κ bit security requires $p^{c/2} \geq 2^{\kappa}$. Birthday resistance for a digest of d field elements similarly requires $p^{d/2} \geq 2^{\kappa}$. Here c is the capacity, and d is the digest length over $\mathbb{F}_p$. Sponge security is also closely related to the hardness of the CICO problem for the underlying permutation [9], which is recalled in the next subsection.

All three permutations are defined over a finite field $\mathbb{F}_p$. We denote the state width by t, and write the state as $x = (x_0, x_1, \ldots, x_{t-1}) \in \mathbb{F}_p^t$. Following the unified notation of Grassi et al. [25], the permutations of Poseidon/Poseidon2 and Neptune can be written as

$$P(x) = E_{r_f + r'_f - 1} \circ \cdots \circ E_{r_f} \circ I_{r_P - 1} \circ \cdots \circ I_0 \circ E_{r_f - 1} \circ \cdots \circ E_0 \circ M'(x),$$

where M' is the initial linear layer, E_i denotes an external full round, and I_i denotes an internal partial round. The total number of external full rounds is $r_F = r_f + r'_f$. Here r_f full rounds are placed before the internal partial rounds, and r'_f full rounds are placed after them. The number of internal partial rounds is denoted by r_P.

For Poseidon and Poseidon2, an external full round $E : \mathbb{F}_p^t \to \mathbb{F}_p^t$ has the form

$$E(x) = M_E \cdot \left((x_0 + c_0)^{\alpha}, (x_1 + c_1)^{\alpha}, \ldots, (x_{t-1} + c_{t-1})^{\alpha}\right)^T,$$

where M_E is the external linear layer, c_i are round constants, and α is the degree of the power S-box. Usually α is chosen as a small positive integer satisfying $\gcd(\alpha, p - 1) = 1$. An internal partial round $I : \mathbb{F}_p^t \to \mathbb{F}_p^t$ generally applies power S-boxes only to the first s coordinates. For Poseidon, Poseidon2, and Neptune, the standard setting is $s = 1$. Thus, only the first coordinate passes through an S-box, and the internal partial round has the form

$$I(x) = M_I \cdot \left((x_0 + c'_0)^{\alpha}, x_1 + c'_1, \ldots, x_{t-1} + c'_{t-1}\right)^T,$$

where M_I is the internal linear layer, and c'_i is a round constant.

The main difference between Poseidon and Poseidon2 lies in the initial linear layer and in the choice of linear layers. In Poseidon, M' is the identity matrix, and the same MDS matrix M is used for both external and internal linear layers. In concrete instantiations, the designers recommend using Cauchy

matrices to construct the MDS linear layer [21]. This matrix must also satisfy conditions that prevent infinite subspace trails. These conditions are related to the minimal polynomials of $M, M^2, \ldots$ [22].

In Poseidon2, one sets $M' = M_E$, and designs M_E and M_I separately. For small state widths, M_E is an MDS matrix. For larger state widths, it satisfies the required branch number conditions. The internal linear layer M_I must still satisfy the conditions that prevent infinite subspace trails [24].

Neptune is a variant of Poseidon2. It is defined for even state widths $t = 2t' \leq 24$. It follows the same overall round structure. Its initial linear layer M' and internal partial rounds I have the same form as in Poseidon2. The difference is the nonlinear layer in the external full rounds [23,25]. In Neptune, the state is grouped into pairs, and a bivariate nonlinear map $S : \mathbb{F}_p^2 \to \mathbb{F}_p^2$ is applied to each pair. Therefore, an external full round of Neptune can be written as

$$E(x) = M_E \cdot \big(S(x_0 + c_0, x_1 + c_1),\, S(x_2 + c_2, x_3 + c_3),\, \ldots,\, S(x_{t-2} + c_{t-2}, x_{t-1} + c_{t-1})\big)^T.$$

The map S is based on a Lai–Massey structure [28], and satisfies $\deg(S) = \deg(S^{-1}) = 4$.

We now briefly recall the round number recommendations given by the designers of Poseidon and Neptune. Section 4.2 uses them as baselines for complexity comparisons. For Poseidon, at a security level of κ bits, under the condition $2^{80} \leq 2^\kappa \leq \min\{2^{256}, p^{t/3}\}$, the designers recommend $r_F = 6$, namely $r_f = 3$ and $r'_f = 3$, and set

$$r_P = \left\lceil \max\left\{ 1 + \frac{\min\{\kappa, \log_2(p)\}}{\log_2(\alpha)} + \log_\alpha(t) - r_F,\, r_{\mathrm{GB}} \right\} \right\rceil.$$

Here α is the S-box degree. The first term corresponds to interpolation attacks, and the second term comes from Gröbner basis attacks. More precisely,

$$r_{\mathrm{GB}} = \max\left\{ \tfrac{\min\{\log_2(p),\kappa\}}{\log_2(\alpha)} - r_F,\, t - 1 + \log_\alpha(2) \cdot \min\left\{ \tfrac{\kappa}{t+1},\, \tfrac{\log_2(p)}{2} \right\} - r_F,\, r'_{\mathrm{GB}} \right\}.$$

The second term corresponds to Gröbner basis attacks using linear subspace trails, while r'_{GB} corresponds to the round-level estimate in [1]. The practical recommendation also adds a security margin: two extra external full rounds and 7.5% extra internal partial rounds. Thus, $r_F = 6$ is increased to $r_F = 8$.

For Neptune, we consider the same security level condition $2^{80} \leq 2^\kappa \leq \min\{2^{256}, p^{t/3}\}$. Under this condition, the designers recommend $r_F = 6$, namely $r_f = 4$ and $r'_f = 2$, and set

$$r_P = \left\lceil \max\left\{ \frac{\min\{\kappa, \log_2(p)\} - 6}{\log_2(\alpha)} + 3 + t + \log_\alpha(t),\, \frac{\kappa}{2} - 4 \cdot t - 2 \right\} \right\rceil.$$

Here α is the degree of the internal power S-box. The first term is used to resist interpolation attacks. It also includes t additional internal rounds because an invariant subspace trail may cover t internal rounds [23]. The security margin of Neptune is 12.5% extra internal partial rounds, with no additional external full rounds.

2.2 The CICO Problem

In the algebraic security analysis of POSEIDON-like permutations, a basic problem is the constrained-input-constrained-output (CICO) problem. Let $P : \mathbb{F}_p^t \to \mathbb{F}_p^t$ be an invertible permutation. Given two integers n_I and n_O, where $0 < n_I, n_O < t, n_I + n_O \leq t$, the traditional CICO$(n_I, n_O)$ problem asks the attacker to find an input state $x \in \mathbb{F}_p^t$ such that n_I coordinates of x are fixed, and n_O coordinates of the output state $y = P(x)$ are fixed. The fixed values can usually be taken as 0, since replacing them by other constants does not affect the variable count, equation count, or degree growth considered below [25,9].

In this formulation, the CICO(n_I, n_O) problem gives a common abstraction for analyzing preimage attacks on sponge hash functions and attacks on compression functions. If P is viewed as a random permutation, then the hardness of the CICO problem is often used to characterize the complexity of constrained state search in the sponge mode.

As shown in [29], the CICO problem can be refined according to the type of output constraint. Preimage attacks with digest constraints are described as constrained input constrained digest (CICD), while CICO problems with capacity constraints are described as constrained input constrained capacity (CICC). In sponge mode with additional message absorption, CICD can be transformed into CICC. This gives an extra constraint budget $Ec = r - \min\{c, d\}$ for a general preimage attack, under the requirement that the resulting ideal remains zero-dimensional and has solutions over the base field. In compression mode, the corresponding budget is $Ec = t - d$.

The same work also points out that the length of a linear subspace trail is limited by two quantities: the structural trail bound $\lceil t/s \rceil - 1$, and the available constraint budget $\lfloor Ec/s \rfloor$. In the standard POSEIDON/POSEIDON2 and NEPTUNE settings, $\lfloor Ec/s \rfloor \leq \lceil t/s \rceil - 1$, so the effective bound in preimage attacks is the constraint budget $\lfloor Ec/s \rfloor$.

2.3 Basic Arithmetic Complexity Models

Firstly, for a general univariate polynomial equation F over $\mathbb{F}_p$, let $\deg(F) = \delta$. The solving complexity is [7]

$$\mathcal{C}_{\text{UNI}}(\delta) = O(\delta \log(\delta)(\log(\delta) + \log(p)) \log\log(\delta)) .$$

However, Ding et al. proposed an acceleration method for solving such polynomial equations over finite fields. It reduces the complexity to [35]

$$\mathcal{C}_{\text{Graeffe}}(\delta) = O(\delta \log(\delta)(\log(p) - \log(\delta) + 1) \log\log(\delta)) .$$

In the later computations, we use this optimized complexity formula.

Next, for a general bivariate polynomial system $\{F, G\} \subset \mathbb{F}_p[x, y]$, let δ_x and δ_y denote the maximum degrees of the two equations in the x- and y-directions, respectively. Without loss of generality, assume $\delta_x \geq \delta_y$. The bottleneck of resultant-based solving is the computation of the resultant [12,34]. Its

complexity is [19]

$$\mathcal{C}_{\mathrm{RES}}(\delta_x, \delta_y) = O(\delta_y \mathcal{M}(\delta_x \delta_y) \log(\delta_x \delta_y)) \,,$$

where $\mathcal{M}$ denotes the arithmetic complexity of multiplying two polynomials. By the Cantor–Kaltofen algorithm [11], one has $\mathcal{M}(\delta) = O(\delta \log(\delta) \log\log(\delta))$.

Finally, we recall the Gröbner basis complexity model used later in the paper. For zero-dimensional ideals in $\mathbb{F}_p[x_1, \ldots, x_n]$ with $n \geq 3$, a Gröbner basis attack usually consists of two stages. The first stage computes a Gröbner basis under grevlex, or another monomial ordering suitable for computation, typically using F4/F5 algorithms. Suppose that the system consists of n variables and n equations $F_1, \ldots, F_n$, and let $\delta_i = \deg(F_i)$. Under the semi-regularity assumption [4,17], the degree of regularity is estimated by the Macaulay bound $d_{\mathrm{reg}} = 1 + \sum_{i=1}^{n} (\delta_i - 1)$. Hence, the complexity of the F4/F5 stage is estimated as [15,16,4,17]

$$\mathcal{C}_{\mathrm{GB}} = O\left(\binom{d_{\mathrm{reg}} + n}{n}^{w} \right) = O\left(\binom{\sum_{i=1}^{n} \delta_i + 1}{n}^{w} \right),$$

where $2 \leq w < 3$. Structured polynomial systems from symmetric cryptography often do not satisfy the regularity assumption [1,13,7,27]. Nevertheless, the regular system estimate is still commonly used as a loose upper bound in design analyses, usually together with extra security margins to compensate for possible non-regular behavior [25].

The second stage of Gröbner basis computation changes the monomial ordering, usually through the FGLM algorithm, in order to obtain a lexicographic Gröbner basis. This makes it possible to recover a univariate representation and enumerate the solutions [14,27]. Let $D_I = \dim_{\mathbb{F}_p}(\mathbb{F}_p[x_1, \ldots, x_n]/I)$ be the vector space dimension of the quotient ring of the zero-dimensional ideal I. The classical FGLM complexity is $\mathcal{C}_{\mathrm{FGLM}} = O(nD_I^3)$, and with fast linear algebra it is often written as $\mathcal{C}_{\mathrm{FGLM}} = O(nD_I^{\omega})$. In general, the Bézout bound $D_I \leq \prod_{i=1}^{n} \delta_i$ gives a conservative estimate. In complexity estimates for subspace attacks on Poseidon and Neptune, the F4/F5 stage is usually treated as the main bottleneck, while the FGLM stage is regarded as secondary [25,27]. We follow this convention in the rest of the paper, and when necessary we also check whether the FGLM stage changes the round number conclusions.

3 Constructing Nonlinear Subspaces for Internal Rounds

As recalled in Sect. 2.2, the work [29] shows that, for preimage attacks on Poseidon, Poseidon2, and Neptune in sponge mode, the effective budget for imposing extra subspace constraints is $Ec = r - \min\{c, d\}$. In compression mode, the corresponding budget is $Ec = t - d$. In more general cases, each internal partial round activates s S-boxes, linear subspace constraints can therefore support a trail of length about $\lfloor Ec/s \rfloor$. This is the starting point of our nonlinear construction.

In this paper, we allow nonlinear constraints on intermediate states. The attacker no longer requires an intermediate state to lie only in a linear subspace. Instead, several low-degree nonlinear relations may constrain the state simultaneously. The discussion in this section is given for Poseidon, but it also applies to Poseidon2 and Neptune. The construction uses only the partial round structure, and does not depend on the concrete choice of the internal round linear layer. Under the same budget of Ec extra constraints, if each internal partial round activates s S-boxes, a linear subspace trail of length $\lfloor Ec/s \rfloor$ can be extended to a nonlinear subspace trail of length $2\lfloor Ec/s \rfloor$. In particular, when $s = 1$, the length is extended from Ec to $2Ec$.

This section is organized as follows. Section 3.1 gives a concrete example and shows how to use a parametric Macaulay matrix to find such an extended subspace trail. Section 3.2 gives an algebraic explanation of the observed phenomenon, and extracts the key conditions needed for the later construction. Section 3.3 then gives a general nonlinear subspace trail construction used in our attack models.

3.1 An Example of Introducing Nonlinear Constraints via Parametric Macaulay Matrices

We consider the following Poseidon instance over the KoalaBear field: $\alpha = 3$, $p = 2130706433 = 2^{31} - 2^{24} + 1$, $t = 5$, $s = 1$, $r = 4$, $c = 1$, and $d = 1$. The polynomial degrees appearing in this example remain far below $p - 2$. By the CICO discussion in Sect. 2.2, the corresponding output constraint and extra constraint budget are $\mathrm{Oc} = 1$ and $Ec = 3$. Therefore, the linear subspace technique can in general introduce only three extra equations to build a linear subspace trail.

We set the starting state of the subspace trail in partial rounds as

$$\mathrm{state}_0 = \left(t_0,\, t_1x_0 + t_2,\, t_3x_1 + t_4,\, t_5x_2 + t_6,\, t_7x_3 + t_8\right)^{\mathsf{T}}.$$

Here we introduce a parameter family $T = \{t_0, \ldots, t_8\}$ in the initial state. These parameters will later be solved through an auxiliary parameter system. The remaining free variables are denoted by $X = \{x_0, x_1, x_2, x_3\}$, and they form the main variables of the attack polynomial system.

The first coordinate is set to the parameter t_0, which gives the first linear constraint f_0. Thus, the first partial round S-box does not increase the degree in X. The remaining coordinates are affine linear forms in X. Hence, the initial state is still linear with respect to the main variables.

Next, we let this state propagate through one partial round. Since X does not appear in the first coordinate, this partial round keeps the degree in the main variables equal to one, while the degree in the parameter variables becomes three. The state becomes

$$\mathrm{state}_1 = \left(g_0^{(1)}(X,T),\, g_1^{(1)}(X,T),\, g_2^{(1)}(X,T),\, g_3^{(1)}(X,T),\, g_4^{(1)}(X,T)\right)^{\mathsf{T}}.$$

Here $g_i^{(1)}(X,T)$ is a polynomial in $\mathbb{F}_p[X,T]$, where $i = 0,1,2,3,4$. To track the degree growth with respect to different variables, we write $\deg_{x_j}(g_i^{(1)})$, $\deg_{t_j}(g_i^{(1)})$, and $\deg_T(g_i^{(1)})$ for the degree in a single main variable, in a single parameter variable, and in the whole parameter family, respectively. We extend the degree function to state vectors coordinatewise, and define $\deg_X(\text{state}_1) = \max_i \deg_X(g_i^{(1)})$. Then $\deg_T(\text{state}_1) = 3$. More precisely, the degree in t_0 is 3, while the degrees in $t_1, \ldots, t_8$ are 1. Moreover, $\deg_X(\text{state}_1) = 1$.

For state_1, we can impose a constraint on the first coordinate polynomial, so that the state degree does not rise to 3. Since $\deg_X(g_0^{(1)}(X,T)) = 1$, we can only reduce this coordinate to degree 0. Thus, we set $g_0^{(1)}(X,T) = t_9$. This is the second extra constraint after f_0, and we denote it by f_1. Then we set $\text{state}_1 = (t_9, g_1^{(1)}, \ldots, g_4^{(1)})^{\mathsf{T}}$, and let it pass through the next partial round to obtain state_2. At this point, $\deg_T(\text{state}_2) = \deg_{t_0}(\text{state}_2) = 3$, and $\deg_X(\text{state}_2) = 1$.

For state_2, which is still linear in X, we do something different from the previous approach. We no longer impose another linear constraint. Instead, we let it pass freely through the third partial round and reach state_3. Then $\deg_T(\text{state}_3) = 9$ and $\deg_X(\text{state}_3) = 3$. According to the previous linearization strategy, the process would stop here. However, if the goal is to block the growth of equation degrees, then we can still impose a constraint on the first coordinate, even though $\deg_X$ has already reached 3. We can force it to degenerate to degree 0, 1, or 2. Clearly, degeneration to degree 1 contains degeneration to degree 0 as a special case. Degeneration to degree 2 would make the next round state have degree 6. We will discuss this type of construction later.

Here we impose the constraint $g_0^{(3)}(X,T) = t_{10}x_0 + t_{11}x_1 + t_{12}x_2 + t_{13}x_3 + t_{14}$. This is the third extra constraint, denoted by f_2. Although $\text{Ec} = 3$ in this instance, the parameter system still has remaining degrees of freedom after the first three constraints. We use these degrees of freedom to introduce one more constraint, so that the degree growth caused by the next S-box is also blocked. After the constrained state_3 passes through the fourth partial round, we obtain state_4. Then $\deg_T(\text{state}_4) = 9$ and $\deg_X(\text{state}_4) = 3$. We further impose the fourth constraint $g_0^{(4)}(X,T) = t_{15}x_0 + t_{16}x_1 + t_{17}x_2 + t_{18}x_3 + t_{19}$, and denote it by f_3. This step introduces five new parameter variables. It reduces the first coordinate polynomial to degree one in X, so that the next active S-box does not increase the degree of the main system.

We have now obtained the ideal $\langle f_1, f_2, f_3 \rangle$ generated by the subspace constraints. If the parameters can be chosen so that $\langle f_1, f_2, f_3 \rangle = \langle f_1, f_2 \rangle$, equivalently $f_3 \in \langle f_1, f_2 \rangle$, then adding f_3 does not change the dimension of the ideal. We still have $\dim V(\langle f_1, f_2, f_3 \rangle) = 2$. After the final output constraint at the digest or capacity position is added, the complete ideal becomes zero-dimensional.

We use a parametric Macaulay matrix to construct such a degenerating subideal. We take the grevlex order $x_3 > x_2 > x_1 > x_0$. The columns of the matrix are monomials of degree at most three. The rows are given by elements of $\langle f_1, f_2, f_3 \rangle$ of degree at most three. Since $\deg_X(f_2) = \deg_X(f_3) = 3$, the polynomials f_2 and f_3 are directly inserted as the last two rows. Since $\deg_X(f_1) = 1$,

the first 15 rows are obtained by multiplying f_1 by all monomials of degree at most two. The resulting matrix is shown in Fig. 1. It has 17 rows and 35 columns.

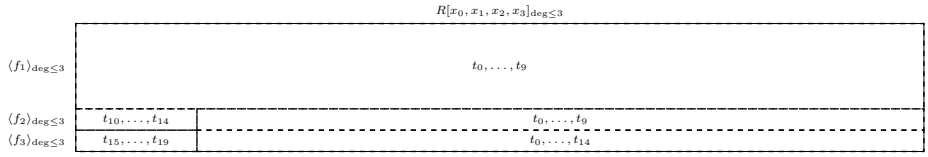

Fig. 1: Structure of the parametric Macaulay matrix truncated at degree 3.

Let this parametric Macaulay matrix be M_T. Then $(M_T)_{f_2,\{1,\ldots,x_3\}}$ denotes the submatrix obtained by taking the row of f_2 and the first five columns. During the construction of the matrix, especially when f_1 is lifted as a generator, the coefficient columns only undergo some permutations. At this point, $f_3 \in \langle f_1, f_2 \rangle$ is equivalent to the last row of M_T lying in the linear span of the preceding rows. Further, this is equivalent to the existence of a nontrivial coefficient vector $c = (c_0, c_1, \ldots, c_{15}, -1)$, where $C = \{c_0, \ldots, c_{15}\}$, such that $c \cdot M_T = 0$. This system gives a nontrivial linear relation involving the last row, and hence shows that $f_3 \in \langle f_1, f_2 \rangle$. The existence of a coefficient vector C satisfying this relation is further discussed in Sect. 3.2.

In Fig. 1, the parameters $t_{15}, \ldots, t_{19}$ appear only in the first five columns of the f_3 row separately. Therefore, we can first handle the last 30 columns. After randomly fixing $t_0, \ldots, t_9$, the first 16 rows restricted to the last 30 columns become a constant block, while the f_3 row contains only $t_{10}, \ldots, t_{14}$ and has degree at most 3. The compatibility conditions are obtained by right multiplying the last 30 column block of M_T by a basis of the right kernel of its constant upper block. The upper rows vanish, while the f_3 row yields one parameter equation for each kernel vector.

Computation shows that the rank of this constant block and the dimension of its kernel are both 15. This is because the coefficients of f_1 do not appear in these nonlinear columns, and the corresponding rows are zero there. Hence, $t_{10}, \ldots, t_{14}$ have to satisfy 15 cubic compatibility equations. A Sage computation shows that the coordinate ring of this five-variable parameter ideal has Krull dimension 2. In other words, it actually gives only three independent constraints. Moreover, a solution over the base field can be found within seconds.

It remains to consider the first 5 columns of M_T. Once the last 30 columns are fixed, the variables in C are also determined. Substituting them into the linear equations from the first 5 columns, we can solve the decoupled linear parameters $t_{15}, \ldots, t_{19}$. Since the 30 nonlinear columns form a rank-deficient block, the variables in C still have some remaining freedom. However, this freedom cannot be used iteratively. We explain the reason later.

Two other ways to model the linear constraints are given in Auxiliary Supporting Material A. The first starts from a three dimensional affine subspace;

see Fig. A.1. The second explicitly keeps both linear constraints f_0 and f_1 in the Macaulay matrix; see Fig. A.2. The corresponding Macaulay matrices have different sizes. However, the number of effective constraints computed from them is still 3. This shows that the kernel dimension of the nonlinear block is only a loose estimate. The true number of effective constraints has to be determined by the parameter ideal.

The process of state degree growth is summarized in Fig. 2. Horizontal arrows mean that $\deg_X$ does not increase after the constraint is imposed. The slanted arrow means free propagation, where the degree is multiplied by α.

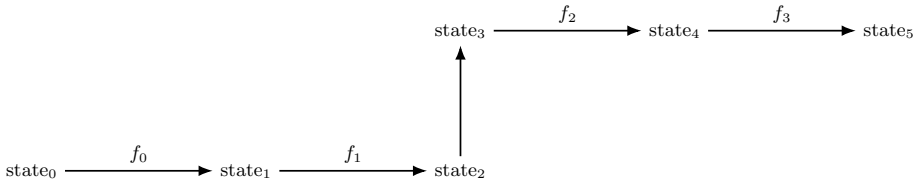

Fig. 2: Constraint placement and state degree growth

The concrete attack pattern shown in Fig. 2 and the corresponding solution for T are listed in Auxiliary Supporting Material A. This nonlinear subspace trail suppresses four degree increases across five internal partial rounds, whereas a traditional linear subspace trail can suppress only three.

3.2 Observations and Algebraic Argument

In this subsection, we first extract several observations from the example in Sect. 3.1. Based on these observations, we use methods similar to [29] to explain the algebraic essence of the parametric Macaulay matrix method.

Observation 1. The construction in Sect. 3.1 lets the state pass freely through one partial round, so that its degree grows from one to three. This step is necessary for extending the trail beyond the linear subspace limit: it allows more than Ec constraints to continue suppressing degree growth in the main variable ideal system.

Linear constraints are too rigid. They cannot introduce coefficient parameters for the main variables, and can only modify constant terms. Geometrically, the parameters can only translate the constraint hyperplanes; they cannot change their orientations or introduce curvature. Thus, the dimension of the ideal is essentially determined by the internal linear layers. We also observed that constraints of degree α^2 or higher do not provide a substantial advantage over constraints of degree α. They mainly transfer the same compatibility problem to a higher degree setting, while making the parameter system more complicated.

Observation 2. The condition $f_3 \in \langle f_1, f_2 \rangle$ in Fig. 1 represents a strong form of algebraic degeneration. It is stronger than the failure of the nonzero divisor property in a semi-regular or regular sequence. For a regular sequence, if $vf_i \in \langle f_1, \ldots, f_{i-1} \rangle$, then one must have $v \in \langle f_1, \ldots, f_{i-1} \rangle$. Equivalently, the image of f_i is a nonzero divisor in the quotient ring $R[X]/\langle f_1, \ldots, f_{i-1} \rangle$. By contrast, the condition $f_3 \in \langle f_1, f_2 \rangle$ means that adding f_3 does not enlarge the ideal at all. Thus, it does not further decrease the Krull dimension of the corresponding quotient ring.

In our attack, we fix $v = 1$. This means that $\langle f_3 \rangle$ is a true subset of $\langle f_1, f_2 \rangle$. In fact, other choices are also possible. For example, one may require $x_0 f_3, x_1 f_3, x_2 f_3, x_3^2 f_3 \in \langle f_1, f_2 \rangle$. In this example, only finitely many low-degree elements in $\langle f_3 \rangle$, such as f_3 and $x_3 f_3$, do not belong to $\langle f_1, f_2 \rangle$. Geometrically, the variety of f_1 and f_2 still covers almost all points of f_3, and the dimension of $R/\langle f_1, f_2, f_3 \rangle$ is still 1.

Proposition 1 (Finite remainders do not change the dimension). *Let $S = R[x_1, \ldots, x_m]$, $I = \langle f_1, \ldots, f_n \rangle$, and suppose that $\dim(S/I) = m - n$. Let $J = \langle v_1, \ldots, v_q \rangle$ be a zero-dimensional ideal, namely $\mathrm{ht}(J) = \dim S$. Assume that $Jf_{n+1} \subseteq I$, equivalently, $v_i f_{n+1} \in I$ holds for all $i = 1, \ldots, q$. Then every positive dimensional prime ideal containing I also contains f_{n+1}. Equivalently, adding f_{n+1} to I does not remove any positive dimensional component of $V(I)$; it can only affect zero-dimensional components.*

Proof. In the quotient ring S/I, the submodule generated by $\bar{f}_{n+1}$ satisfies $J\bar{f}_{n+1} = 0$. Hence the map $S/J \to S\bar{f}_{n+1}$, $\bar{a} \mapsto a\bar{f}_{n+1}$, is surjective. Since J is zero-dimensional, S/J has only finitely many standard monomials. Therefore, $S\bar{f}_{n+1}$ is also finite dimensional. Equivalently, except for finitely many low-degree multiples, all sufficiently high-degree multiples of f_{n+1} can be expressed by $f_1, \ldots, f_n$. Thus, adding f_{n+1} only removes a zero-dimensional remainder, and does not change the positive dimensional part of the ideal.

We stress that, if one uses a more complicated model in which finitely many low-degree elements of $\langle f_3 \rangle$ do not belong to $\langle f_1, f_2 \rangle$, then one must ensure that $v_1 f_3, v_2 f_3, v_3 f_3, v_4 f_3 \in \langle f_1, f_2 \rangle$ hold simultaneously. Otherwise, if only a condition such as $x_0 f_3 \in \langle f_1, f_2 \rangle$ holds, this nonregular behavior may increase the degree at which the relevant Hilbert function vanishes. The Macaulay bound can then be lower than the true highest degree reached in the Gröbner basis computation. Constructing such constraints in a more general way may require more parameters and several Macaulay matrices. Moreover, after degree lifting, the conditions in the higher degree space become more complicated. At present, we have not found a stronger result along this direction. Whether more general ideal degeneration conditions can improve the attack remains an open question.

Algebraic Interpretation of Parametric Macaulay Matrices We now turn the above observations into a criterion that can be used later. Among the first Ec subspace constraints, suppose that the first l are linear constraints.

Following [29], these linear constraints can be absorbed by restricting the initial state to a suitable affine subspace. The remaining $Ec - l$ constraints are degree α nonlinear constraints and generate the basic ideal

$$I_{\text{basic}} = \langle f_l, \ldots, f_{Ec-1} \rangle \subset R[X].$$

In the rest of this paper, we work only with the truncated degree range up to α, without degree lifting. Accordingly, let $V_{\leq\alpha} = \{g \in R[X] : \deg(g) \leq \alpha\}$. We decompose this space as $V_{\leq\alpha} = H \oplus L$, where H is spanned by monomials of degrees $2, \ldots, \alpha$, and L is spanned by monomials of degrees 0 and 1. Let π_H and π_L denote the projections onto H and L, respectively. We also write $I^{(\alpha)}_{\text{basic}} = I_{\text{basic}} \cap V_{\leq\alpha}$ for the degree at most α part of the basic ideal.

Proposition 2 (high-degree projection criterion without degree lifting). *In the above modeling without degree lifting, assume that the degree-α homogeneous components of $\pi_H(f_l), \ldots, \pi_H(f_{E_c-1})$ form a regular sequence. Let a later added constraint in the internal rounds be $f = h + \ell \in V_{\leq\alpha}$, where $h \in H$ and $\ell \in L$. The high-degree part $h = \pi_H(f)$ is determined by the previous intermediate state after the S-box and the linear diffusion. Then there exists a choice of the low-degree parameters such that f degenerates into the basic ideal within degree at most α if and only if $\pi_H(f) \in \pi_H\left(I^{(\alpha)}_{\text{basic}}\right)$. In other words, the discrepancy in the low-degree space L can be balanced by newly introduced parameters. The real condition for maintaining the trail is whether the projection of the new constraint to H has already been quotiented out by the components of the basic constraints in H.*

Proof. Since we do not allow degree lifting here, all constraints being compared are restricted to $V_{\leq\alpha}$. The linear constraints have already been absorbed into the initial state by starting from a linear subspace. Hence, in this truncated space, the basic ideal is generated only by several degree α constraints that reduce to degree one. A later added constraint has the same form.

Under the regularity assumption, the first syzygy module of the degree-α homogeneous components is generated by the Koszul syzygies. Let $g \in I^{(\alpha)}_{\text{basic}}$ and write $g = \sum_{i=l}^{E_c-1} a_i f_i$. If some a_i has positive degree, let $d = \max_i \deg(a_i)$. Since $\deg(g) \leq \alpha$, the homogeneous component of degree $\alpha + d$ gives a syzygy among the degree-α homogeneous components of the basic constraints. Subtracting the corresponding Koszul syzygies of the full generators f_i lowers d without changing g. Repeating this reduction yields a representation with scalar coefficients. Therefore, $I^{(\alpha)}_{\text{basic}} = \text{span}_R\{f_l, \ldots, f_{\lfloor Ec/s \rfloor - 1}\}$.

For each basic constraint f_i, let $h_i = \pi_H(f_i)$ and $\ell_i = \pi_L(f_i)$. If the new constraint $f = h + \ell$ can be expressed by the basic constraints without degree lifting, then there exist scalars λ_i such that $f - \sum_{i=l}^{Ec-1} \lambda_i f_i = 0$. Taking the H projection gives $h = \sum_{i=l}^{Ec-1} \lambda_i h_i$, namely $\pi_H(f) \in \pi_H\left(I^{(\alpha)}_{\text{basic}}\right)$.

Conversely, if this high-degree projection relation holds, then $f - \sum_{i=l}^{Ec-1} \lambda_i f_i$ lies entirely in L. Since $\pi_L(f)$ can be fully controlled by the newly introduced

affine parameters, these parameters can always be chosen so that $f - \sum_{i=l}^{Ec-1} \lambda_i f_i = 0$. Hence, whether the new constraint degenerates into the basic ideal is determined only by whether its H projection lies in the span of the H projections of the basic constraints.

We use the similar linear-algebra technique from the work [29], and decompose the MDS matrix used in the internal rounds of Poseidon as $\left(\begin{smallmatrix} A & B \\ C & D \end{smallmatrix}\right)$, where A and D are the $s \times s$ and $(t-s) \times (t-s)$ blocks, respectively. The decomposition of the matrix is identical in Poseidon2 and Neptune.

In the instance of Sect. 3.1, we have $t = 5$ and $s = 1$, so A is a scalar and D is a 4×4 block. Let the last four coordinates of the initial state be X_0. Following the linear subspace construction in [29], we start from the constraint $BX_0 = 0$. In a linear subspace trail, constant terms can be separated from the main variable system and treated independently. We therefore restrict the following discussion to homogeneous linear constraints on the main variables. Hence, we write $\text{state}_0 = (t_0, X_0)^\mathsf{T}$ and $\text{state}_1 = (t_0', X')^\mathsf{T}$ with $X' = DX_0$.

Figure 3 gives the state transition chain under this notation. After entering the nonlinear trail, before each intermediate state passes through the next S-box, its first coordinate is replaced by a new u_i. These u_i are the only nonlinear sources in the partial rounds. In particular, the newly introduced u_3 in f_3 plays the role of balancing the low-degree space L in Proposition 2.

$$
\begin{array}{ccccc}
S_0 & \xrightarrow{f_0} & S_1 & \xrightarrow{f_1} & S_2 \\
\begin{pmatrix} t_0 \\ X_0 \end{pmatrix} & & \begin{pmatrix} t_0' \\ X' \end{pmatrix} & & \begin{pmatrix} A(t_0'+rc_1)^3 + BX' = u_0 - rc_2 \\ C(t_0'+rc_1)^3 + DX' \end{pmatrix} \\
\to & & S_3 & \xrightarrow{f_2} & S_4 \\
 & & \begin{pmatrix} Au_0^3 + BC(t_0'+rc_1)^3 + BDX' = u_1 - rc_4 \\ Cu_0^3 + DC(t_0'+rc_1)^3 + D^2X' \end{pmatrix} & & \begin{pmatrix} Au_1^3 + BCu_0^3 + BDC(t_0'+rc_1)^3 + BD^2X' = u_2 - rc_5 \\ Cu_1^3 + DCu_0^3 + D^2C(t_0'+rc_1)^3 + D^3X' \end{pmatrix} \\
 & \xrightarrow{f_3} & & & S_5 \\
 & & & & \begin{pmatrix} Au_2^3 + BCu_1^3 + BDCu_0^3 + BD^2C(t_0'+rc_1)^3 + BD^3X' = u_3 - rc_6 \\ Cu_2^3 + DCu_1^3 + D^2Cu_0^3 + D^3C(t_0'+rc_1)^3 + D^4X' \end{pmatrix}
\end{array}
$$

Fig. 3: State transition chain when the degeneration constraints f_0, f_1, f_2, f_3 are imposed successively.

Remark. If $s \neq 1$, then u_i^3 here means cubing each component of u_i separately. For simplicity of presentation, round constants in coordinates other than the first one are omitted from the diagram. These constants only contribute to the low-degree space L. Since they enter only through linear operations, they can always be separated and treated independently. Therefore, they do not affect the conclusions below.

In this instance, after the linear constraints f_0 and f_1 are absorbed, the basic ideal in Proposition 2 is $\langle f_2 \rangle$. To introduce f_3 without changing the Krull

dimension of the ideal, we need $f_3 \in \langle f_2 \rangle$, equivalently $\langle f_2, f_3 \rangle = \langle f_2 \rangle$. From the block matrix expansion in Fig. 3, we obtain

$$f_2 = Au_0^3 + BC(t_0' + rc_1)^3 + BDX' - (u_1 - rc_4),$$

$$f_3 = Au_1^3 + BCu_0^3 + BDC(t_0' + rc_1)^3 + BD^2X' - (u_2 - rc_5).$$

Observation 2 has already explained that this paper does not consider further degree lifting. Therefore, for the two cubic constraints f_2 and f_3, requiring f_2 to divide f_3 can only lead to a scalar multiple relation.

In this example, $\alpha = 3$. Hence H is spanned by quadratic and cubic monomials, and L is spanned by constant and linear monomials. Once the initial affine linear form is fixed, u_0 is already fixed by X_0 and the internal linear layer. In contrast, u_1 is introduced by the constraint f_2 within the Ec budget, and can be chosen freely. By Proposition 2, whether f_3 can be included in $\langle f_2 \rangle$ without degree lifting depends only on whether the high-degree projection of f_3 can be represented by the high-degree projection of f_2. Since both are cubic constraints, this condition degenerates to a scalar multiple relation between their high-degree parts. There are only two possibilities: either u_1 is a constant, or u_1 is proportional to u_0 as an affine linear form.

The first case is to take u_1 as a constant. In this case, the high-degree parts of f_2 and f_3 are reduced to Au_0^3 and BCu_0^3, respectively. Thus, their high-degree projections can be canceled by $-BCA^{-1}f_2 + f_3$. After moving to the low-degree space, Proposition 2 shows that the remaining terms can be balanced by the newly introduced u_2. The required homogeneous linear part is

$$(u_2)_{\deg=1} = (-BCA^{-1}BD + BD^2)X' = (-BCA^{-1}BD + BD^2)DX_0.$$

On the other hand, $(u_0)_{\deg=1} = BX' = BDX_0$, where $X_0 \in W_1$. For a random Cauchy matrix and sufficiently large characteristic in Poseidon, these two block matrix expressions are generally not parallel. A similar obstruction is expected for Poseidon2 and Neptune, since the required proportionality again imposes strong compatibility conditions on the resulting affine forms. Hence, it is hard for u_2 to align with u_0 again.

The second case is to make u_1 linearly dependent on u_0. That is, there exists a diagonal $s \times s$ matrix K such that $u_1 = Ku_0$. Then the high-degree projection of f_3 becomes $(AK^3 + BC)u_0^3$, and hence can be canceled by $-(AK^3 + BC)A^{-1}f_2 + f_3$. The corresponding low-degree balancing condition is

$$(u_2)_{\deg=1} = (BD^2 - (AK^3 + BC)A^{-1}BD)X' + (AK^3 + BC)A^{-1}K(u_0)_{\deg=1}.$$

Even if we take into account that X' comes from a special three dimensional subspace related to D, the above expression still contains several block matrix products. In general, it also does not make u_2 align with u_0 again.

In summary, the true requirement in the high-degree space in Sect. 3.1 is that u_1 and u_0 are proportional as affine linear forms. Once this condition is satisfied, the L space part can be balanced by u_2. For the three main variables x_0, x_1, x_2, the linear dependence between u_1 and u_0 gives exactly three effective constraints.

This explains why the parameter ideal on the high-degree space in Sect. 3.1 stably gives three constraints. At the same time, the two cases above both show that u_2 and u_0 are generally no longer proportional as affine linear forms. If one further imposes a cubic constraint f_4 on state_5, then the high-degree part contains u_2^3, which cannot be quotiented out by $\langle f_2 \rangle$. Therefore, along the path constructed in Sect. 3.1 and in this subsection, one cannot deactivate another partial round S-box and extend the trail by one more round.

Proposition 2 motivates the construction in the next subsection. Based on this criterion, we give a general nonlinear subspace trail construction without relying on parametric Macaulay matrix computations.

3.3 A General Nonlinear Subspace Trail Construction

This subsection gives the general nonlinear subspace trail construction used in our attack models. The construction does not introduce any additional affine constraints on the starting state. It first leaves $u_0, \ldots, u_{\lfloor Ec/s \rfloor - 1}$ undetermined, fixes the middle source to a constant, and aligns the subsequent nonlinear sources with $u_1, \ldots, u_{\lfloor Ec/s \rfloor - 1}$. The balancing conditions determine $u_1, \ldots, u_{\lfloor Ec/s \rfloor - 1}$ together with the final source, while u_0 remains free. Consequently, the high-degree projections of the additional $\lfloor Ec/s \rfloor$ constraints lie in $\pi_H\left(I_{\text{basic}}^{(\alpha)}\right)$. The following discussion is for general instances. The notation t, s, α, Ec is the same as above.

As in Sect. 3.2, we use u_i to pack the first s coordinates before each power S-box. Thus, u_i is the only nonlinear source in each partial round. As above, we omit round constants in coordinates outside the first s positions. The initial state is taken as $\text{state}_0 = (u_0 - rc_0, X)^{\mathsf{T}}$, where u_0 is an undetermined affine linear vector. The coordinates of u_0 and X together form t linearly independent affine linear forms in a family of t main variables. For example, when $s = 1$, one may take $u_0 = x_{t-1}$ and $X = (x_0, x_1, \ldots, x_{t-3}, x_{t-2})^{\mathsf{T}}$. After the first partial round, we introduce the first constraint $f_0 : \; Au_0^{\alpha} + BX = u_1 - rc_1$, where u_1 is undetermined.

The recursion is then completely similar. At the first s coordinates of state_{i+1}, we introduce the constraint

$$f_i : \quad Au_i^{\alpha} + \sum_{j=0}^{i-1} BD^{i-1-j} Cu_j^{\alpha} + BD^i X = u_{i+1} - rc_{i+1},$$

where $i \geq 1$. This iterative process continues until $\text{state}_{2\lfloor Ec/s \rfloor}$ and the constraint $f_{2\lfloor Ec/s \rfloor - 1}$ imposed on it. Except that state_0 is linear, all later states remain of degree α. The path is shown in Fig. 4.

When $s \neq 1$, each f_i actually packs s independent constraints on the first s coordinates. The first $\lfloor Ec/s \rfloor$ constraints are freely added by the attacker. Here f_i degenerates the i-th appearing α degree polynomial to the affine vector $u_{i+1} - rc_{i+1}$, and these constraints generate the basic ideal $I_{\text{basic}} = \langle f_0, \ldots, f_{\lfloor Ec/s \rfloor - 1} \rangle$. When each f_i is unpacked into s constraints, their degree-α homogeneous components have a block lower triangular form. More precisely, the degree-α part of

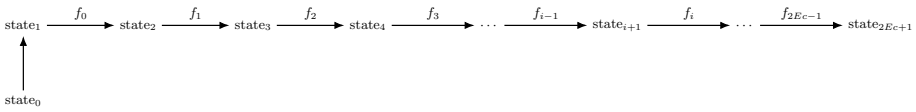

Fig. 4: State transition path corresponding to a nonlinear subspace trail of length $2Ec$ when $s = 1$.

f_i involves only $u_0^\alpha, \ldots, u_i^\alpha$, and the coefficient block of the newly appearing term u_i^α is A. Treating the u_i as vectors of indeterminates, the coefficient matrix of these homogeneous components with respect to the monomial blocks $u_0^\alpha, u_1^\alpha, \ldots$ is block lower triangular with A on every diagonal block. Since A is nonsingular, this matrix is invertible. Therefore, these homogeneous components form a regular sequence. Thus, the regularity condition in Proposition 2 is satisfied.

Let its α degree truncation be $I_{\text{basic}}^{(\alpha)} = \text{span}_R\{f_0, \ldots, f_{\lfloor Ec/s \rfloor - 1}\}$. In order not to reduce the dimension of the coordinate ring of the main system, the later constraints $f_{\lfloor Ec/s \rfloor}, \ldots, f_{2\lfloor Ec/s \rfloor - 1}$ should all fall into I_{basic} without degree lifting. By Proposition 2, it is enough to quotient out the high-degree projections by $\pi_H(I_{\text{basic}}^{(\alpha)})$ and then choose the affine parameters so that the remaining linear terms also vanish modulo $I_{\text{basic}}^{(\alpha)}$.

Step 1. *Align the two halves of the substitution chain.* Set $u_{\lfloor Ec/s \rfloor} = 1$. For $j = 1, \ldots, \lfloor Ec/s \rfloor - 1$, set $u_{\lfloor Ec/s \rfloor + j} = u_j$. More generally, these equalities may be replaced by nonzero scalar multiples. For simplicity, we take all such scalars to be 1 in this paper. Under this choice, the high-degree parts of $f_{\lfloor Ec/s \rfloor}, \ldots, f_{2\lfloor Ec/s \rfloor - 1}$ lie in the span of $\{u_0^\alpha, \ldots, u_{\lfloor Ec/s \rfloor - 1}^\alpha\}$, and no new high-degree source is introduced.

Step 2. *Use the triangular structure of the basic ideal to quotient out the high-degree terms.* Let $H_{\text{basic}} = \pi_H(I_{\text{basic}}^{(\alpha)})$, namely the high-degree space spanned by $H_{f_0}, \ldots, H_{f_{\lfloor Ec/s \rfloor - 1}}$. Since the high-degree projection of f_i in the recursion only involves $u_0^\alpha, \ldots, u_i^\alpha$, these projections have a triangular form: H_{f_0} only contains u_0^α, H_{f_1} only contains u_0^α, u_1^α, and so on. Therefore, for any later constraint f_k, where $k = \lfloor Ec/s \rfloor, \ldots, 2\lfloor Ec/s \rfloor - 1$, one can gradually quotient out $u_{\lfloor Ec/s \rfloor - 1}^\alpha, \ldots, u_0^\alpha$ as

$$\langle f_k \rangle_{\deg \le \alpha} \,/\, I_{\text{basic}}^{(\alpha)} = \langle f_k \rangle_{\deg \le \alpha} \,/\, \langle f_{\lfloor Ec/s \rfloor - 1} \rangle_{\deg \le \alpha} \,/\, \cdots / \, \langle f_0 \rangle_{\deg \le \alpha}.$$

After the high-degree projections are quotiented out, these later constraints leave affine linear remainders in X and the undetermined u_i. Under the substitutions in Step 1, u_0 remains free, whereas $u_1, \ldots, u_{\lfloor Ec/s \rfloor - 1}$ enter the balancing system. The last constraint $f_{2\lfloor Ec/s \rfloor - 1}$ introduces the additional affine vector $u_{2\lfloor Ec/s \rfloor}$.

Step 3. *Determine the earlier nonlinear sources by solving the balancing system in L.* These affine linear constraints form a linear system with $\lfloor Ec/s \rfloor$ equations. After collecting the terms in U on the left and the terms in X on the right, we formally write it as $M_u U = M_x X$, where $U = (u_1, \ldots, u_{\lfloor Ec/s \rfloor - 1},$

$u_{2\lfloor Ec/s\rfloor})^{\mathsf{T}}$. In this system, $u_{2\lfloor Ec/s\rfloor}$ appears only in the last equation. The remaining $\lfloor Ec/s\rfloor - 1$ variables $u_1, \ldots, u_{\lfloor Ec/s\rfloor-1}$ are constrained by the same number of equations. The matrices M_u and M_x are determined by block products of A, B, C, D, and their explicit expansions are cumbersome. In Auxiliary Supporting Material B, we show that the balancing matrix M_u is generically nonsingular for the linear layer considered here. Therefore, the system admits a unique solution for U, given by $U = M_u^{-1} M_x X$.
Therefore, one can solve $u_1, \ldots, u_{\lfloor Ec/s\rfloor-1}, u_{2\lfloor Ec/s\rfloor}$ as affine linear expressions in X. And $u_{\lfloor Ec/s\rfloor}, \ldots, u_{2\lfloor Ec/s\rfloor-1}$ are aligned with $1, u_1, \ldots, u_{\lfloor Ec/s\rfloor-1}$ according to the substitution chain above. Finally, we obtain a nonlinear subspace trail of length $2\lfloor Ec/s\rfloor$, and it satisfies $I_{\text{basic}} = \langle f_0, \ldots, f_{2\lfloor Ec/s\rfloor-1}\rangle$
Hence, after the intermediate state passes through $2\lfloor Ec/s\rfloor+1$ partial rounds, its degree grows only from degree one to degree α.

Even and odd constructions. The construction above has a closely related variant in which u_0 is also included in the alignment. Let $q = \lfloor Ec/s\rfloor$. In the construction described above, u_0 remains an independent affine vector, while $u_q = 1, u_{q+j} = u_j, 1 \le j \le q-1$. Thus, q vector constraints are available to generate I_{basic}, and $q-1$ pairs of nonlinear sources are aligned. This gives a nonlinear subspace trail of length $2q$.

Alternatively, one may spend s degrees of freedom to include u_0 in the alignment. In this case, u_0 is determined by the balancing system rather than remaining independent, so only $Ec - s$ extra constraints remain available for constructing the basic ideal. Hence the number of basic vector constraints becomes $\lfloor \frac{Ec-s}{s} \rfloor = q - 1$.

The corresponding alignment is $u_{q-1} = 1, u_{q+j} = u_j, 0 \le j \le q-2$. Therefore, although the basic ideal has one fewer vector generator, the alignment now includes u_0 and still contains $q-1$ pairs. The resulting trail has length $2\lfloor \frac{Ec-s}{s} \rfloor + 1 = 2q - 1$.

Accordingly, we call the construction developed above the *even construction*, and this variant the *odd construction*. In terms of indices, the even construction uses $I_{\text{basic}} = \langle f_0, \ldots, f_{q-1}\rangle$ and ends with f_{2q-1} and u_{2q}, whereas the odd construction uses $I_{\text{basic}} = \langle f_0, \ldots, f_{q-2}\rangle$ and ends with f_{2q-2} and u_{2q-1}.

The odd construction has a more symmetric algebraic structure. In Auxiliary Supporting Material B, we first establish the generic nonsingularity of its balancing matrix M_u and then extend the analysis to the even construction. Since the odd construction is always one partial round shorter, the complexity comparisons in the remainder of this paper use only the even construction. Accordingly, the even construction is the main nonlinear subspace construction used in our attack models.

Observation 3. (1) A natural way to continue the extension is to tune the scalar coefficients in the substitution chain and the value of $u_{\lfloor Ec/s\rfloor}$, so that the final $u_{2\lfloor Ec/s\rfloor}$ also lies in $\text{span}\{1, u_1, \ldots, u_{\lfloor Ec/s\rfloor-1}\}$. However, this requires an additional rank condition that is not enforced by the balancing system in L. The entries of the matrix defining this additional rank condition are given by

repeated block matrix products. For a general MDS matrix over a large characteristic field, one should not expect such rank drops to occur systematically. We did not observe such exploitable linear dependence in our Poseidon experiments, nor in experiments with the internal linear layers of Poseidon2 and Neptune.

(2) This paper only deals with constraints that degenerate degree α to degree 1. If one degenerates to a polynomial of degree lower than α but still nonlinear, then the degree growth may also be restricted. However, the H/L decomposition is no longer clean, and usually more constrained rounds are needed to obtain the same effect. Whether there exists a truly better variant of this type remains an open problem.

(3) The construction in Sect. 3.3 does not explicitly use a parametric Macaulay matrix. However, for the nonlinear degeneration variants in (2) above or the degree-lifted modeling in Observation 2, parametric Macaulay matrices still provide an algorithmic search method. In other words, Sects. 3.3 can be viewed as analytically tractable special cases from this parameter system perspective.

A concrete instance. Finally, we provide a concrete instance motivated by the 2026 Poseidon bounty program for the CICO-2 setting. This instance is not a complete attack, but an explicit nonlinear subspace trail for the internal partial rounds. It is constructed over the KoalaBear field with $t = 16$ and $\alpha = 3$. Since two input coordinates and two output coordinates are constrained, the available extra constraint budget is $Ec = 16 - 2 - 2 = 12$. Applying the construction of Sect. 3.3, we obtain a nonlinear subspace trail of length $2Ec = 24$. The verification code and the explicit polynomial path are available in the online repository.

4 Attack Models

This section compares several competitive algebraic modeling strategies for preimage attacks on Poseidon and Neptune. We evaluate the gain of our nonlinear subspace model under a unified complexity metric. Following the analysis of Grassi et al., the main bottleneck in solving these Poseidon-like systems under the standard Gröbner basis workflow is usually the F4/F5 computation under grelex or grevlex order. The FGLM stage, which converts the basis to lexicographic order, is usually less dominant. Therefore, we mainly compare the complexity of the first Gröbner basis stage. When the model gives a zero-dimensional univariate or bivariate ideal, we instead use the Graeffe transform or resultant based solving, since these methods have a clear advantage over Gröbner basis solving in the univariate and bivariate settings [25].

For the main system ideals generated by the model in Sect. 3.3 when $\alpha = 3$, we also tested semi-regularity following the method of [25]. The experiments suggest that these systems are typically semi-regular. Thus, the following complexity estimates treat them as semi-regular sequences.

4.1 Attack Models

This subsection lists the attack models to be compared. Section 4.2 then computes the number of rounds covered by these models under the same complexity budget. All models below are for preimage attacks. According to the refined CICO analysis in [29], the CICD problem in preimage attacks can be transformed into the CICC problem. Therefore, we take $\min\{c, d\} = d$ without loss of generality. Then $Ec = r - \min\{c, d\} = r - d = t - (c + d)$. In compression mode, one has $Ec = t - \min\{c, d\} = t - d$. We discuss this difference in Sect. 4.3.

Start-in-the-Middle (SitM) is used only for NEPTUNE, because inverting the external rounds of POSEIDON causes a degree explosion [25]. The other models discussed below apply to Poseidon, POSEIDON2, and NEPTUNE. POSEIDON can additionally use input subspaces for round skipping, which is not discussed separately in the following comparison.

We do not separately consider four known strategies because of the reasons below. First, input subspaces are avoided by the initial linear mixing layer of POSEIDON2 [7,24]. Second, infinite subspace trails are shown in [29] to be unlikely over large characteristic fields. Third, round-level modeling trades a large number of variables for low-degree equations, and existing experiments show that it does not give a competitive attack [1]. Fourth, the early linear subspace estimate [20] was later shown to be imprecise, since one cannot simply subtract the subspace length from the number of internal rounds when estimating degree growth [25].

Besides the nonlinear subspace construction in Sect. 3.3, we also compare a modeling strategy that has received less attention. When guiding an intermediate state into a subspace, this strategy does not introduce variable substitution. Instead, it directly constrains a high-degree intermediate state to lie in a subspace. We give the concrete forms in the models below.

(1) Basic Attack (Inputs–Outputs)

The basic attack does not construct any subspace. It uses the following model:

$$X \in \mathbb{F}_p^r \xrightarrow{E^{r_f'} \circ I^{r_P} \circ E^{r_f} \circ M'(\cdot)} \text{digest} \in \mathbb{F}_p^d$$

This gives d equations and introduces r variables.

For POSEIDON/POSEIDON2, all equations have degree $\alpha^{r_P + r_F}$. For NEPTUNE, all equations have degree $4^{r_F} \cdot \alpha^{r_P}$. For both systems, if $d = 1$, we use $\mathcal{C}_{\text{UNI}}$ to estimate the complexity; if $d = 2$, we use $\mathcal{C}_{\text{RES}}$; and if $d \geq 3$, we use $\mathcal{C}_{\text{GB}}$.

(2) Forward+Substitution+Linear subspace

Model 2 follows the Forward+Substitution+linear subspace strategy in [25]. It first propagates forward and substitutes variables, and then forces the intermediate state to enter a linear subspace. The model is

$$X \in \mathbb{F}_p^r \xrightarrow{I^{\tau} \circ E^{r_f} \circ M'(\cdot)} \text{lin-}S \xrightarrow{E^{r_f'} \circ I^{r_P - \tau - Ec}} \text{digest} \in \mathbb{F}_p^d$$

Here the coset is fixed as a constant, and lin-S denotes a linear subspace. We use τ to denote the number of internal partial rounds before entering the subspace. This model gives t equations of degree $\zeta^{r_f}\alpha^{\tau}$ at lin-S, and d equations of

degree $\zeta^{r_f'}\alpha^{r_P-\tau-Ec}$ at the digest side. Here $\zeta=\alpha$ for Poseidon, and $\zeta=4$ for NEPTUNE. It introduces $t+d$ variables. Since $t+d>2$, we use Gröbner basis computation:

$$\begin{aligned}\min_{0\le\tau\le r_P-Ec}\mathcal{C}_{\mathrm{GB}} &= O\left(\binom{1+t\zeta^{r_f}\alpha^{\tau}+d\zeta^{r_f'}\alpha^{r_P-\tau-Ec}}{t+d}^{w}\right)\\ &\ge O\left(\binom{1+2\sqrt{dt}\zeta^{(r_f+r_f')/2}\alpha^{(r_P-Ec)/2}}{t+d}^{w}\right).\end{aligned}$$

(3) Forward+Substitution+Nonlinear subspace

Model 3 replaces the linear subspace in Model 2 by the nonlinear subspace constructed in Sect. 3.3:

$$X\in\mathbb{F}_p^r \xrightarrow{I^{\tau}\circ E^{r_f}\circ M'(\cdot)} \text{nonlin-}S \xrightarrow{E^{r_f'}\circ I^{r_P-\tau-2Ec-1}} \text{digest}\in\mathbb{F}_p^d$$

This model gives t equations of degree $\zeta^{r_f}\alpha^{\tau}$ at nonlin-S. The nonlinear subspace constraints give Ec equations of degree α. The digest side gives d equations of degree $\zeta^{r_f'}\alpha^{r_P-\tau-2Ec}$. The total number of variables is $r+t=2t-c$. Hence,

$$\begin{aligned}\min_{0\le\tau\le r_P-(2Ec+1)}\mathcal{C}_{\mathrm{GB}} &= O\left(\binom{1+t\zeta^{r_f}\alpha^{\tau}+d\zeta^{r_f'}\alpha^{r_P-\tau-2Ec}+Ec\cdot\alpha}{2t-c}^{w}\right)\\ &\ge O\left(\binom{1+2\sqrt{dt}\zeta^{(r_f+r_f')/2}\alpha^{(r_P-2Ec)/2}+Ec\cdot\alpha}{2t-c}^{w}\right).\end{aligned}$$

(4) Forward+Nonsubstitution+Linear subspace

Model 4 does not replace variables in the middle rounds. Instead, it directly imposes Ec linear subspace constraints on a high-degree intermediate state, and then gives d final equations at the digest side. In this case, the optimal choice is $\tau=0$, and the path is

$$X\in\mathbb{F}_p^r \xrightarrow{E^{r_f}\circ M'(\cdot)} \text{lin-}C \xrightarrow{E^{r_f'}\circ I^{r_P-Ec}} \text{digest}\in\mathbb{F}_p^d$$

Here lin-C means that the high-degree state is directly constrained. Since $r=d+Ec$, only instances with $t=3$ may use resultant solving. Other instances use Gröbner basis computation:

$$\mathcal{C}_{\mathrm{GB}} = O\left(\binom{1+d\zeta^{r_F}\alpha^{r_P-Ec}+Ec\zeta^{r_f}}{r}^{w}\right).$$

(5) Forward+Nonsubstitution+Nonlinear subspace

Model 5 directly constrains a high-degree intermediate state to lie in a nonlinear subspace.

$$X\in\mathbb{F}_p^r \xrightarrow{E^{r_f}\circ M'(\cdot)} \text{nonlin-}C \xrightarrow{E^{r_f'}\circ I^{r_P-2Ec-1}} \text{digest}\in\mathbb{F}_p^d$$

As in Model 4, except for instances with $t = 3$, where resultant solving may be used, all other instances use Gröbner basis computation:

$$\mathcal{C}_{\text{GB}} = O\left(\binom{1 + d\zeta^{r_F}\alpha^{r_P - 2Ec} + Ec\alpha\zeta^{r_f}}{r}^{w}\right).$$

(6) SitM+Linear subspace

SitM modeling is used only for NEPTUNE. This model introduces variables at the beginning of the internal rounds. It then goes forward to satisfy the digest constraints and backward to satisfy the capacity constraints, while using a linear subspace:

$$\text{capacity} \in \mathbb{F}_p^c \xleftarrow{(M')^{-1} \circ (E^{r_f})^{-1}(\cdot)} X \in \text{lin-}S \xrightarrow{E^{r'_f} \circ I^{r_P - Ec}(\cdot)} \text{digest} \in \mathbb{F}_p^d$$

The input capacity side gives c equations of degree ζ^{r_f}. The output digest side gives d equations of degree $\zeta^{r'_f}\alpha^{r_P - Ec}$. When $c + d = 2$, we use resultant solving, with complexity

$$\mathcal{C}_{\text{RES}} = 4 \cdot \left(4^{r'_f} \cdot \alpha^{r_P - t + 2}\right)^3 \cdot \left(\log\left(4^{r'_f} \cdot \alpha^{r_P - t + 2}\right)\right)^2 \cdot \left(1 + \log_2\left(\log_2\left(4^{r'_f} \cdot \alpha^{r_P - t + 2}\right)\right)\right).$$

When $c + d > 2$, we use Gröbner basis computation:

$$\mathcal{C}_{\text{GB}} = O\left(\binom{1 + c \cdot 4^{r_f} + d \cdot 4^{r'_f}\alpha^{r_P - Ec}}{c + d}^{w}\right).$$

Following the treatment of Grassi et al. [25], if the digest side equations reach degree $p - 2$ in some instances, we estimate their degree as $p - 2$.

(7) SitM+Nonlinear subspace

Model 7 adds Ec nonlinear subspace constraints to Model 6:

$$\text{capacity} \in \mathbb{F}_p^c \xleftarrow{(M')^{-1} \circ (E^{r_f})^{-1}(\cdot)} X \in \text{nonlin-}S \xrightarrow{E^{r'_f} \circ I^{r_P - 2Ec - 1}(\cdot)} \text{digest} \in \mathbb{F}_p^d$$

Since $c + d + Ec = t > 2$, this model uses Gröbner basis computation:

$$\mathcal{C}_{\text{GB}} = O\left(\binom{1 + c \cdot 4^{r_f} + d \cdot 4^{r'_f}\alpha^{r_P - 2Ec} + Ec \cdot \alpha}{t}^{w}\right).$$

4.2 Experimental Results and Discussion

This subsection compares the number of internal rounds covered by different attack models for POSEIDON2 and NEPTUNE under a unified complexity model. We recall all results of POSEIDON2 can be extended to POSEIDON by using $r_f - 1$, since the first full round of POSEIDON can be skipped for free.

The main results are shown in Fig. 5 and Fig. 6. Here the round number refers to the minimum number of internal partial rounds required for the corresponding algebraic solving complexity to reach or exceed 2^{128}. Throughout this comparison, we fix the security target at 128 bits and keep the number of external full rounds unchanged. We now summarize the main observations from these experimental comparisons.

(1) The nonlinear subspace model covers more internal rounds than the previous linear subspace model in almost all instances under the 2^{128} complexity bound. The improvement is usually close to a factor of two. For some POSEIDON/POSEIDON2 parameters, such as $p \approx 2^{256}$, $c = d = 1$, and $\alpha = 7$, the new model even reaches or exceeds the number of internal rounds recommended by the designers.

(2) In some instances, nonlinear subspace modeling outperforms the basic attack, and this advantage becomes more visible as the state width t grows. The basic attack is not very sensitive to t. By contrast, in several parameter families, the capacity c and digest length d remain small as t increases, leaving more extra constraints Ec to the attacker. The nonlinear subspace model exploits these constraints more efficiently and bypasses more internal partial rounds. As a result, the advantage gained from reducing the ideal degree outweighs the advantage of Graeffe acceleration or resultant solving over Gröbner basis solving.

(3) The experiments show that variable substitution has almost no influence on the final number of covered rounds. In most instances, the substitution and nonsubstitution models yield nearly identical results. Thus, under the current complexity estimates, variable substitution is not the main factor determining the round number. Nevertheless, the systems obtained after substitution may contain many redundant syzygies, which is also observed in round-level modeling [1]. Therefore, whether substitution models are always preferable in practical F4/F5 computation still requires further experimental verification.

(4) For NEPTUNE, the SitM modeling strategy is generally stronger than the Forward strategy. The relative advantage between linear and nonlinear subspace models depends on the state width t. For small or medium size parameters, the linear subspace model may perform better because the corresponding resultant computation is cheaper. As t grows, however, the stronger degree reduction achieved by nonlinear subspaces becomes the dominant effect, and the nonlinear model shows a clearer advantage.

We also check whether the FGLM stage changes the round-number conclusions. The results for the Forward+Substitution+Nonlinear subspace model are given in Table C.1 of Auxiliary Supporting Material C. In almost all instances, FGLM gives the same round numbers as the F4/F5 estimate, and only slightly lower values for a few small-variable cases. The similar behavior is observed for the other models. Since $\mathcal{C}_{\mathrm{FGLM}}$ is estimated with $\omega = 3$ while $\mathcal{C}_{\mathrm{GB}}$ is estimated with $\omega = 2$, this comparison is already conservative. We therefore treat the F4/F5 stage as the main bottleneck.

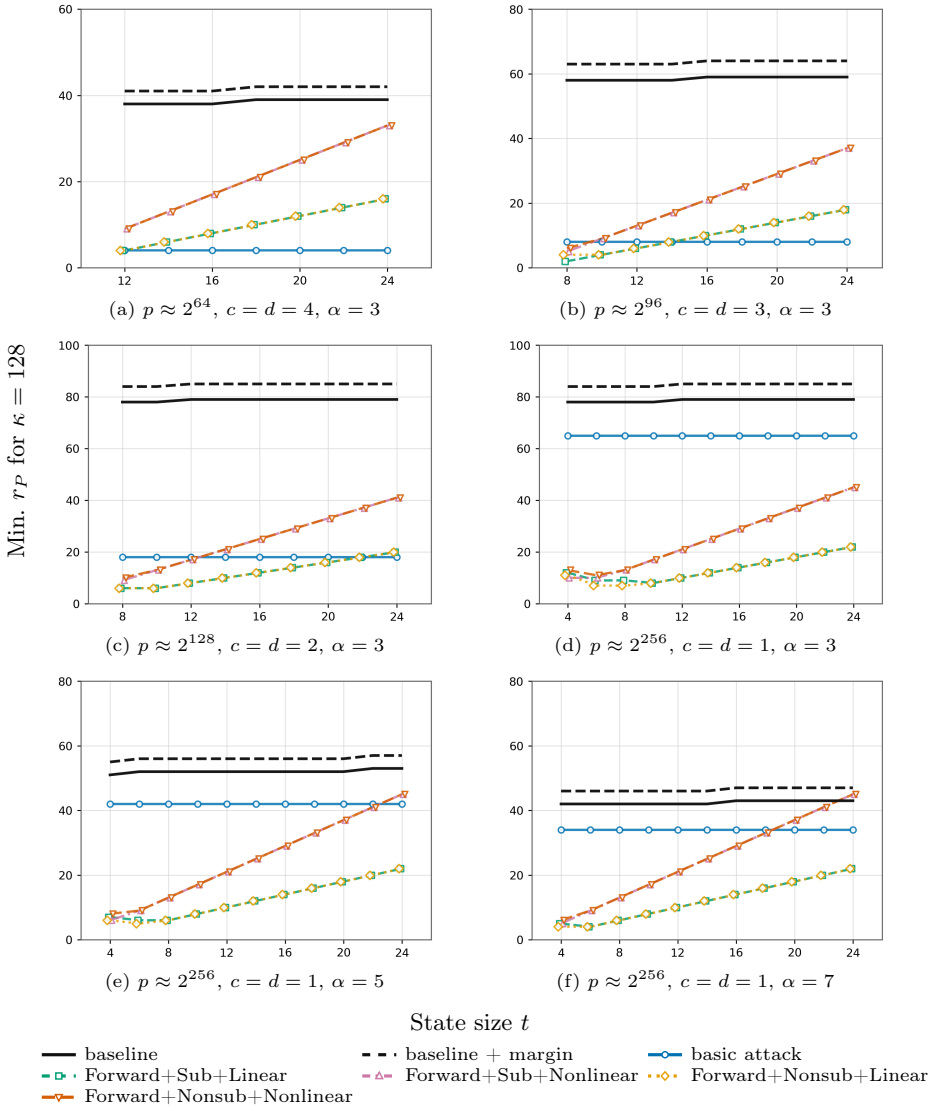

Fig. 5: Minimum number of internal rounds r_P required to achieve $\kappa = 128$-bit security for Poseidon2 over $\mathbb{F}_p$ in sponge mode, plotted against the state size t. The capacity c and digest length d are set to $\lceil \log_p(2^{2\kappa}) \rceil$. We fix $r_F = 6$ and compare the results with the designers' recommendation for r_P, both with and without the 7.5% security margin. Curves based on linear or nonlinear subspace models exploit the maximum available subspace constraint budget $Ec = t - (c + d)$. Here $\omega = 2$.

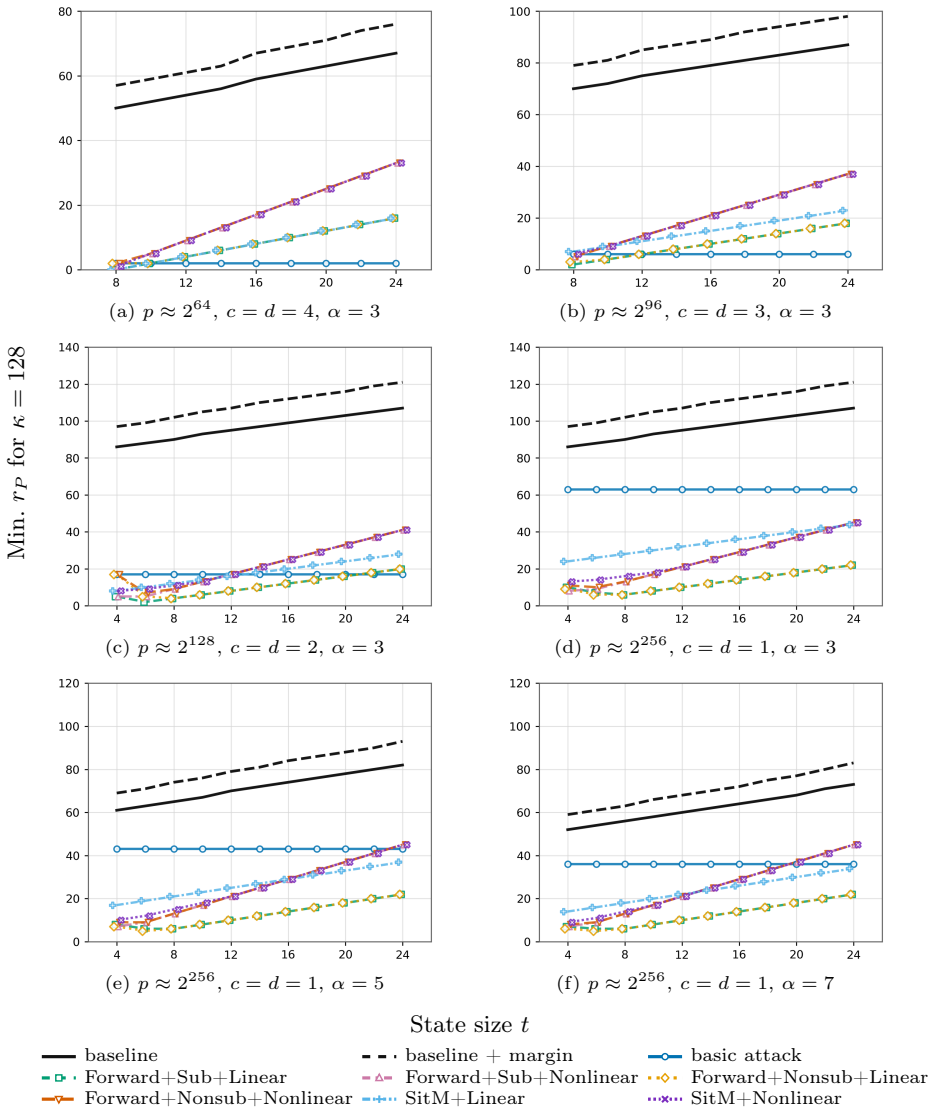

Fig. 6: Minimum number of internal rounds r_P required to achieve $\kappa = 128$-bit security for NEPTUNE over $\mathbb{F}_p$ in sponge mode, plotted against the state size t. The capacity c and digest length d are set to $\lceil \log_p(2^{2\kappa}) \rceil$. We compare the results with the designers' recommendation for r_P, both with and without the 12.5% security margin. Curves based on linear or nonlinear subspace models exploit the maximum available subspace constraint budget $Ec = t - (c + d)$. Here $\omega = 2$.

4.3 Discussion on Compression Mode

Compared with sponge mode, compression mode has no input side capacity constraint. Although only small state Poseidon2 instances currently have a formal compression mode specification, we still compute the compression mode complexity for the Poseidon/Poseidon2 and Neptune instances considered above. The results show that nonlinear subspace modeling reaches or exceeds the designers' recommended number of rounds in three cases. In one of them, it even exceeds the recommendation with the security margin. These cases are listed in Table C.2 of Auxiliary Supporting Material C. They may serve as useful references for future studies of compression mode designs.

5 Conclusion

This paper studied algebraic attack modeling for Poseidon/Poseidon2 and Neptune from a linear-algebraic and computational-algebraic perspective. We generalized linear subspace trails to nonlinear subspace constraints, and introduced parametric Macaulay matrices as a computational framework for constructing such constraints. Our construction extends a linear subspace trail of length $\lfloor Ec/s \rfloor$ to a nonlinear subspace trail of length $2\lfloor Ec/s \rfloor$.

The complexity estimates show that this longer trail directly leads to more analyzable internal rounds. In most instances, the nonlinear subspace model covers approximately twice as many partial rounds as the linear subspace model, and in several concrete cases it reaches or exceeds the designers' recommended number of internal rounds. We also compared modeling variants with and without variable substitution. Overall, nonlinear subspaces provide an effective modeling method for algebraic security analysis of Poseidon-like hash functions.

References

1. Ashur, T., Buschman, T., Mahzoun, M.: Algebraic cryptanalysis of the HADES design strategy: Application to Poseidon and Poseidon2. In: ACISP 2024. LNCS, vol. 14896, pp. 225–244. Springer, Cham (2024). https://doi.org/10.1007/978-981-97-5028-3_12
2. Bak, A., Bariant, A., Boeuf, A., Briaud, P., Øygarden, M., Phanse, A.: The algebraic CheapLunch: Extending FreeLunch attacks on arithmetization-oriented primitives beyond CICO-1. Cryptology ePrint Archive, Report 2025/2040 (2025). https://eprint.iacr.org/2025/2040
3. Bak, A., Bariant, A., Boeuf, A., Hostettler, M., Jazeron, G.: Claiming bounties on small scale Poseidon and Poseidon2 instances using resultant-based algebraic attacks. Cryptology ePrint Archive, Report 2026/150 (2026). https://eprint.iacr.org/2026/150
4. Bardet, M., Faugère, J.-C., Salvy, B.: On the complexity of the F5 Gröbner basis algorithm. J. Symbolic Comput. **70**, 49–70 (2015). https://doi.org/10.1016/j.jsc.2014.09.025

5. Bariant, A., Boeuf, A., Lemoine, A., Manterola Ayala, I., Øygarden, M., Perrin, L., Raddum, H.: The algebraic FreeLunch: Efficient Gröbner basis attacks against arithmetization-oriented primitives. In: CRYPTO 2024. LNCS, vol. 14923, pp. 139–173. Springer, Cham (2024). https://doi.org/10.1007/978-3-031-68385-5_5
6. Bariant, A., Boeuf, A., Briaud, P., Hostettler, M., Øygarden, M., Raddum, H.: Improved resultant attack against arithmetization-oriented primitives. In: CRYPTO 2025. LNCS, vol. 16004, pp. 335–367. Springer, Cham (2025). https://doi.org/10.1007/978-3-032-01901-1_11
7. Bariant, A., Bouvier, C., Leurent, G., Perrin, L.: Algebraic attacks against some arithmetization-oriented primitives. IACR Trans. Symmetric Cryptol. **2022**(3), 73–101 (2022). https://doi.org/10.46586/tosc.v2022.i3.73-101
8. Bertoni, G., Daemen, J., Peeters, M., Van Assche, G.: On the indifferentiability of the sponge construction. In: EUROCRYPT 2008. LNCS, vol. 4965, pp. 181–197. Springer, Berlin, Heidelberg (2008). https://doi.org/10.1007/978-3-540-78967-3_11
9. Bertoni, G., Daemen, J., Peeters, M., Van Assche, G.: Cryptographic sponge functions. Version 0.1 (2011). https://keccak.team/files/CSF-0.1.pdf
10. Beyne, T., Canteaut, A., Dinur, I., Eichlseder, M., Leander, G., Leurent, G., Naya-Plasencia, M., Perrin, L., Sasaki, Y., Todo, Y., Wiemer, F.: Out of oddity: New cryptanalytic techniques against symmetric primitives optimized for integrity proof systems. In: CRYPTO 2020. LNCS, vol. 12172, pp. 299–328. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-56877-1_11
11. Cantor, D.G., Kaltofen, E.L.: On fast multiplication of polynomials over arbitrary algebras. Acta Inform. **28**(7), 693–701 (1991). https://doi.org/10.1007/BF01178683
12. Cox, D.A., Little, J., O'Shea, D.: *Using Algebraic Geometry*. 2nd edn. Graduate Texts in Mathematics, vol. 185. Springer, New York (2005). https://doi.org/10.1007/b138611
13. Dubois, V., Gama, N.: The degree of regularity of HFE systems. In: ASIACRYPT 2010. LNCS, vol. 6477, pp. 557–576. Springer, Berlin, Heidelberg (2010). https://doi.org/10.1007/978-3-642-17373-8_32
14. Faugère, J.-C., Gianni, P.M., Lazard, D., Mora, T.: Efficient computation of zero-dimensional Gröbner bases by change of ordering. J. Symbolic Comput. **16**(4), 329–344 (1993). https://doi.org/10.1006/jsco.1993.1021
15. Faugère, J.-C.: A new efficient algorithm for computing Gröbner bases (F4). J. Pure Appl. Algebra **139**(1–3), 61–88 (1999). https://doi.org/10.1016/S0022-4049(99)00005-5
16. Faugère, J.-C.: A new efficient algorithm for computing Gröbner bases without reduction to zero (F5). In: ISSAC 2002, pp. 75–83. ACM, New York (2002). https://doi.org/10.1145/780506.780516
17. Faugère, J.-C., Bardet, M., Salvy, B.: On the complexity of Gröbner basis computation of semi-regular overdetermined algebraic equations. In: International Conference on Polynomial System Solving, pp. 71–74 (2004).
18. Faugère, J.-C., Perret, L.: Algebraic attacks against STARK-friendly ciphers. Cryptology ePrint Archive, Report 2020/948 (2020). https://eprint.iacr.org/2020/948
19. von zur Gathen, J., Gerhard, J.: *Modern Computer Algebra*. 3rd edn. Cambridge University Press, Cambridge (2013). https://doi.org/10.1017/CBO9781139856065

20. Grassi, L., Lüftenegger, R., Rechberger, C., Rotaru, D., Schofnegger, M.: On a generalization of substitution-permutation networks: The HADES design strategy. In: EUROCRYPT 2020. LNCS, vol. 12106, pp. 674–704. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-45724-2_23
21. Grassi, L., Khovratovich, D., Rechberger, C., Roy, A., Schofnegger, M.: Poseidon: A new hash function for zero-knowledge proof systems. In: 30th USENIX Security Symposium (USENIX Security 21), pp. 519–535. USENIX Association (2021). https://www.usenix.org/conference/usenixsecurity21/presentation/grassi
22. Grassi, L., Rechberger, C., Schofnegger, M.: Proving resistance against infinitely long subspace trails: How to choose the linear layer. IACR Trans. Symmetric Cryptol. **2021**(2), 314–352 (2021). https://doi.org/10.46586/tosc.v2021.i2.314-352
23. Grassi, L., Onofri, S., Pedicini, M., Sozzi, L.: 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). https://doi.org/10.46586/tosc.v2022.i3.20-72
24. Grassi, L., Khovratovich, D., Schofnegger, M.: Poseidon2: A faster version of the Poseidon hash function. In: AFRICACRYPT 2023. LNCS, vol. 14064, pp. 177–203. Springer, Cham (2023). https://doi.org/10.1007/978-3-031-37679-5_8
25. Grassi, L., Koschatko, K., Rechberger, C.: Poseidon and Neptune: Gröbner basis cryptanalysis exploiting subspace trails. IACR Trans. Symmetric Cryptol. **2025**(2), 34–86 (2025). https://doi.org/10.46586/tosc.v2025.i2.34-86
26. Keller, N., Rosemarin, A.: Mind the middle layer: The HADES design strategy revisited. In: EUROCRYPT 2021. LNCS, vol. 12697, pp. 35–63. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-77886-6_2
27. Koschatko, K., Lüftenegger, R., Rechberger, C.: Exploring the six worlds of Gröbner basis cryptanalysis: Application to Anemoi. IACR Trans. Symmetric Cryptol. **2024**(4), 138–190 (2024). https://doi.org/10.46586/tosc.v2024.i4.138-190
28. Lai, X., Massey, J.L.: A proposal for a new block encryption standard. In: EUROCRYPT 1990. LNCS, vol. 473, pp. 389–404. Springer, Berlin, Heidelberg (1991). https://doi.org/10.1007/3-540-46877-3_35
29. Li, E., Wang, G.: Revisiting linear subspace trails in Poseidon2 and Neptune. Cryptology ePrint Archive, Report 2026/967 (2026). https://eprint.iacr.org/2026/967
30. Merz, S.-P., Rodríguez García, À.: Skipping class: Algebraic attacks exploiting weak matrices and operation modes of Poseidon2(b). Cryptology ePrint Archive, Report 2026/306 (2026). https://eprint.iacr.org/2026/306
31. Sanso, A., Vitto, G.: Attacking Poseidon via Graeffe-based root-finding over NTT-friendly fields. Cryptology ePrint Archive, Report 2025/937 (2025). https://eprint.iacr.org/2025/937
32. Sanso, A., Vitto, G.: Top Gun: Degree annihilation attacks on Poseidon. Cryptology ePrint Archive, Report 2026/1254 (2026). https://eprint.iacr.org/2026/1254
33. Vitto, G.: Slipway: Accessing finite subspace trails in Poseidon. Cryptology ePrint Archive, Report 2026/1579 (2026). https://eprint.iacr.org/2026/1579
34. Yang, H.-S., Zheng, Q.-X., Yang, J., Liu, Q.-F., Tang, D.: A new security evaluation method based on resultant for arithmetic-oriented algorithms. In: ASIACRYPT 2024. LNCS, vol. 15490, pp. 457–489. Springer, Singapore (2024). https://doi.org/10.1007/978-981-96-0941-3_15

35. Zhao, Z., Ding, J.: Breaking Poseidon challenges with Graeffe transforms and complexity analysis by FFT lower bounds. Cryptology ePrint Archive, Report 2025/950 (2025). https://eprint.iacr.org/2025/950

Auxiliary Supporting Material A: Additional Parametric Macaulay Matrices and a Concrete Nonlinear Subspace Trail Instance

Auxiliary Parametric Macaulay matrix diagrams. The following two diagrams correspond to the alternative modeling choices discussed in Sect. 3.1 for the state width $t = 5$. The first diagram corresponds to the case where the initial state of the subspace construction starts from a three dimensional affine linear subspace. In this case, only the nonlinear constraints f_2 and f_3 appear in the parametric Macaulay matrix. The second diagram corresponds to the case where the initial state starts from the full five dimensional affine linear space. The two linear constraints f_0 and f_1 are then kept explicitly in the parametric Macaulay matrix together with f_2 and f_3.

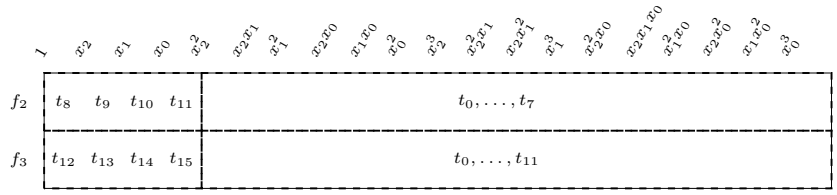

Fig. A.1: Parametric Macaulay matrix for a subspace construction starting from a three dimensional affine linear subspace. Only the nonlinear constraints f_2 and f_3 are kept explicitly.

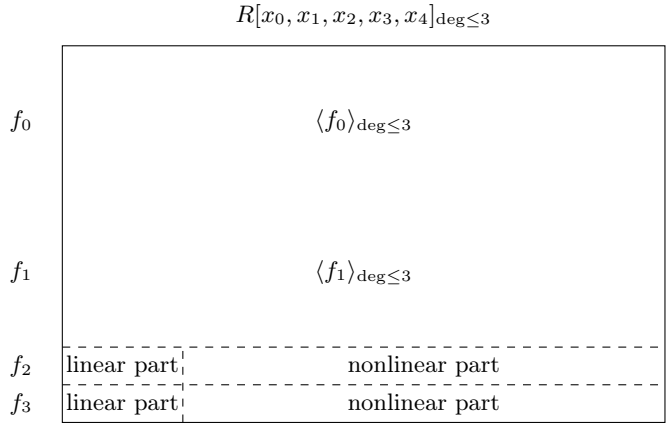

Fig. A.2: Parametric Macaulay matrix for a subspace construction starting from the full five dimensional affine linear space. The linear constraints f_0 and f_1 are kept explicitly together with f_2 and f_3.

A Nonlinear Subspace Trail Instance for Internal Rounds of Poseidon

We give a concrete instance corresponding to the parametric Macaulay matrix in Fig. 1 of Sect. 3.1. The computation is carried out over the KoalaBear prime field, and we use the setting $Ec = 3$. The Cauchy MDS matrix is sampled at random and is chosen to have no invariant subspace relevant to the linear trail analysis. The round constants are also sampled uniformly over the base field. The parameter values T listed below are obtained according to the construction described in Sect. 3.1.

Global parameters.

- Field characteristic: $p = 2130706433$.
- State width: $t = 5$.
- S-box exponent: $\alpha = 3$.

Solved parameters $t_0, \ldots, t_{19}$.

$t_0 = 482429642$ $t_1 = 1264188353$ $t_2 = 172669890$ $t_3 = 1608911998$ $t_4 = 909765259$
$t_5 = 991176468$ $t_6 = 1293612732$ $t_7 = 779174805$ $t_8 = 1483728965$ $t_9 = 1070554810$
$t_{10} = 410783974$ $t_{11} = 586002162$ $t_{12} = 1039982828$ $t_{13} = 795339404$ $t_{14} = 1166762331$
$t_{15} = 1206996979$ $t_{16} = 1131330311$ $t_{17} = 1951226827$ $t_{18} = 1371824262$ $t_{19} = 0$

Cauchy MDS matrix.

$$M = \begin{bmatrix} 736667229 & 970160856 & 850438748 & 1830105517 & 1753672653 \\ 1342567709 & 149224043 & 1090950053 & 198697406 & 641895687 \\ 857864942 & 2038965227 & 1746376617 & 502711643 & 1439873814 \\ 803378364 & 985639787 & 802424001 & 688794579 & 674416884 \\ 48882600 & 1710554494 & 855454602 & 1417219922 & 698222071 \end{bmatrix}$$

Round constants.

Round	Constant vector
0	1813382118, 827307999, 1627694678, 1911784257, 903170602
1	86939546, 556019485, 2073320061, 1097954097, 1043521778
2	869589436, 1971893209, 1683194652, 1782095536, 651359108
3	2078334659, 1023484162, 768905175, 1252803391, 1914826684
4	1950157577, 469102188, 1083806779, 299088013, 605242169

Auxiliary Supporting Material B: Rank Analysis of the Balancing Matrix in the Nonlinear Subspace Construction

This supporting material gives the rank analysis of the balancing matrix M_u arising from the nonlinear subspace construction in Sect. 3.3. We first rewrite the successive elimination of the basic constraints as a block Toeplitz elimination problem. This gives an explicit expression for the affine remainders of the later constraints and reduces the nonsingularity of M_u to that of a smaller structured matrix. We first analyze the odd construction, whose pairing structure gives a particularly regular matrix representation, and then extend the argument to the even construction used in the attack models. Finally, we prove generic nonsingularity for the linear layer structures relevant to the constructions considered in this paper.

B.1 Block Toeplitz representation of the basic constraints

We use the notation of Sect. 3.3. Let m denote the number of basic vector constraints. We keep m independent of the concrete constraint budget in the derivation below; it is instantiated according to the odd or even construction when the corresponding balancing system is considered.

Partition the internal linear layer according to the first s active coordinates as

$$M = \begin{pmatrix} A & B \\ C & D \end{pmatrix},$$

where

$$A \in \mathbb{F}_p^{s\times s}, \qquad B \in \mathbb{F}_p^{s\times(t-s)}, \qquad C \in \mathbb{F}_p^{(t-s)\times s}, \qquad D \in \mathbb{F}_p^{(t-s)\times(t-s)}.$$

We assume that A is nonsingular and define its Schur complement

$$S = D - CA^{-1}B.$$

For the nonlinear sources and affine outputs, write

$$z_i = u_i^{\alpha}, \qquad \eta_i = u_{i+1} - rc_{i+1}.$$

Define the block sequence

$$\Gamma_0 = A, \qquad \Gamma_k = BD^{k-1}C \quad (k \geq 1).$$

Then the recursive constraint in Sect. 3.3 can be written as

$$f_i = \sum_{j=0}^{i} \Gamma_{i-j} z_j + BD^i X - \eta_i.$$

The basic ideal is generated by $f_0, \ldots, f_{m-1}$.

The degree α source terms $z_0, \ldots, z_{m-1}$ in these constraints have the block coefficient matrix

$$\mathcal{T}_m = \begin{pmatrix} A & 0 & 0 & \cdots & 0 \\ BC & A & 0 & \cdots & 0 \\ BDC & BC & A & \cdots & 0 \\ BD^2C & BDC & BC & \ddots & \vdots \\ \vdots & \vdots & \vdots & \ddots & A \end{pmatrix}.$$

Since A is nonsingular, $\mathcal{T}_m$ is nonsingular.

Proposition 3. *Let f be any later constraint whose degree α part lies in the span of $z_0, \ldots, z_{m-1}$. The affine remainder obtained by successively eliminating $z_{m-1}, \ldots, z_0$ with the basic constraints $f_{m-1}, \ldots, f_0$, as in Sect. 3.3, is identical to the affine remainder obtained by solving the block system defined by $\mathcal{T}_m$ and eliminating all z_i simultaneously.*

Proof. Write the degree α part of f as

$$R_0 z_0 + \cdots + R_{m-1} z_{m-1}.$$

To eliminate all z_i, one has to find block coefficients $L_0, \ldots, L_{m-1}$ such that

$$f - \sum_{i=0}^{m-1} L_i f_i$$

contains no z_i. Since the coefficient matrix of $z_0, \ldots, z_{m-1}$ in the basic constraints is $\mathcal{T}_m$, the coefficients L_i are the unique solution of the corresponding linear system. The successive elimination used in Sect. 3.3 is precisely block forward substitution in this lower triangular system. Multiplication by $\mathcal{T}_m^{-1}$ gives the same unique solution. Hence both procedures cancel the same degree α terms using the same linear combination of the basic constraints, and therefore leave the same affine remainder.

Proposition 3 allows us to replace the successive elimination in the construction by the more convenient matrix operation $\mathcal{T}_m^{-1}$.

Lemma 1. *The inverse of $\mathcal{T}_m$ is again block lower triangular Toeplitz. If G_k denotes the block on its kth subdiagonal, then*

$$G_0 = A^{-1}, \qquad G_k = -A^{-1} B S^{k-1} C A^{-1} \quad (k \geq 1).$$

Equivalently,

$$\mathcal{T}_m^{-1} = \begin{pmatrix} A^{-1} & 0 & 0 & \cdots & 0 \\ -A^{-1}BCA^{-1} & A^{-1} & 0 & \cdots & 0 \\ -A^{-1}BSCA^{-1} & -A^{-1}BCA^{-1} & A^{-1} & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ * & * & * & \cdots & A^{-1} \end{pmatrix}.$$

Proof. The Toeplitz blocks of $\mathcal{T}_m^{-1}$ satisfy

$$AG_0 = I_s, \qquad AG_k + \sum_{j=1}^{k} BD^{j-1}CG_{k-j} = 0 \quad (k \geq 1).$$

Thus $G_0 = A^{-1}$. For $k \geq 1$, substituting

$$G_k = -A^{-1}BS^{k-1}CA^{-1}$$

into the recurrence reduces its left hand side to

$$B\left(D^{k-1} - S^{k-1} - \sum_{j=1}^{k-1} D^{j-1}(D-S)S^{k-j-1}\right)CA^{-1}.$$

Since $D - S = CA^{-1}B$, the expression in parentheses vanishes by the matrix identity

$$D^r - S^r = \sum_{j=0}^{r-1} D^{r-1-j}(D-S)S^j,$$

which does not require D and S to commute. This proves the claimed formula.

B.2 Normalized generators of the basic ideal

We next use Lemma 1 to replace the basic constraints by an equivalent set of generators in which each generator contains only one degree α source term. Let

$$F_{\text{basic}} = (f_0, \ldots, f_{m-1})^{\mathsf{T}}, \qquad (g_0, \ldots, g_{m-1})^{\mathsf{T}} = \mathcal{T}_m^{-1} F_{\text{basic}}.$$

Since $\mathcal{T}_m$ is nonsingular, the constraints $g_0, \ldots, g_{m-1}$ generate the same ideal as $f_0, \ldots, f_{m-1}$.

Lemma 2. *For $0 \leq i \leq m-1$, the normalized generator g_i is*

$$g_i = z_i + A^{-1}BS^iX - \sum_{j=0}^{i} G_{i-j}\eta_j,$$

where $G_0 = A^{-1}$ and $G_k = -A^{-1}BS^{k-1}CA^{-1}$ for $k \geq 1$. In particular, z_i is the only degree α source term appearing in g_i.

Proof. By construction, multiplication by $\mathcal{T}_m^{-1}$ transforms the degree α coefficient matrix of the basic constraints into the identity, so the degree α part of g_i is exactly z_i.

It remains to determine the coefficient of X. From the ith block row of $\mathcal{T}_m^{-1}F_{\text{basic}}$, this coefficient is

$$\sum_{j=0}^{i} G_{i-j}BD^j.$$

Using the expression for G_k in Lemma 1 and the identity $D - S = CA^{-1}B$, the same matrix power difference identity gives

$$\sum_{j=0}^{i} G_{i-j} B D^j = A^{-1} B S^i.$$

The affine terms are transformed by the same block row and therefore contribute $-\sum_{j=0}^{i} G_{i-j}\eta_j$. Combining these terms gives the claimed expression.

Thus, reduction modulo the basic ideal can be performed by eliminating each z_i directly with the corresponding normalized generator g_i. This form will be used below to derive the balancing matrix for the odd construction.

B.3 Reduction of the later constraints in the odd construction

We first consider the odd construction. Recall that its nonlinear sources satisfy $z_m = 1$ and $z_{m+1+j} = z_j$ for $0 \le j \le m-1$.

Consider a later constraint f_{m+k}, where $0 \le k \le m$. After applying the above pairing, the coefficient of z_i, for $0 \le i \le m-1$, is

$$R_{k,i} = \Gamma_{m+k-i} + \mathbf{1}_{\{i \le k-1\}} \Gamma_{k-1-i}.$$

The first term is the original contribution of z_i, while the second term comes from its paired copy z_{m+1+i}. Hence,

$$f_{m+k} = \sum_{i=0}^{m-1} R_{k,i} z_i + \Gamma_k + BD^{m+k} X - \eta_{m+k} + c_k,$$

where Γ_k is the contribution of the fixed source $z_m = 1$ and $c_k \in \mathbb{F}_p^s$ is the constant contribution of the fixed source $z_m = 1$.

We now eliminate $z_0, \ldots, z_{m-1}$ using the normalized generators of Lemma 2. Define

$$\overline{f}_{m+k} = f_{m+k} - \sum_{i=0}^{m-1} R_{k,i} g_i.$$

By construction, $\overline{f}_{m+k}$ contains no degree α source term and is affine linear in X and the remaining u_i.

Lemma 3. *For $0 \le k \le m$, the reduced constraint is*

$$\begin{aligned}\overline{f}_{m+k} = {} & \left(BD^k S^m - \mathbf{1}_{\{k \ge 1\}} BD^{k-1}\right) X - \eta_{m+k} \\ & + \sum_{j=0}^{m-1} \left(BD^k S^{m-j-1} CA^{-1} + \mathbf{1}_{\{k \ge 1,\ j=k-1\}} I_s\right) \eta_j + \varepsilon_k.\end{aligned}$$

where $\varepsilon_k \in \mathbb{F}_p^s$ is a constant vector independent of X and of the undetermined nonlinear sources.

Proof. After substituting the expression for g_i, the coefficient of η_j in $\overline{f}_{m+k}$ is

$$\sum_{i=j}^{m-1} R_{k,i} G_{i-j}.$$

Using the definition of $R_{k,i}$, split this sum according to its two terms. For the first part, Lemma 1 and the matrix power difference identity give

$$\sum_{i=j}^{m-1} \Gamma_{m+k-i} G_{i-j} = BD^k S^{m-j-1} C A^{-1}.$$

The second part is a coefficient of the Toeplitz convolution between Γ_i and G_i. Since these two block sequences are inverse to each other, it equals I_s when $j = k-1$ and vanishes otherwise. Therefore,

$$\sum_{i=j}^{m-1} R_{k,i} G_{i-j} = BD^k S^{m-j-1} C A^{-1} + \mathbf{1}_{\{k\geq 1,\ j=k-1\}} I_s.$$

Similarly, the coefficient of X is

$$BD^{m+k} - \sum_{i=0}^{m-1} R_{k,i} A^{-1} B S^i.$$

The contribution from the first term of $R_{k,i}$ reduces to $BD^k S^m$. For $k \geq 1$, the contribution from the second term is

$$\sum_{i=0}^{k-1} \Gamma_{k-1-i} A^{-1} B S^i = BD^{k-1},$$

again by the same Toeplitz convolution identity. This gives

$$BD^k S^m - \mathbf{1}_{\{k\geq 1\}} BD^{k-1}.$$

All remaining terms are independent of X and of the undetermined nonlinear sources and are collected into ε_k.

B.4 Reduction of the balancing matrix in the odd construction

We now collect the coefficients of the undetermined affine vectors in the reduced constraints of Lemma 3. Recall that $\eta_j = u_{j+1} - rc_{j+1}$ and that the odd construction uses $u_m = 1$ and $u_{m+1+j} = u_j$ for $0 \leq j \leq m-1$.

For $k = 0$, the term $-\eta_m$ contributes $-u_0$. Hence u_0 appears in the first reduced constraint with coefficient $-I_s$.

For $1 \leq k \leq m-1$, the term $-\eta_{m+k}$ contributes $-u_k$. On the other hand, the additional identity block in Lemma 3, corresponding to $j = k-1$, contributes $+u_k$ through η_{k-1}. These two terms cancel. The remaining coefficients are therefore obtained from

$$BD^k S^{m-j-1} C A^{-1} \eta_j.$$

Writing $\ell = j + 1$, the coefficient of u_ℓ is

$$(N_{\mathrm{o}})_{k,\ell} = BD^k S^{m-\ell} CA^{-1}, \qquad 1 \le k, \ell \le m-1.$$

Finally, for $k = m$, the term $-\eta_{2m}$ introduces the new affine vector $-u_{2m+1}$, which appears only in the last reduced constraint.

Let

$$U_{\mathrm{o}} = (u_0, u_1, \ldots, u_{m-1}, u_{2m+1})^{\mathsf{T}}.$$

With the reduced constraints ordered by $k = 0, 1, \ldots, m$, the balancing matrix has the block form

$$M_u^{\mathrm{o}} = \begin{pmatrix} -I_s & * & 0 \\ 0 & N_{\mathrm{o}} & 0 \\ 0 & * & -I_s \end{pmatrix},$$

where

$$N_{\mathrm{o}} = \left(BD^k S^{m-\ell} CA^{-1}\right)_{1 \le k, \ell \le m-1}.$$

Proposition 4. *The balancing matrix M_u^{o} is nonsingular if and only if N_{o} is nonsingular. Moreover,*

$$\det M_u^{\mathrm{o}} = \det N_{\mathrm{o}}.$$

Proof. The two boundary diagonal blocks of M_u^{o} are both $-I_s$. Hence its determinant is the product of the determinants of the three diagonal blocks,

$$\det M_u^{\mathrm{o}} = \det(-I_s)^2 \det N_{\mathrm{o}} = \det N_{\mathrm{o}}.$$

The equivalence follows immediately.

Thus, the rank analysis of the balancing system in the odd construction reduces entirely to the smaller matrix N_{o}. The code for constructing N_{o} from a given matrix M and checking whether it is full rank is available at online

B.5 Krylov factorization of the core matrix

Let $n = m - 1$. Define

$$\mathcal{O}_D = \begin{pmatrix} BD \\ BD^2 \\ \vdots \\ BD^n \end{pmatrix} \in \mathbb{F}_p^{ns \times (t-s)}$$

and

$$\mathcal{K}_S = \left(S^n CA^{-1}\ S^{n-1} CA^{-1} \cdots SCA^{-1}\right) \in \mathbb{F}_p^{(t-s) \times ns}.$$

Then the block definition of N_{o} gives

$$N_{\mathrm{o}} = \mathcal{O}_D \mathcal{K}_S.$$

The first factor is the observability type Krylov matrix associated with (B, D), while the second is the controllability type Krylov matrix associated with (S, C). This factorization makes the dimensional restriction on the balancing system explicit.

Proposition 5. *A necessary condition for N_{o}, and hence M_u^{o}, to be nonsingular is*

$$ns \le t - s,$$

where $n = m - 1$.

Proof. If N_{o} is nonsingular, then $\operatorname{rank} N_{\mathrm{o}} = ns$. Since $N_{\mathrm{o}} = \mathcal{O}_D \mathcal{K}_S$ and both factors pass through a space of dimension $t - s$,

$$\operatorname{rank} N_{\mathrm{o}} \le t - s.$$

Therefore $ns \le t - s$.

When $t - s = ns$, both Krylov factors are square and

$$\det N_{\mathrm{o}} = \det(\mathcal{O}_D) \det(\mathcal{K}_S).$$

When $t - s > ns$, full row rank of $\mathcal{O}_D$ and full column rank of $\mathcal{K}_S$ are not sufficient by themselves. The product can still lose rank if

$$\ker(\mathcal{O}_D) \cap \operatorname{im}(\mathcal{K}_S) \neq \{0\}.$$

Therefore, nonsingularity of the original linear layer, or even the MDS property, does not by itself imply nonsingularity of N_{o}.

MDS does not imply nonsingularity of N_{o}. As a concrete example, consider the Cauchy matrix over $\mathbb{F}_{101}$

$$M = \begin{pmatrix} 14 & 96 & 44 & 33 \\ 91 & 75 & 29 & 84 \\ 83 & 95 & 97 & 59 \\ 58 & 45 & 2 & 8 \end{pmatrix}.$$

It is generated by

$$(x_1, x_2, x_3, x_4) = (94, 39, 57, 83), \qquad (y_1, y_2, y_3, y_4) = (72, 27, 69, 56),$$

through $M_{ij} = (x_i + y_j)^{-1}$. Since the Cauchy parameters satisfy the standard distinctness and nonzero denominator conditions, M is MDS. However, for $s = 1$ and $m = 3$, the corresponding core matrix is

$$N_{\mathrm{o}} = \begin{pmatrix} 33 & 53 \\ 33 & 53 \end{pmatrix}, \qquad \operatorname{rank} N_{\mathrm{o}} = 1.$$

Hence the MDS property alone is insufficient to guarantee the nonsingularity required by the balancing construction.

B.6 The Poseidon2 internal linear layer

We now specialize the preceding analysis to the internal linear layer of Poseidon2. Since the standard construction has $s = 1$, write

$$A = a, \qquad B = \mathbf{1}^{\mathsf{T}}, \qquad C = \mathbf{1}, \qquad D = \Lambda + \mathbf{1}\mathbf{1}^{\mathsf{T}},$$

where $\Lambda = \mathrm{diag}(\lambda_1, \ldots, \lambda_{t-1})$. If the full internal matrix is written as $\mathrm{diag}(\lambda_0, \lambda_1, \ldots, \lambda_{t-1}) + \mathbf{1}\mathbf{1}^{\mathsf{T}}$, then $a = \lambda_0 + 1$. Its Schur complement is

$$S = D - a^{-1}\mathbf{1}\mathbf{1}^{\mathsf{T}} = \Lambda + \beta\mathbf{1}\mathbf{1}^{\mathsf{T}}, \qquad \beta = 1 - a^{-1}.$$

Thus, both D and S belong to the same diagonal plus rank one matrix family.

Let $n = m - 1$ and reverse the column order of N_{o}, obtaining $\widetilde{N}_{\mathrm{o}}$. Then

$$(\widetilde{N}_{\mathrm{o}})_{r,k} = a^{-1}\mathbf{1}^{\mathsf{T}} D^r S^k \mathbf{1}, \qquad 1 \le r, k \le n.$$

Define

$$V_D = (D\mathbf{1}, D^2\mathbf{1}, \ldots, D^n\mathbf{1}), \qquad V_S = (S\mathbf{1}, S^2\mathbf{1}, \ldots, S^n\mathbf{1}).$$

Since D and S are symmetric,

$$\widetilde{N}_{\mathrm{o}} = a^{-1} V_D^{\mathsf{T}} V_S.$$

Hence the core matrix is a cross Gram matrix between two Krylov sequences. Column reversal only changes the determinant by

$$\det N_{\mathrm{o}} = (-1)^{n(n-1)/2} \det \widetilde{N}_{\mathrm{o}}.$$

This structure does not imply nonsingularity for every admissible choice of the diagonal parameters. For example, over $\mathbb{F}_{101}$, take

$$(\lambda_0, \lambda_1, \lambda_2, \lambda_3) = (7, 83, 66, 15).$$

Then the corresponding internal matrix, as well as D and S, are nonsingular, but for $m = 3$ one obtains

$$N_{\mathrm{o}} = \begin{pmatrix} 39 & 67 \\ 95 & 13 \end{pmatrix}, \qquad \det N_{\mathrm{o}} = 0.$$

Thus, pairwise distinct diagonal parameters and nonsingularity of the internal linear layer do not by themselves guarantee the nonsingularity of the balancing matrix.

Corollary 1. *Suppose $n = t - 1$. If $a \neq 0$, $\det D \neq 0$, $\det S \neq 0$, and the diagonal parameters $\lambda_1, \ldots, \lambda_{t-1}$ are pairwise distinct, then N_{o} is nonsingular.*

Proof. For $T_\gamma = \Lambda + \gamma\mathbf{1}\mathbf{1}^{\mathsf{T}}$, each vector $T_\gamma^k\mathbf{1}$ is obtained by evaluating a monic polynomial of degree k at $\lambda_1, \ldots, \lambda_{t-1}$. Therefore the corresponding Krylov matrix differs from a Vandermonde matrix only by an invertible triangular change of basis. Consequently,

$$\det V_D = \det D \prod_{i<j} (\lambda_j - \lambda_i), \qquad \det V_S = \det S \prod_{i<j} (\lambda_j - \lambda_i).$$

Hence

$$\det N_{\mathrm{o}} = (-1)^{n(n-1)/2} a^{-n} \det D \det S \prod_{i<j} (\lambda_j - \lambda_i)^2,$$

which is nonzero under the stated assumptions.

The full size case above is useful for exposing the algebraic structure, although the parameter range relevant to our attacks is generally truncated, with $n < t-1$. In that case, the Cauchy Binet formula gives

$$\det \widetilde{N}_{\mathrm{o}} = a^{-n} \sum_{\substack{J \subseteq \{1,\ldots,t-1\} \\ |J|=n}} \det V_D[J,:] \det V_S[J,:].$$

The different terms in this sum may cancel, which explains why pairwise distinct diagonal parameters alone are insufficient in the truncated case. Nevertheless, the determinant is not identically zero.

Lemma 4. *For fixed $a \neq 0$ and $n \leq t-1$, the determinant of N_{o} is not the zero polynomial in the diagonal parameters $\lambda_1, \ldots, \lambda_{t-1}$.*

Proof. Set $\lambda_i = \rho \xi_i$ and $\Xi = \mathrm{diag}(\xi_1, \ldots, \xi_{t-1})$. Then

$$D = \rho \Xi + \mathbf{1}\mathbf{1}^{\mathsf{T}}, \qquad S = \rho \Xi + \beta \mathbf{1}\mathbf{1}^{\mathsf{T}}.$$

For $1 \leq r, k \leq n$, the highest degree term in ρ is

$$(\widetilde{N}_{\mathrm{o}})_{r,k} = a^{-1} \rho^{r+k} \sum_{i=1}^{t-1} \xi_i^{r+k} + \text{terms of lower degree in } \rho.$$

Therefore the highest degree coefficient of $\det \widetilde{N}_{\mathrm{o}}$ is

$$a^{-n} \det \left(\sum_{i=1}^{t-1} \xi_i^{r+k} \right)_{1 \leq r,k \leq n}.$$

By the Cauchy Binet formula,

$$\det \left(\sum_{i=1}^{t-1} \xi_i^{r+k} \right)_{1 \leq r,k \leq n} = \sum_{\substack{J \subseteq \{1,\ldots,t-1\} \\ |J|=n}} \left(\prod_{i \in J} \xi_i^2 \right) \Delta(\xi_J)^2.$$

Choose $\xi_1, \ldots, \xi_n$ pairwise distinct and nonzero, and set $\xi_{n+1} = \cdots = \xi_{t-1} = 0$. Then only the term $J = \{1, \ldots, n\}$ remains and it is nonzero. Hence the highest degree coefficient is not identically zero, so neither is $\det N_{\mathrm{o}}$.

Corollary 2. *For the Poseidon2 internal linear layer, the balancing matrix M_u^{o} of the odd construction is generically nonsingular whenever the dimensional condition $(m-1)s \leq t - s$ is satisfied. Equivalently, the parameter choices for which M_u^{o} is singular lie in a proper algebraic subset of the admissible parameter space.*

Proof. By Proposition 4, M_u^{o} is nonsingular exactly when N_{o} is nonsingular. Lemma 4 shows that $\det N_{\mathrm{o}}$ is not identically zero. Its zero set is therefore a proper algebraic subset.

B.7 Extension to the even construction

We finally consider the even construction used in the attack models. The basic ideal remains $I_{\mathrm{basic}} = \langle f_0, \dots, f_{m-1} \rangle$, so the block Toeplitz matrix $\mathcal{T}_m$, its inverse, and the normalized generators of Lemma 2 remain unchanged.

The only difference is the pairing of the nonlinear sources. In the even construction, $z_m = 1$ and $z_{m+i} = z_i$ for $1 \le i \le m-1$. Thus, unlike the odd construction, z_0 has no paired copy in the second half of the path.

For the later constraint f_{m+k}, where $0 \le k \le m-1$, the coefficient of z_i becomes

$$R^{\mathrm{e}}_{k,i} = \Gamma_{m+k-i} + \mathbf{1}_{\{1 \le i \le k\}} \Gamma_{k-i}.$$

Applying the same normalized elimination as above gives the same Krylov contribution as in the odd construction, except for an additional boundary term caused by the missing copy of z_0.

Let $n = m-1$ and order the affine variables as $U_{\mathrm{e}} = (u_1, \dots, u_{m-1}, u_{2m})^{\mathsf{T}}$. Define the $ns \times ns$ core matrix

$$(N_{\mathrm{e}})_{k,\ell} = BD^k S^{m-\ell} C A^{-1} - \mathbf{1}_{\{\ell=1\}} \Gamma_k A^{-1}, \qquad 0 \le k \le m-2, \quad 1 \le \ell \le m-1.$$

With this ordering, the balancing matrix has the form

$$M^{\mathrm{e}}_u = \begin{pmatrix} N_{\mathrm{e}} & 0 \\ * & -I_s \end{pmatrix}.$$

Proposition 6. *The balancing matrix M^{e}_u is nonsingular if and only if N_{e} is nonsingular. Moreover,*

$$\det M^{\mathrm{e}}_u = (-1)^s \det N_{\mathrm{e}}.$$

Proof. The last unpaired affine vector u_{2m} occurs only in the final reduced constraint, with coefficient $-I_s$. The remaining affine variables have coefficient matrix N_{e}, which gives the displayed block triangular form. The determinant identity and the equivalence follow immediately.

The relation with the odd construction is especially transparent if we separate the Krylov part of N_{e}. Define

$$\widehat{N}_{\mathrm{e}} = \left(BD^k S^{m-\ell} C A^{-1} \right)_{\substack{0 \le k \le m-2 \\ 1 \le \ell \le m-1}}.$$

Then

$$\widehat{N}_{\mathrm{e}} = \begin{pmatrix} B \\ BD \\ \vdots \\ BD^{n-1} \end{pmatrix} \left(S^n C A^{-1} \; S^{n-1} C A^{-1} \; \cdots \; S C A^{-1} \right),$$

and

$$N_{\mathrm{e}} = \widehat{N}_{\mathrm{e}} - \begin{pmatrix} \Gamma_0 A^{-1} \\ \Gamma_1 A^{-1} \\ \vdots \\ \Gamma_{n-1} A^{-1} \end{pmatrix} \left(I_s \; 0 \; \cdots \; 0 \right).$$

Hence the even core matrix differs from the corresponding Krylov matrix only by a correction of rank at most s. For POSEIDON2, where $s = 1$, this is a rank one boundary correction.

The generic rank analysis therefore proceeds as in the odd construction. For the POSEIDON2 structure, set again $\lambda_i = \rho\xi_i$. The Krylov part of N_{e} has the same highest degree behavior in ρ as before, whereas the boundary correction has strictly lower degree and therefore cannot cancel the highest degree term of the determinant. The same Cauchy Binet argument consequently shows that $\det N_{\mathrm{e}}$ is not identically zero.

Corollary 3. *For the POSEIDON2 internal linear layer, the balancing matrix M_u^{e} of the even construction is generically nonsingular whenever the dimensional condition $(m-1)s \le t - s$ is satisfied.*

Proof. The highest degree term of $\det N_{\mathrm{e}}$ under the scaling $\lambda_i = \rho\xi_i$ is the same nonzero Hankel type determinant obtained in the odd construction. The additional boundary correction has lower degree in ρ and does not affect this term. Thus $\det N_{\mathrm{e}}$ is not identically zero. Proposition 6 then gives the claim for M_u^{e}.

B.8 Algebraic interpretation of generic nonsingularity

We finally clarify the meaning of generic nonsingularity used above. Consider the entries of the linear layer

$$M = \begin{pmatrix} A & B \\ C & D \end{pmatrix}$$

as algebraic parameters, restricted to the open set $\det A \neq 0$. Since $S = D - CA^{-1}B$, every entry of the core matrices N_{o} and N_{e} is a rational function of these parameters.

After clearing powers of $\det A$, their determinants can therefore be written as

$$\det N_{\mathrm{o}} = \frac{\Psi_{\mathrm{o}}(A,B,C,D)}{(\det A)^{\nu_{\mathrm{o}}}}, \qquad \det N_{\mathrm{e}} = \frac{\Psi_{\mathrm{e}}(A,B,C,D)}{(\det A)^{\nu_{\mathrm{e}}}},$$

for some polynomials $\Psi_{\mathrm{o}}, \Psi_{\mathrm{e}}$ and nonnegative integers $\nu_{\mathrm{o}}, \nu_{\mathrm{e}}$.

Proposition 7. *Suppose that the determinant of a core matrix N arising from the balancing construction is not identically zero on the parameter family under consideration. Then the set of parameters for which N is nonsingular is a nonempty Zariski open subset of that family. Equivalently, rank deficiency is confined to a proper algebraic subset.*

Proof. On the open set $\det A \neq 0$, the condition $\det N = 0$ is equivalent, after clearing denominators, to the vanishing of a polynomial Ψ. If $\det N$ is not identically zero, then Ψ is not the zero polynomial. Hence its zero set is a proper algebraic subset, and its complement is a nonempty Zariski open set.

Here, "generic" is used only in this algebraic sense. It means that the property holds outside a proper algebraic subset. It does not by itself give a numerical probability bound for uniformly sampled parameters over a finite field.

Combining Proposition 7 with Lemma 4 and Corollary 3 gives the following conclusion.

Corollary 4. *For the Poseidon2 internal linear layer, both the odd and even balancing matrices are generically nonsingular whenever their corresponding dimensional conditions are satisfied.*

For concrete finite field instances, generic nonsingularity does not replace an instance level rank check. The balancing matrices used in the experimental parameter sets can be verified directly over their defining fields.

Auxiliary Supporting Material C: Additional Internal Round Number Estimates

This supplementary material collects two additional tables used in the discussion of Sect. 4.2. The first table reports, under the same complexity budget of 2^{128}, the numbers of internal partial rounds required for the Gröbner basis and FGLM complexity estimates to reach this bound in the Forward+Substitution+Nonlinear subspace model. The second table lists compression mode instances in which nonlinear subspace modeling reaches or exceeds the designers' recommended number of internal rounds.

Table C.1: Number of internal rounds at which the GB and FGLM complexities reach 2^{128} for the Forward+Substitution+Nonlinear subspace model. The entries are written as $\mathcal{C}_{\mathrm{GB}}/\mathcal{C}_{\mathrm{FGLM}}$. We consider Poseidon2 over $\mathbb{F}_p$ in sponge mode, with capacity c and digest length d set to $\lceil \log_p(2^{2\kappa}) \rceil$ and $\kappa = 128$. We fix $r_F = 6$ and do not include the security margin. The left value uses the estimated GB complexity $\mathcal{C}_{\mathrm{GB}}(d_{\mathrm{MAC}}, \omega = 2)$, while the right value uses the conjectured FGLM complexity $\mathcal{C}_{\mathrm{FGLM}}(d_I, \omega = 3)$.

Instance	$t = 4$	$t = 6$	$t = 8$	$t = 10$	$t = 12$	$t = 14$	$t = 16$	$t = 18$	$t = 20$	$t = 22$	$t = 24$
a: $p \approx 2^{64}$, $c = d = 4$, $\alpha = 3$					9/9	13/13	17/17	21/21	25/25	29/29	33/33
b: $p \approx 2^{96}$, $c = d = 3$, $\alpha = 3$			5/5	9/9	13/13	17/17	21/21	25/25	29/29	33/33	37/37
c: $p \approx 2^{128}$, $c = d = 2$, $\alpha = 3$			9/9	13/13	17/17	21/21	25/25	29/29	33/33	37/37	41/41
d: $p \approx 2^{256}$, $c = d = 1$, $\alpha = 3$	10/9	10/10	13/13	17/17	21/21	25/25	29/29	33/33	37/37	41/41	45/45
e: $p \approx 2^{256}$, $c = d = 1$, $\alpha = 5$	6/5	9/9	13/13	17/17	21/21	25/25	29/29	33/33	37/37	41/41	45/45
f: $p \approx 2^{256}$, $c = d = 1$, $\alpha = 7$	5/5	9/9	13/13	17/17	21/21	25/25	29/29	33/33	37/37	41/41	45/45

Table C.2: Compression mode instances in which nonlinear subspace modeling reaches or exceeds the designers' recommended number of internal rounds. Results of Poseidon in the table are valid for both Poseidon and Poseidon2. The free full round skipping available for Poseidon is not included here.

System	Case	p	t	c	d	α	Ec	Model	r_P
Poseidon	a	2^{64}	24	0	4	3	20	baseline	39
Poseidon	a	2^{64}	24	0	4	3	20	baseline margin	42
Poseidon	a	2^{64}	24	0	4	3	20	basic attack	4
Poseidon	a	2^{64}	24	0	4	3	20	Forward+Sub+Linear	20
Poseidon	a	2^{64}	24	0	4	3	20	Forward+Sub+Nonlinear	41
Poseidon	a	2^{64}	24	0	4	3	20	Forward+Nonsub+Linear	20
Poseidon	a	2^{64}	24	0	4	3	20	Forward+Nonsub+Nonlinear	41
Poseidon	f	2^{256}	22	0	1	7	21	baseline	43
Poseidon	f	2^{256}	22	0	1	7	21	baseline margin	47
Poseidon	f	2^{256}	22	0	1	7	21	basic attack	34
Poseidon	f	2^{256}	22	0	1	7	21	Forward+Sub+Linear	21
Poseidon	f	2^{256}	22	0	1	7	21	Forward+Sub+Nonlinear	43
Poseidon	f	2^{256}	22	0	1	7	21	Forward+Nonsub+Linear	21
Poseidon	f	2^{256}	22	0	1	7	21	Forward+Nonsub+Nonlinear	43
Poseidon	f	2^{256}	24	0	1	7	23	baseline	43
Poseidon	f	2^{256}	24	0	1	7	23	baseline margin	47
Poseidon	f	2^{256}	24	0	1	7	23	basic attack	34
Poseidon	f	2^{256}	24	0	1	7	23	Forward+Sub+Linear	23
Poseidon	f	2^{256}	24	0	1	7	23	Forward+Sub+Nonlinear	47
Poseidon	f	2^{256}	24	0	1	7	23	Forward+Nonsub+Linear	23
Poseidon	f	2^{256}	24	0	1	7	23	Forward+Nonsub+Nonlinear	47