Revisiting Linear Subspace Trails in Poseidon/Poseidon2 and Neptune

Enyan Li and Gaoli Wang

Software Engineering Institute, East China Normal University

Shanghai, China

lienyan6@gmail.com, glwang@sei.ecnu.edu.cn

Abstract

Algebraic attacks are an important class of cryptanalytic techniques for Poseidon/Poseidon2 and Neptune. These designs use partial S-box activation in internal rounds to reduce arithmetization cost, and this structure makes linear subspace trails an important technique in their algebraic analysis. If the state at some internal round is restricted to a suitable linear subspace, then several subsequent internal rounds can be linearized and therefore do not increase the degree of the polynomial system used to model the attack. Recent studies have explored such linearization and round-skipping techniques for algebraic attacks on Poseidon/Poseidon2 and Neptune. These techniques reduce the complexity of attacks by lowering the degree of the polynomial systems that model the attacked rounds. Consequently, a precise bound on the length of linear subspace trails is needed for a more accurate assessment of the algebraic security of these designs.

We revisit infinite and finite linear subspace trails in Poseidon-like designs. First, motivated by previous work that relates infinitely long linear subspace trails to invariant subspaces of the linear layer, we revisit this phenomenon for the Cauchy MDS matrices used in Poseidon. We further give a quantitative heuristic estimate for the probability that such invariant subspace conditions occur under random parameter choices. Second, we analyze finite linear subspace trails for partial rounds in a general setting, where the state width is t and each round activates s S-box coordinates. Under the rank growth condition stated in this paper, when no such invariant subspace exists, a finite trail has length at most $\lceil t/s \rceil - 1$. For Poseidon, $s = 1$, this gives at most $t - 1$ consecutive linearized internal partial rounds. For the internal linear layers of Poseidon2 and Neptune, a similar conclusion applies. More precisely, if repeated diagonal entries in the lower-right block do not give rise to infinitely long linear subspace trails, then the maximum length of a finite linear subspace trail is exactly $t - 1$. Third, for preimage attacks in sponge mode with rate r, capacity c, and digest size d, the available extra constraint budget is $Ec = r - \min\{c, d\}$. For Poseidon/Poseidon2 and Neptune, the finite trail bound and this constraint budget together determine how many internal partial rounds can be linearized in the corresponding attack model.

Keywords: Poseidon/Poseidon2; Neptune; arithmetization-oriented primitives; algebraic cryptanalysis; linear subspace trails; CICO.

1 Introduction

1.1 Background and Motivation

Arithmetization-friendly hash functions are designed to admit low-cost arithmetic circuit implementations over large prime fields. This requirement is central in zero-knowledge proof systems, where a hash function is implemented as an arithmetic circuit rather than through bit-oriented operations. In applications such as SNARKs, STARKs, verifiable computation, and related zero-knowledge systems, the relevant cost is largely determined by the number of finite field multiplications, some of which arise from representing the nonlinear layers of the hash function. This motivates designs that reduce the number of S-box evaluations while using linear layers to provide diffusion. Poseidon/Poseidon2 and Neptune are representative designs of this type [14, 16, 17]. This line of designs has also been extended to binary field settings, leading to PoseidonB and Poseidon2B [22]. Ongoing efforts to standardize such designs in blockchain and zero-knowledge ecosystems, including the Ethereum proposal for a Poseidon precompile [4], further illustrate the interest in their practical deployment. Poseidon-like designs follow the HADES strategy, using external rounds with full S-box layers and internal rounds with partial S-box activation. This design reduces arithmetization cost, but the sparse use of S-boxes also enables linear subspace trails in algebraic attacks. When the state at some internal round is restricted to a suitable linear subspace, several subsequent partial rounds can be linearized and hence do not increase the degree of the resulting polynomial system [9].

Previous work has studied how to choose linear layers to exclude infinitely long subspace trails. Furthermore, finite linear trails have been used to improve algebraic attacks on Poseidon/Poseidon2 and Neptune [15, 18]. These results motivate a more direct study of the possible length of finite linear subspace trails. We address this question for a generalized Poseidon-like partial round structure, in which each internal round activates s S-box coordinates rather than the single active coordinate used in Poseidon/Poseidon2 and Neptune.

1.2 Related Work

Subspace attacks play an important role in the cryptanalysis of HADES-like hash functions with partial S-box activation in their internal rounds. A linear layer may permit subspace trails, along which the state can be linearized over several rounds. As a consequence, the corresponding rounds do not increase the degree of the resulting algebraic system. This idea was used against HadesMiMC and has been further developed in subsequent analyses [9, 13, 21].

For Poseidon-like designs, linear subspace trails appear from two different perspectives. From the design perspective, infinitely long trails must be excluded. They correspond to invariant subspaces arising from the internal linear layer. This requirement leads to algebraic conditions involving powers of the linear layer and its minimal polynomial structure [15]. From the attack perspective, excluding infinitely long trails does not remove finite linear trails. Such trails may still linearize several internal rounds and thereby reduce the degree of the polynomial system used in a preimage attack. The original Poseidon analysis already incorporated such considerations [14], and a later revisit of Poseidon/Poseidon2 and Neptune gave more refined complexity estimates for attacks based on linear subspace trails [18].

Faugère and Perret analyzed Gröbner basis attacks on STARK-friendly ciphers and hash functions such as Jarvis and Friday, and estimated the degree of the final univariate polynomial obtained from the lexicographic Gröbner basis [12]. A related line of work focuses on round-skipping in the algebraic modeling of such primitives. Bariant et al. showed that careful modeling of the round function can improve algebraic attacks on arithmetization-oriented primitives. For Poseidon, they observed that the first round can be skipped in the algebraic attack [6]. Merz and Rodríguez García later extended this round-skipping viewpoint to Poseidon2 and Poseidon2b. Using the special structure of the external and internal linear layers, they showed that the algebraic modeling can skip several initial external rounds and, in some cases, a few subsequent internal rounds. This further reduces the complexity of algebraic attacks on Poseidon2 and Poseidon2b [25].

Several works study Poseidon and related arithmetization-friendly primitives through other algebraic cryptanalysis techniques. Ashur et al. revisited Gröbner basis modeling for such designs and observed that some parameter regimes may lead to optimistic security estimates [1]. More recent works study how the choice of monomial order affects the dominant computational cost, including FreeLunch and CheapLunch methods [2, 7]. Other recent analyses include improvements based on resultant [8, 26] and Graeffe approaches [27]. Koschatko et al. also provided a more systematic complexity analysis of Gröbner basis attacks on Anemoi, offering a useful reference point for similar analyses of other arithmetization-oriented primitives [23]. More recently, Bak et al. proposed resultant-based CICO-1 and CICO-2 attack frameworks for arithmetization-oriented permutations with α-inverse S-boxes, such as Anemoi, Rescue, and Griffin [3]. Their method improves the solving step for several concrete CICO-1 and CICO-2 instances.

Building on the framework of linear subspace trails and invariant subspace conditions developed in [15], this paper revisits the internal round structure of Poseidon/Poseidon2 and Neptune from this perspective. We first re-examine the algebraic conditions under which infinitely long linear subspace trails may arise. We then provide a more explicit and detailed analysis of finite linear subspace trails, in a generalized setting where each internal partial round activates s S-box coordinates. The resulting bounds clarify how many internal rounds can be linearized in algebraic attacks on Poseidon-like designs.

1.3 Our Contributions

The contributions of this paper are threefold.

1. **Invariant subspace conditions for infinitely long linear trails.** We revisit the algebraic conditions under which infinitely long linear subspace trails can arise for the MDS layers used in the Poseidon design. For the Cauchy MDS matrices recommended in the design, we reformulate the relevant invariant subspace condition in terms of their characteristic polynomials. We then give a qualitative discussion of why the required polynomial condition is unlikely over fields of large characteristic, by comparing it with the behavior expected for random monic polynomials. This complements the existing criteria for excluding infinitely long subspace trails [15].

2. **A finite trail bound and its realization under preimage attacks.** We formulate the linear trail analysis for the general case where each internal partial round activates s S-box coordinates. Under the rank growth condition made explicit later in the paper, and excluding the existence of infinitely long trails, we show that a

finite linear trail has length at most $\lceil t/s \rceil - 1$. For the Poseidon setting $s = 1$, this gives at most $t - 1$ consecutive linearized internal partial rounds.

We further relate this structural bound to the additional linear constraints available in preimage attacks. In sponge mode with rate r, capacity c, and digest size d, this budget is $Ec = r - \min\{c, d\} = t - \min\{2c, c + d\}$. Since each linearized internal partial round consumes s such constraints, the realizable trail length in the corresponding preimage problem is also bounded by $\lfloor Ec/s \rfloor$. Thus, for Poseidon, the number of internal partial rounds that can actually be linearized is jointly limited by the finite trail bound induced by the MDS matrix and the constraint budget depending on the sponge mode.

3. **Linear trail analysis for the internal layers of Poseidon2 and Neptune.** We extend the preceding analysis from the MDS layers used in Poseidon to the internal linear layers used in Poseidon2 and Neptune. For these layers, the existence of infinitely long linear subspace trails is governed in a more explicit way by the diagonal parameters of the internal matrix. In particular, repeated diagonal entries give rise to persistent invariant directions and hence to infinitely long linear subspace trails. After excluding the case of repeated diagonal entries, the finite trail behavior becomes more regular: in the usual setting $s = 1$, the maximum length of a finite linear subspace trail is exactly $t-1$, compatible with the conclusion obtained for Poseidon.

1.4 Organization

The rest of the paper is organized as follows. Section 2 introduces the notation for Poseidon/Poseidon2 and Neptune and recalls the traditional CICO formulation. Section 3 studies the algebraic conditions corresponding to infinitely long and finite linear subspace trails. Section 4 extends the preceding conclusions, which are mainly developed for the MDS linear layer of Poseidon, to the linear layers used in the internal rounds of Poseidon2 and Neptune. Section 5 concludes the paper.

2 Preliminaries

2.1 The Structure of Poseidon and Neptune

Poseidon/Poseidon2 and Neptune use sponge mode [14, 16, 17]. In addition, Poseidon2 is also allowed in compression mode, although the only formally specified compression instance has state size $t = 3$ [17, 18].

If the underlying permutation used in sponge mode is modeled as a random permutation, then a sponge mode with capacity c is indifferentiable from a random oracle up to about $p^{c/2}$ queries [10]. Hence κ-bit security usually requires $p^{c/2} \geq 2^{\kappa}$. If the digest consists of d field elements, birthday resistance also requires $p^{d/2} \geq 2^{\kappa}$. Here c denotes the capacity, and d denotes the digest length. Beyond these generic sponge security requirements, sponge hash security can also be phrased through the CICO problem for the underlying permutation [11]. The next subsection recalls the traditional CICO formulation.

Poseidon/Poseidon2 and Neptune are arithmetization-oriented permutations over a finite field $\mathbb{F}_p$. We denote the state size by t, and write a state as $x = (x_0, x_1, \ldots, x_{t-1}) \in$

$\mathbb{F}_p^t$. Poseidon/Poseidon2 and Neptune follow the HADES structure. They consist of full rounds at the beginning and end of the permutation, with partial rounds in the middle. In the unified notation of [18], the permutation 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. The maps E_i are external full rounds, and the maps I_i are internal partial rounds. The total number of full rounds is $r_F = r_f + r'_f$. The first r_f full rounds come before the partial rounds, and the last r'_f full rounds come after them. The number of internal partial rounds is denoted by r_P.

In Poseidon and Poseidon2, an external full round applies a power S-box to every coordinate and then applies an external linear layer:

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

Here M_E is the external linear layer, c_i are round constants, and α is the S-box degree. Usually α is a small positive integer with $\gcd(\alpha, p-1) = 1$. Internal partial rounds apply nonlinear S-boxes only to a few coordinates. For the discussion in Section 3, we consider the general case in which each internal partial round activates the first s coordinates.

The standard Poseidon and Poseidon2 instances considered here use $s = 1$. In this case, a partial round has the form

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

Here M_I is the internal linear layer and $c' = (c'_0, \ldots, c'_{t-1}) \in \mathbb{F}_p^t$ is a fixed round constant vector.

Poseidon and Poseidon2 differ in their choice of linear layers. Poseidon takes M' to be the identity matrix, while the external and internal linear layers use an MDS matrix. The concrete Poseidon instances considered here use Cauchy matrices for this MDS layer [14]. To rule out infinitely long subspace trails, including trails with bounded period, the linear layer must satisfy additional algebraic conditions. Grassi et al. [15] give a sufficient condition for excluding infinitely long subspace trails of period at most ℓ: the minimal polynomials of $M, M^2, \ldots, M^\ell$ should all have maximal degree t and be irreducible over the base field.

Poseidon2 instead separates the external linear layer M_E from the internal linear layer M_I. The specified state widths start with the small cases $t = 2, 3$, while the remaining choices are multiples of 4. For the small-width cases $t = 2$ and $t = 3$, the designers set the external matrix equal to the internal matrix after requiring the latter to be MDS. For $t = 4t' \leq 24$ with $t' \in \mathbb{N}$, the external matrix is built from a fixed 4×4 MDS matrix by a Kronecker product construction [17, 25]. The internal matrix has the form

$$M_I = \Lambda + \mathbf{1}\mathbf{1}^\mathsf{T},$$

where Λ is diagonal and $\mathbf{1}$ denotes the all-one column vector. Thus, unlike Poseidon, Poseidon2 uses different structured matrices in the external and internal rounds. In particular, the internal layer still has to resist infinitely long subspace trails [17].

Neptune is defined for even state sizes $t = 2t' \leq 24$. Its initial linear layer M' and its internal round linear layer are the same as those of Poseidon2, while the nonlinear layer in the external full rounds is changed [16, 18]. Instead of applying a power map to

every coordinate separately, NEPTUNE groups the state into pairs and applies a bivariate map $S : \mathbb{F}_p^2 \to \mathbb{F}_p^2$ to each pair. Thus its external full round 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 the Lai–Massey structure [24], and satisfies $\deg(S) = \deg(S^{-1}) = 4$.

We finally recall the recommended round numbers for POSEIDON/POSEIDON2 and NEPTUNE. These parameters serve only as background for the concrete instances discussed later. For POSEIDON/POSEIDON2, under security level κ bits and $2^{80} \leq 2^\kappa \leq \min\{2^{256}, p^{t/3}\}$, the recommendation is $r_F = 6$, namely $r_f = 3$ and $r'_f = 3$, and [18]

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

Here α is the S-box degree. The first term protects against interpolation attacks [20]. The second term concerns Gröbner basis attacks. More precisely,

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

The second term in this maximum protects against GB attacks exploiting linear subspace trails. The term r'_{GB} corresponds to the analysis of [18]. The final recommended parameters also include a security margin: two additional external full rounds and 7.5% more internal partial rounds. This raises $r_F = 6$ to $r_F = 8$.

For NEPTUNE, under the same security condition $2^{80} \leq 2^\kappa \leq \min\{2^{256}, p^{t/3}\}$, the recommendation is $r_F = 6$, namely $r_f = 4$ and $r'_f = 2$, and [16, 18]

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

Here α is the degree of the internal power S-box. The first term protects against interpolation attacks. Since an invariant subspace trail may cover t internal rounds, this term includes an additional t internal rounds [16]. NEPTUNE's security margin consists of 12.5% additional internal partial rounds. It does not add external full rounds.

2.2 The Traditional CICO Problem

The degrees of freedom analysis in this paper starts from the traditional CICO (constrained-input-constrained-output) problem. Let $P : \mathbb{F}_p^t \to \mathbb{F}_p^t$ be an invertible permutation. Let n_I and n_O be positive integers such that $0 < n_I, n_O < t$ and $n_I + n_O \leq t$. The traditional CICO(n_I, n_O) problem asks for 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. Thus the attacker must satisfy coordinate constraints at both the input and output sides.

For the discussion below, we take the fixed coordinates to be zero. Since our analysis only concerns how many linear constraints are available and how they are constructed, their concrete fixed values are irrelevant.

At the mode level, CICO(n_I, n_O) gives a common abstraction for sponge preimage attacks and compression function attacks. Input side constraints may come from capacity coordinates that are fixed in the mode. Output side constraints come from digest coordinates, capacity coordinates, or feedforward conditions. Section 3 further separates

the positions of the output constraints. Constraining digest coordinates gives the CICD case, while constraining capacity coordinates gives the CICC case. This distinction is essential for computing the extra constraint budget Ec, because the two types of output constraints may leave different numbers of degrees of freedom to the attacker.

3 Linear Subspace Trails for the Poseidon MDS Layer

We briefly recall the CICO formulation of preimage problems in sponge mode, which will serve as the basis for the later discussion. A sponge mode preimage problem is a CICO problem in which the input capacity is fixed. We call the case where the prescribed output coordinates are digest coordinates a constrained-input-constrained-digest (CICD) problem, and the case where the prescribed output coordinates are capacity coordinates a constrained-input-constrained-capacity (CICC) problem. Let Oc denote the number of prescribed output constraints. Thus $\mathrm{Oc} = d$ for CICD and $\mathrm{Oc} = c$ for CICC. Starting from the $r = t - c$ free rate variables, the attacker may further impose extra constraints in order to construct a linear subspace trail. We denote their number by Ec. Under the usual semi-regularity assumption [5], after imposing Oc output equations and Ec additional equations, the expected dimension of the ideal is $r - \mathrm{Oc} - Ec$. Hence, while keeping the system zero-dimensional without overconstraining it, the available extra constraint budget is $Ec = r - \mathrm{Oc}$. When additional absorbing steps are allowed in the sponge preimage model, a CICD instance can be rewritten as a CICC instance: one guesses the capacity of a later state, inverts one permutation call, and uses two message absorptions to absorb the resulting rate difference. Therefore, for a preimage problem in sponge mode, one may always choose the output model with fewer constraints, so that $\mathrm{Oc} = \min\{c, d\}$ and

$$Ec = r - \min\{c, d\} = t - \min\{2c, c + d\}.$$

Compression mode has no capacity part and does not admit this inversion based reduction. Thus, for a preimage problem with digest length d, the corresponding budget is simply $Ec = t - d$. In summary, the extra constraint budget relevant to sponge mode preimage attacks is $r - \min\{c, d\}$ rather than the more restrictive quantity $r - d = t - (c + d)$.

3.1 An Algebraic Form of Linear Subspace Trails

Grassi et al. [15] related subspace trails of SPN hash functions over $\mathbb{F}_p^t$ to minimal polynomial conditions of the linear layer and its powers. In particular, their analysis relates infinitely long subspace trails to the existence of a nontrivial invariant subspace of the linear layer contained in the inactive coordinate subspace. Motivated by this viewpoint, this section specializes the discussion to Poseidon-like hash functions and reformulates the relevant trail conditions through a block matrix formulation. For infinitely long trails, we analyze the corresponding invariant subspace conditions and, under a random choice for the linear layer, estimate how unlikely these conditions are to occur. For finite trails, we give a more explicit algebraic analysis of the rank growth constraints that determine their possible length.

We start with the block form of the linear layer in one partial round. Let r_p denote a single partial round map, and reserve r_P for the total number of partial rounds. Write the internal linear layer as

$$M_I = \begin{pmatrix} A & B \\ C & D \end{pmatrix} \in \mathbb{F}_p^{t \times t},$$

where M_I is an MDS matrix over $\mathbb{F}_p$. The blocks have sizes $A \in \mathbb{F}_p^{s\times s}$, $B \in \mathbb{F}_p^{s\times(t-s)}$, $C \in \mathbb{F}_p^{(t-s)\times s}$, and $D \in \mathbb{F}_p^{(t-s)\times(t-s)}$. Thus A, B, C, D are the four natural blocks induced by the first s coordinates and the remaining $t-s$ coordinates. We write an intermediate state as $X = (x, X')^\mathsf{T}$, where $x \in \mathbb{F}_p^{s\times 1}$ and $X' \in \mathbb{F}_p^{(t-s)\times 1}$. In each partial round, the first s coordinates are the coordinates to which the α-power S-box is applied.

Suppose that, at the input of a partial round, the first s coordinates are constants. Then the active S-boxes act only on constants and do not increase the degree of the variable-dependent system. Write such a state as $Y = (c_0, \ldots, c_{s-1}, g_1(X'), \ldots, g_{t-s}(X'))^\mathsf{T}$, where $c_0, \ldots, c_{s-1} \in \mathbb{F}_p$ and the g_i are polynomials in X'. For each i, write $g_i(X') = \widetilde{g}_i(X') + \gamma_i$, where $\gamma_i \in \mathbb{F}_p$ and $\widetilde{g}_i$ has zero constant term. Then Y is the sum of the variable-dependent vector $(0, \ldots, 0, \widetilde{g}_1(X'), \ldots, \widetilde{g}_{t-s}(X'))^\mathsf{T}$ and a constant vector. Since linear layers and round constants act affinely, these two parts can be tracked separately. Hence, as long as the system follows a linear subspace trail, we may absorb round constants into the constants γ_i and omit them from the variable-dependent subsystem.

The same separation can be repeated over several consecutive partial rounds. Thus, when linear constraints maintain a trail, we may assume without loss of generality that their constant terms are zero. Choosing different constants only changes the inhomogeneous terms of the corresponding affine constraints and does not affect the possible trail length. In what follows, we therefore analyze homogeneous linear constraints and omit the round constants from the variable-dependent part of the system.

Next, we recall a standard criterion for MDS matrices.

Lemma 1 ([19])**.** *Let $M \in \mathbb{F}_p^{m\times m}$. Then M is MDS if and only if every square submatrix of M has full rank.*

This immediately gives the following consequence.

Corollary 1. *Let M be an MDS matrix over $\mathbb{F}_p$. Then every square submatrix of M is itself an MDS matrix.*

In particular, the lower-right block D in the above decomposition is an MDS matrix. This fact will be used later.

We now restate linear subspace trails in elementary linear algebraic terms [15, 18]. After removing the constant part, a partial round first applies the α-power map to the first s coordinates and then applies the linear layer M. We call a trail maintained by linear constraints in this setting a linear subspace trail.

Proposition 1 (linear algebraic form of linear subspace trails)**.** *After separating off the constant part, let the variable-dependent part of a partial round input be $X = (x, X')^\mathsf{T}$, and let S_i be the i-th linearized state. For any $\ell \geq 0$, a linear subspace trail of length $\ell+1$ maintained by $\ell+1$ groups of linear constraints exists if and only if*

$$x = 0, \qquad BD^iX' = 0 \quad (0 \leq i \leq \ell - 1),$$

where the second set of conditions is empty when $\ell = 0$.

The constraint $x = 0$ makes the first s coordinates of S_0 constant. The remaining ℓ groups of constraints force $S_1, \ldots, S_\ell$ to stay in the same linear subspace. Under these conditions, $S_i = (0, D^iX')^\mathsf{T}$ for $0 \leq i \leq \ell$, and the next state is $S_{\ell+1} = (BD^\ell X', D^{\ell+1}X')^\mathsf{T}$. This state is still linear in the variables. If $BD^\ell X' \neq 0$, then the first s coordinates of $S_{\ell+1}$ contain linear terms, and the partial S-boxes are activated in the next internal round.

Moreover, an infinitely long linear subspace trail exists if and only if $\ker(B)$ contains a nontrivial D-invariant subspace.

Proof. The first partial S-box introduces no variable degree only if $x = 0$. Under this condition, the power map has no effect on the variable-dependent part in the first s coordinates. The linear layer gives $M(0, X')^\mathsf{T} = (BX', DX')^\mathsf{T}$. The next state remains in the subspace with zero first block if and only if $BX' = 0$. In that case, $S_1 = (0, DX')^\mathsf{T}$.

Now assume inductively that $S_i = (0, D^iX')^\mathsf{T}$. After one more partial round, the variable-dependent part becomes $(BD^iX', D^{i+1}X')^\mathsf{T}$. Therefore the next state remains in the same linear subspace if and only if $BD^iX' = 0$. This proves the finite constraint chain.

For an infinite trail, we need $BD^iX' = 0$ for all $i \geq 0$. Equivalently, the cyclic subspace $\langle X', DX', D^2X', \ldots \rangle$ is contained in $\ker(B)$. This cyclic subspace is D-invariant. Conversely, if $\ker(B)$ contains a nonzero D-invariant subspace W, then any nonzero $X' \in W$ satisfies $D^iX' \in W \subseteq \ker(B)$ for all $i \geq 0$. This gives an infinitely long linear subspace trail. □

Throughout this discussion, all imposed linear constraints are homogeneous. The reason is the separation above: constant terms can be processed independently in the constant subsystem. Therefore setting the right-hand side to zero does not affect the length of the linear subspace trail. Proposition 1 converts the cryptographic trail condition into a linear algebraic condition. Finite trails correspond to the constraint chain $BD^iX' = 0$. Infinite trails correspond to the existence of a nontrivial D-invariant subspace inside $\ker(B)$.

3.2 Algebraic Conditions for Infinitely Long Linear Subspace Trails

3.2.1 Invariant Subspaces

We now describe this condition more explicitly. Let $V = \mathbb{F}_p^{t-s}$, and regard V as an $\mathbb{F}_p[x]$-module where the polynomial variable x acts through the linear map D. A subspace $W \subseteq V$ is D-invariant if and only if W is a submodule of this $\mathbb{F}_p[x]$-module. Thus the invariant subspace structure of D is determined by the structure of V as a finitely generated $\mathbb{F}_p[x]$-module. Equivalently, it is determined by the invariant factors in the rational canonical form.

The primary decomposition theorem first gives a coarse decomposition from the coprime factors of the minimal polynomial. If $m_D(x) = \prod_i q_i(x)^{e_i}$, where the q_i are pairwise coprime and irreducible, then $V = \bigoplus_i V_i$, where $V_i = \ker q_i(D)^{e_i}$ is a D-invariant primary component. The minimal polynomial records only the corresponding irreducible factors and their largest powers. The sizes and multiplicities of the rational blocks inside each V_i, and hence the finer invariant subspace structure, are still determined by the invariant factors or elementary divisors. When the characteristic polynomial is squarefree, this structure becomes simpler: the primary components can be read directly from the irreducible factors of the characteristic polynomial. We will use the following elementary fact repeatedly.

Lemma 2. *Let A be a linear map on a finite dimensional vector space V, and let $f(x) \in \mathbb{F}_p[x]$. Then both $\ker(f(A))$ and $\mathrm{Im}(f(A))$ are A-invariant subspaces.*

Proof. Take $\eta \in \ker(f(A))$. Then $f(A)\eta = 0$, and hence $f(A)(A\eta) = Af(A)\eta = 0$, because A commutes with $f(A)$. Thus $A\eta \in \ker(f(A))$. Similarly, if $\xi \in \mathrm{Im}(f(A))$, then $\xi = f(A)v$ for some $v \in V$. Therefore $A\xi = Af(A)v = f(A)Av \in \mathrm{Im}(f(A))$. □

Let $f_D(\lambda)$ be the characteristic polynomial of D, so that $\deg f_D = \dim V = t - s$. If f_D is irreducible over $\mathbb{F}_p$, then the minimal polynomial of D equals f_D. Hence V is cyclic as an $\mathbb{F}_p[\lambda]$-module and $V \cong \mathbb{F}_p[\lambda]/(f_D)$. Since (f_D) is a maximal ideal, this quotient is a field and therefore has no proper nonzero $\mathbb{F}_p[\lambda]$-submodules. Equivalently, D has no proper nonzero invariant subspace. We now turn to the base field root case, which is the case needed below.

One simple root. Assume that $f_D(\lambda)$ has exactly one linear factor over the base field and has no repeated root, namely $f_D(\lambda) = (\lambda - \lambda_*)g(\lambda)$, where $g(\lambda)$ is coprime to $\lambda - \lambda_*$. By the Cayley–Hamilton theorem, $(D - \lambda_* I)g(D) = 0$. Hence $\mathrm{Im}(g(D)) \subseteq \ker(D - \lambda_* I)$ and $\mathrm{Im}(D - \lambda_* I) \subseteq \ker(g(D))$. Conversely, Bezout's identity gives polynomials $a(\lambda), b(\lambda) \in \mathbb{F}_p[\lambda]$ such that $a(\lambda)(\lambda - \lambda_*) + b(\lambda)g(\lambda) = 1$. Substituting D gives $a(D)(D - \lambda_* I) + b(D)g(D) = I$. If $v \in \ker(D - \lambda_* I)$, then $v = b(D)g(D)v \in \mathrm{Im}(g(D))$. If $v \in \ker(g(D))$, then $v = a(D)(D - \lambda_* I)v \in \mathrm{Im}(D - \lambda_* I)$. Therefore

$$\mathrm{Im}(g(D)) = \ker(D - \lambda_* I), \qquad \mathrm{Im}(D - \lambda_* I) = \ker(g(D)).$$

By Lemma 2, both spaces are D-invariant. Since λ_* is a simple root, $\dim \ker(D - \lambda_* I) = 1$. The rank-nullity theorem gives $\dim \mathrm{Im}(D - \lambda_* I) = t - s - 1$. Thus D yields two natural invariant subspaces: the one-dimensional eigenspace $W_1 = \ker(D - \lambda_* I) = \mathrm{Im}(g(D))$, and the codimension-one invariant subspace $W_2 = \mathrm{Im}(D - \lambda_* I) = \ker(g(D))$.

Multiple roots and repeated roots. If f_D has several distinct simple roots over the base field, then each root gives a one-dimensional eigenspace. These eigenspaces, and their direct sums, are D-invariant. They may therefore become candidates for infinitely long linear subspace trails. If repeated roots occur, the characteristic polynomial alone no longer determines the invariant subspace structure. The corresponding elementary divisors are then needed. However, the experiments and estimates below indicate that repeated root events have negligible probability for the large characteristic parameters considered in this paper.

Since M is MDS and $t \geq 2s$, every $s \times s$ submatrix of B has full rank. It follows that $\mathrm{rank}(B) = s$ and $\dim \ker(B) = (t - s) - s = t - 2s$. Therefore an infinitely long trail can occur only if one of the invariant subspaces generated by D is contained entirely in $\ker(B)$. The next subsections explain this event through a random characteristic polynomial model and experimental evidence.

3.2.2 Random Cauchy Matrices and Random Monic Polynomials

The invariant subspaces of the square matrix D are controlled by its invariant factors and elementary divisors. From the polynomial viewpoint, they are closely related to how the characteristic polynomial factors over the base field. We therefore study the possible invariant subspaces of D through the coefficients and base field roots of the characteristic polynomials of MDS matrices.

Assume that the Cauchy matrix used in Poseidon is generated at random, and let its order be n. Here random generation means that the defining elements are sampled uniformly subject to the usual Cauchy admissibility conditions: the x_i's and y_j's are distinct and all $x_i - y_j$ are nonzero. The differences $x_i - y_j$ generated inside the matrix are uniformly distributed over $\mathbb{F}_p \setminus \{0\}$. Although these differences are correlated, each

individual difference is mapped by inversion bijectively to $\mathbb{F}_p \setminus \{0\}$, which gives the entries of the Cauchy matrix.

We next compare random Cauchy matrices with random monic polynomials through three finite field experiments: entry distributions of the matrices, coefficient distributions of their characteristic polynomials, and base field root-count distributions. All experiments below are over $\mathbb{F}_{101}$. The Cauchy matrix order ranges from 2 to 12, and each setting uses trials = 1000. The following tables and figures summarize the resulting data. For root counts, we compare the observed distributions with the truncated Poisson(1) model, whose relevance will be explained after the experimental results.

Table 1: Chi-square tests over $\mathbb{F}_{101}$ for the entry distribution of random Cauchy matrices against the uniform distribution on $\mathbb{F}_p \setminus \{0\}$

n	χ^2	p-value
2	88.650	0.7627
3	82.333	0.8869
4	90.963	0.7053
5	108.912	0.2329
6	71.844	0.9818
7	76.192	0.9570
8	93.853	0.6273
9	74.721	0.9673
10	92.538	0.6635
11	70.896	0.9853
12	82.478	0.8847

Table 2: Uniformity tests over $\mathbb{F}_{101}$ for the coefficients of characteristic polynomials of random Cauchy matrices

n	degree	χ^2	p-value
2	0	123.726	0.0540
2	1	93.426	0.6657
3	0	98.678	0.5186
3	1	108.374	0.2666
3	2	107.364	0.2893
4	0	128.978	0.0271
4	1	137.664	0.0075
4	2	61.308	0.9992
4	3	72.216	0.9836
5	0	133.220	0.0148
5	1	117.464	0.1120
5	2	139.280	0.0058
5	3	102.112	0.4225
5	4	98.274	0.5301
6	0	91.608	0.7134
6	1	98.274	0.5301
6	2	93.426	0.6657
6	3	106.354	0.3131
6	4	95.648	0.6045
6	5	125.544	0.0429
7	0	85.144	0.8554
7	1	99.688	0.4900
7	2	97.870	0.5416
7	3	92.012	0.7030
7	4	105.546	0.3328
7	5	95.042	0.6214
7	6	80.700	0.9216

n	degree	χ^2	p-value
8	0	112.616	0.1831
8	1	105.748	0.3279
8	2	89.588	0.7630
8	3	78.478	0.9451
8	4	92.618	0.6872
8	5	110.192	0.2283
8	6	108.172	0.2710
8	7	127.160	0.0347
9	0	98.678	0.5186
9	1	99.284	0.5014
9	2	103.324	0.3899
9	3	138.472	0.0066
9	4	87.770	0.8038
9	5	108.576	0.2621
9	6	105.950	0.3229
9	7	96.052	0.5931
9	8	114.636	0.1503
10	0	115.040	0.1443
10	1	99.082	0.5071
10	2	97.466	0.5531
10	3	111.808	0.1974
10	4	102.112	0.4225
10	5	109.990	0.2324
10	6	90.194	0.7485
10	7	87.972	0.7994
10	8	120.292	0.0816
10	9	113.222	0.1728

n	degree	χ^2	p-value
11	0	118.272	0.1025
11	1	85.750	0.8443
11	2	98.476	0.5244
11	3	100.294	0.4729
11	4	82.518	0.8977
11	5	106.152	0.3180
11	6	108.778	0.2578
11	7	93.022	0.6765
11	8	101.506	0.4392
11	9	110.394	0.2243
11	10	98.274	0.5301
12	0	115.040	0.1443
12	1	107.162	0.2940
12	2	90.194	0.7485
12	3	109.788	0.2365
12	4	95.446	0.6101
12	5	78.276	0.9469
12	6	94.436	0.6382
12	7	106.354	0.3131
12	8	83.124	0.8888
12	9	89.790	0.7582
12	10	69.186	0.9919
12	11	96.254	0.5874

Table 1 shows that the entries of randomly generated Cauchy matrices are compatible with the uniform distribution on $\mathbb{F}_p \setminus \{0\}$. The corresponding p-values provide no sufficient evidence to reject this uniformity hypothesis.

Table 3: Chi-square tests over $\mathbb{F}_{101}$ of root-count distributions of random Cauchy matrices and random monic polynomials against the truncated Poisson(1) distribution

n	Cauchy χ^2	Cauchy p-value	random monic χ^2	random monic p-value
2	776.205	0.0000	760.725	0.0000
3	398.152	0.0000	347.461	0.0000
4	101.677	0.0000	89.909	0.0000
5	37.354	0.0000	24.634	0.0002
6	16.642	0.0107	12.348	0.0546
7	8.543	0.2872	9.087	0.2465
8	1.857	0.9851	5.869	0.6619
9	7.922	0.5420	7.096	0.6271
10	1.972	0.9965	16.273	0.0921
11	6.892	0.8077	12.900	0.2999
12	3.886	0.9854	3.807	0.9867

Table 2 compares the coefficients of characteristic polynomials of random Cauchy matrices with the coefficients of random monic polynomials. Some degrees have relatively small p-values, reflecting nonlinear dependencies among characteristic polynomial coefficients. Still, most p-values are well above 0.05. Overall, the experiment does not easily distinguish characteristic polynomials of random Cauchy matrices from random monic polynomials at the level of individual coefficient distributions.

Figure 1 and Table 3 show the base field root-count statistics for characteristic polynomials of random Cauchy matrices and random monic polynomials of the same degree. Although Cauchy matrices are not random square matrices with independent entries, their characteristic polynomials exhibit root-count behavior very close to that of random monic polynomials. The plotted curve is the truncated Poisson distribution with parameter 1. For $n = 10, 11, 12$, both classes of polynomials closely follow this curve. This is consistent with the Poisson-type behavior expected for the number of roots of random polynomials. The agreement is already visible over the small test field $\mathbb{F}_{101}$. For the much larger prime fields used in practical Poseidon instances, finite field effects that may appear over $\mathbb{F}_{101}$ are expected to be less significant. Therefore, the random polynomial heuristic is at least as natural in the parameter regime of practical interest. For small n, the deviations from the truncated Poisson(1) distribution are more visible because low degree polynomials satisfy elementary algebraic restrictions; for example, a cubic polynomial cannot have exactly two distinct roots over the base field.

Taken together, Tables 1 and 2, Figure 1, and Table 3 support the two empirical assumptions used below: the coefficients of characteristic polynomials of random Cauchy matrices behave approximately uniformly, and the base field factorization behavior of such characteristic polynomials is close to that of random monic polynomials.

Heuristic assumption. In the probability estimates below, we approximate the root count and simple root behavior of the characteristic polynomial of a random Cauchy matrix D by those of a random monic polynomial of the same degree.

Fix $x_0 \in \mathbb{F}_p$, and let $f(x) = x^n + a_{n-1}x^{n-1} + \cdots + a_0$ be a random monic polynomial of degree n. The condition $f(x_0) = 0$ is a nontrivial linear constraint on the coefficient vector, and hence holds with probability $1/p$. If x_0 is a repeated root, then one also needs $f'(x_0) = 0$. For $n \geq 2$, these two coefficient constraints are linearly independent, so the probability that both hold is $1/p^2$. Let $R(f) = \#\{x_0 \in \mathbb{F}_p : f(x_0) = 0,\ f'(x_0) = 0\}$.

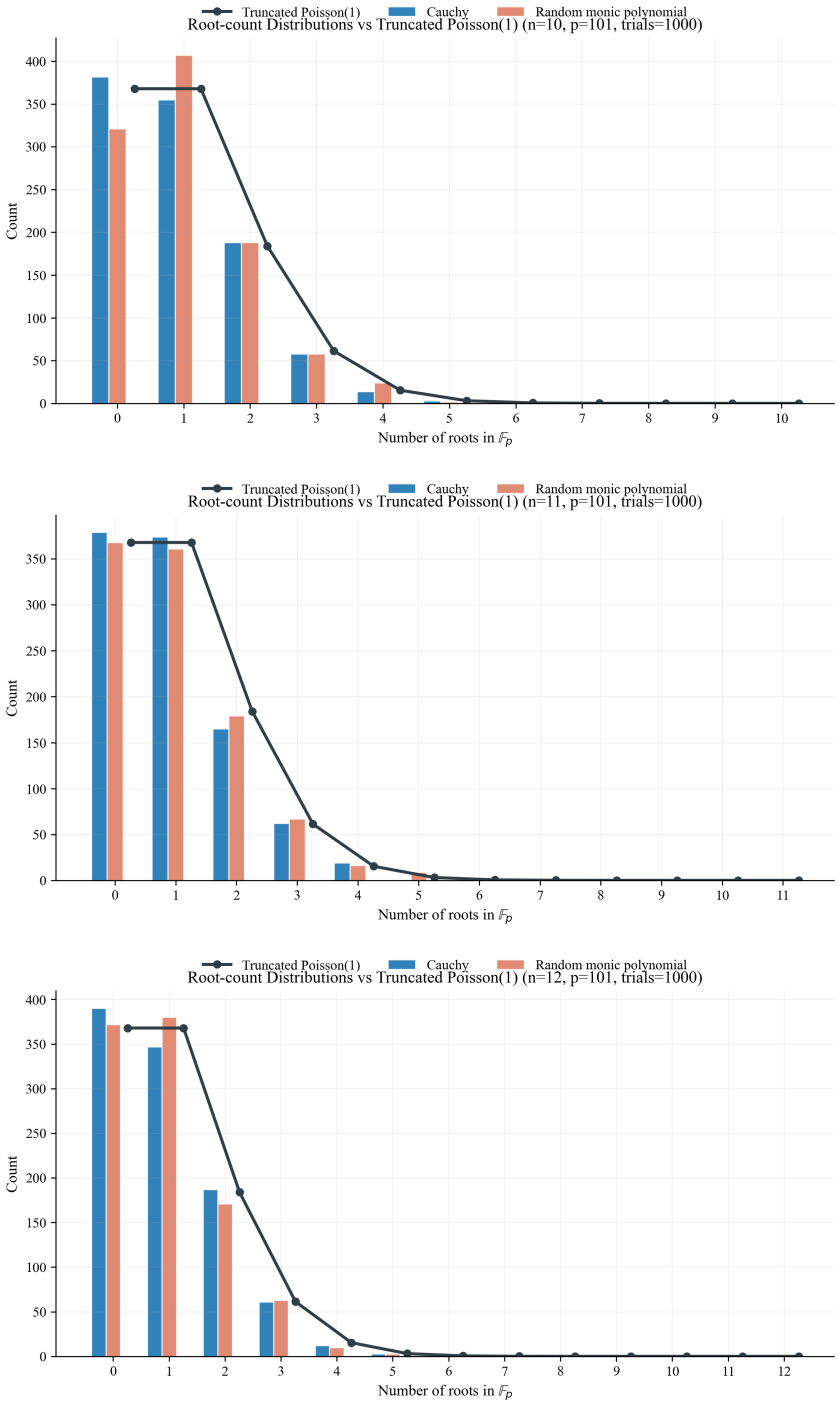

Figure 1: Root-count distributions over $\mathbb{F}_{101}$ for $n = 10, 11, 12$, compared with the truncated Poisson(1) distribution.

Summing over all $x_0 \in \mathbb{F}_p$ gives

$$\mathbb{E}[R(f)] = \sum_{x_0 \in \mathbb{F}_p} \Pr\big(f(x_0) = 0,\ f'(x_0) = 0\big) = p \cdot \frac{1}{p^2} = \frac{1}{p}.$$

Thus, for the practical POSEIDON parameters, the characteristic p is so large that repeated roots can be ignored. We therefore restrict the discussion to squarefree polynomials.

Next consider the distribution of distinct root tuples for a random polynomial. Let $N(f)$ be the number of roots of f in $\mathbb{F}_p$. Then the expected number of roots is $\mathbb{E}[N(f)] = \sum_{a \in \mathbb{F}_p} \Pr(f(a) = 0) = 1$. For pairs of distinct roots, one has

$$\mathbb{E}\left[\binom{N(f)}{2}\right] = \sum_{a<b} \Pr\big(f(a) = 0,\ f(b) = 0\big).$$

Since evaluating at a and b gives two independent linear constraints, the probability in the summand is $1/p^2$. Hence $\mathbb{E}\left[\binom{N(f)}{2}\right] = \binom{p}{2}p^{-2} \approx 1/2$. Similarly, for any fixed k, $\mathbb{E}\left[\binom{N(f)}{k}\right] = \binom{p}{k}p^{-k}$. If we further regard the values obtained by evaluating the characteristic polynomial at different field elements as approximately independent, then the number of base field roots is close to p Bernoulli trials with success probability $1/p$. This is a binomial distribution with parameters $(p, 1/p)$, which converges to Poisson(1) as $p \to \infty$. This agrees with the statistics shown in Figure 1.

In summary, for the POSEIDON instances, the matrix D is extremely unlikely to have repeated eigenvalues. The typical cases are that its characteristic polynomial has no root over the base field, exactly one root, or a small number of distinct roots.

3.2.3 Probability of Infinitely Long Trails in the Random Model

We now return to the probability of infinitely long trails. By Section 3.1, if such a trail exists, then some nontrivial invariant subspace generated by D must be contained in $\ker(B)$. By the heuristic assumption in Section 3.2.2, over very large characteristic fields, the factorization behavior of the characteristic polynomial of a random Cauchy matrix is close to that of a random monic polynomial. Under this empirical assumption, we estimate the probability of an infinite trail event.

We now specialize the preceding discussion to the simple root case of Section 3.2.1, where a base field root of the characteristic polynomial yields a one-dimensional invariant subspace W_1 and a D-invariant subspace W_2 of codimension one. First consider a one-dimensional invariant subspace $W_1 = \langle \eta_0 \rangle$, where $\eta_0 \neq 0$. Let $\ker(B) = \langle \eta_1, \eta_2, \ldots, \eta_{t-2s} \rangle$. The event $W_1 \subseteq \ker(B)$ is equivalent to expressing η_0 linearly in terms of $\eta_1, \ldots, \eta_{t-2s}$. Equivalently, there exists a nonzero coefficient vector $\mathbf{a} \in \mathbb{F}_p^{t-2s} \setminus \{0\}$ such that $(\eta_1, \eta_2, \ldots, \eta_{t-2s})\mathbf{a} = \eta_0$. In the random model, this is an overdetermined linear system in $\mathbf{a}$. For fixed nonzero $\mathbf{a}$ and nonzero η_0, if the row constraints are treated as approximately independent, then each row equation $(\eta_1, \eta_2, \ldots, \eta_{t-2s})_i \mathbf{a} = (\eta_0)_i$ holds with probability about $1/p$. We exclude the trivial case $\eta_0 = 0$. Since η_0 then ranges over $p^{t-s} - 1$ nonzero vectors in V, while there are $p^{t-2s} - 1$ nonzero choices of $\mathbf{a}$, the expected number of nonzero solutions is

$$\mathbb{E}[\#\{\mathbf{a} \neq 0 : (\eta_1, \ldots, \eta_{t-2s})\mathbf{a} = \eta_0\}] = \frac{p^{t-2s} - 1}{p^{t-s} - 1} \approx p^{-s}.$$

Now consider the higher dimensional invariant subspace W_2. When $s = 1$, one has $\dim(W_2) = t - s - 1 = t - 2 = \dim \ker(B)$. Thus $W_2 \subseteq \ker(B)$ forces $W_2 = \ker(B)$. This is stricter than requiring a single eigenvector to lie in $\ker(B)$, so its probability is lower. When $s > 1$, $\dim(W_2) = t - s - 1 > t - 2s = \dim \ker(B)$, and hence $W_2 \subseteq \ker(B)$ is impossible.

If the characteristic polynomial has several distinct linear factors over the base field, then a corresponding one-dimensional eigenspace can yield the required invariant subspace condition only when it is contained in $\ker(B)$. For any fixed eigenspace, the estimate above applies. If there are m such eigenspaces, a union bound gives an upper estimate of at most mp^{-s}.

Thus, for the practical Poseidon parameters, the probability that an infinitely long trail is caused by a cyclic or invariant subspace is expected to be $O(1/p^s)$. This is a very rare event. This also explains why such infinitely long trails are rarely observed in experiments.

3.3 Maximum Length of Finite Linear Subspace Trails

We now consider the finite case: if no infinite trail exists, how far can a finite subspace trail extend? To separate this structural question from the degrees of freedom analysis above, this subsection studies only the longest subspace trail allowed by the Cauchy matrix. It does not ask whether the attacker actually has enough freedom to impose all these linear constraints. Equivalently, we do not yet discuss whether, after adding all possible linear equations in a CICO instance, the resulting ideal still has a root over the base field.

Proposition 2 (Intersection characterization of finite linear subspace trails)**.** *Let $W_i = D^{-(i-1)} \ker(B) \subseteq \mathbb{F}_p^{t-s}$ for $i \geq 1$. For any $r \geq 2$, assuming that the initial constraint $x = 0$ is already imposed, the partial S-boxes in the first r internal partial rounds are not activated by higher degree terms in the variable-dependent system if and only if*

$$X' \in U_r := \bigcap_{i=1}^{r-1} W_i.$$

Here $x = 0$ is the initial subspace constraint on the first s active coordinates. It already guarantees that the partial S-box is not activated for $r = 1$. From $U_2 = W_1$ onward, the spaces W_i describe the successive linear conditions that the inactive coordinate vector X' must satisfy.

Proof. Proposition 1 writes a linear subspace trail as the constraint chain $x = 0$ and $BD^iX' = 0$. Here $x = 0$ gives the first linearized state. The condition $BD^iX' = 0$ is equivalent to $D^iX' \in \ker(B)$, or $X' \in D^{-i} \ker(B)$. For $r \geq 2$, taking $i = 0, \ldots, r-2$ gives $X' \in \bigcap_{i=1}^{r-1} W_i$. Conversely, if X' lies in this intersection, then all these constraints hold simultaneously. Hence the first r internal partial rounds propagate along the same linear subspace trail. □

Lemma 3 (Dimension drop in general position)**.** *Let $U_r = \bigcap_{i=1}^{r-1} W_i$ for $r \geq 2$. Assume that the subspace chain induced by the random Cauchy matrix is in general position, and that no infinitely long invariant subspace occurs. Then, as long as the right-hand side is positive, one has*

$$\dim(U_r) = t - rs.$$

Equivalently, each additional finite linear subspace trail round decreases the intersection dimension by exactly s.

Proof. We start from the initial case. The MDS property of the block matrix gives $\dim\ker(B) = t - 2s$. Let $W_1 = \ker(B)$. The condition $X' \in W_1$ is the extra condition for a two-round trail. Hence $\dim(U_2) = t - 2s$. For this space to be nontrivial, one needs $t > 2s$. Since the smallest Poseidon instance has $t = 3$ and $s = 1$, standard Poseidon parameters always admit at least a two-round subspace trail.

Now consider three rounds. Let $W_2 = D^{-1}\ker(B)$. Since D is invertible, $\dim(W_2) = t-2s$, and a three-round trail exists if and only if $X' \in U_3 = W_1 \cap W_2$. Let $\{\eta_1, \ldots, \eta_{t-2s}\}$ be a basis of W_1. Then $\{D^{-1}\eta_1, \ldots, D^{-1}\eta_{t-2s}\}$ is a basis of W_2. Put these basis vectors as columns in

$$M_1 = \big(\eta_1, \ldots, \eta_{t-2s}, D^{-1}\eta_1, \ldots, D^{-1}\eta_{t-2s}\big) \in \mathbb{F}_p^{(t-s)\times 2(t-2s)}.$$

Then $\operatorname{rank}(M_1) = \dim(W_1 + W_2)$. The dimension formula gives

$$\dim(W_1 \cap W_2) = \dim(W_1) + \dim(W_2) - \dim(W_1 + W_2) = 2(t-2s) - \operatorname{rank}(M_1).$$

Thus a nontrivial three-round subspace trail requires $\dim(W_1 \cap W_2) \geq 1$.

If $t \leq 3s$, then $2(t-2s) \leq t - s$. In this case, M_1 has at most as many columns as rows. If the columns are full rank, then $\operatorname{rank}(M_1) = 2(t-2s)$, so $\dim(W_1 \cap W_2) = 0$. The intersection contains only the zero vector and cannot support a nontrivial trail.

If $t \leq 3s$ but M_1 is not full column rank, then its columns satisfy an extra linear dependence. Let $N = 2(t-2s)$ be the number of columns and $m = t - s$ be the ambient dimension. In the random model, these N columns can be treated as random vectors in general position. If the first $N-1$ columns are linearly independent, then the probability that the last column lies in their span is at most

$$\frac{p^{N-1}}{p^m} = p^{N-1-m} = p^{t-3s-1}.$$

When $t \leq 3s$, this is at most p^{-1}. Hence the rank-defect event has probability $O(1/p)$. It is negligible over the large characteristic fields of interest. Thus, when $t \leq 3s$, a three-round subspace trail typically does not exist.

If $t > 3s$, then M_1 has more columns than rows. Hence $\operatorname{rank}(M_1) \leq t - s < 2t - 4s$. Substituting into the dimension formula gives

$$\dim(W_1 \cap W_2) \geq 2(t-2s) - (t-s) = t - 3s > 0.$$

Thus, when $t > 3s$, the nontrivial intersection needed for a three-round trail must exist.

Under the general position assumption, one has $\operatorname{rank}(M_1) = t-s$, and hence $\dim(W_1 \cap W_2) = t-3s$. If this equality fails, then the columns of M_1 satisfy an additional nontrivial linear relation. For a random matrix over a finite field, such a rank defect occurs with probability $O(1/p)$. This is the same type of rare algebraic degeneration as the extra relations discussed for invariant subspaces. We therefore exclude this exceptional rank drop from the heuristic random model. For almost all randomly generated Cauchy matrices, we may take $\dim(W_1 \cap W_2) = t-3s$. Compared with the two-round space W_1, the three-round space loses exactly another s dimensions.

By Proposition 2, extending the trail to four rounds requires $X' \in U_4 = (W_1 \cap W_2) \cap W_3$, where $W_3 = D^{-2}\ker(B)$. The dimension formula gives

$$\dim\big((W_1 \cap W_2) \cap W_3\big) = \dim(W_1 \cap W_2) + \dim(W_3) - \dim\big((W_1 \cap W_2) + W_3\big).$$

Here $\dim(W_1 \cap W_2) = t - 3s$ and $\dim(W_3) = t - 2s$. Therefore, for the left-hand side to be positive, one needs at least

$$\dim\big((W_1 \cap W_2) + W_3\big) < 2t - 5s.$$

When $t \leq 4s$, the right-hand side is at most $t - s$. If the corresponding column-space matrix is full rank, the intersection dimension drops to zero. If it is rank deficient, this is again a degeneration with probability $O(1/p)$. Thus a nontrivial four-round trail typically appears only when $t > 4s$. In general position, $\dim\big((W_1 \cap W_2) + W_3\big) = t - s$, and hence

$$\dim\big((W_1 \cap W_2) \cap W_3\big) = (t - 3s) + (t - 2s) - (t - s) = t - 4s.$$

The general induction step is identical. Suppose $r \geq 2$, let $U_r = \bigcap_{i=1}^{r-1} W_i$, and assume that $\dim(U_r) = t - rs > 0$. To extend by one more round, we consider $U_{r+1} = U_r \cap W_r$. The dimension formula gives

$$\dim(U_r \cap W_r) = \dim(U_r) + \dim(W_r) - \dim(U_r + W_r).$$

In general position, the sum reaches the largest dimension allowed by the ambient space $\mathbb{F}_p^{t-s}$, namely $t - s$. Therefore

$$\dim(U_{r+1}) = (t - rs) + (t - 2s) - (t - s) = t - (r + 1)s.$$

Once this value is no longer positive, the nontrivial intersection disappears in general position. The partial S-box in the next round must then be activated. This proves $\dim(U_r) = t - rs$ until the right-hand side becomes nonpositive. □

The proof of Lemma 3 also shows what happens when the dimension inequality is not tight at an intermediate intersection. In that case, the relevant subspaces satisfy additional nontrivial linear relations, so the rank drop is larger than in the general position case. In the large characteristic random Cauchy model, such events are very rare. There is one useful extreme case. If $W_2 = W_1$, then $D^{-1}\ker(B) = \ker(B)$, so $\ker(B)$ itself is a D^{-1}-invariant subspace. This gives an infinitely long trail and represents the extreme case where the whole space $\ker(B)$ is invariant under D^{-1}. For a fixed η_i, the probability of $D^{-1}\eta_i \in W_1$ is approximately $(p^{t-2s} - 1)/(p^{t-s} - 1) \approx p^{-s}$. If these events are further treated as independent, then

$$\Pr[W_2 \subseteq W_1] \approx p^{-s(t-2s)}.$$

By the preceding subsection, for the codimension-one invariant subspace W_2 to have dimension $t - 2s$, one needs $s = 1$. In that case the estimate becomes $p^{-(t-2)}$, which is consistent with the infinite trail probability estimate above.

Theorem 1 (Maximum length of finite linear subspace trails)**.** *Assume that no infinitely long linear subspace trail exists, and that the above subspace chain is in general position during the intersection process. If each internal partial round activates s S-boxes, then the maximum length of a finite linear subspace trail is*

$$r_{\max} = \left\lceil \frac{t}{s} \right\rceil - 1.$$

Proof. By Proposition 2, the first r internal partial rounds remain on a linear subspace trail if and only if $U_r = \bigcap_{i=1}^{r-1} W_i$ is nontrivial. By Lemma 3, in general position $\dim(U_r) = t-rs$. Hence the maximum possible r satisfies $t-rs > 0$ and $t-(r+1)s \leq 0$, or equivalently $t/s - 1 \leq r < t/s$. Since r is an integer, $r_{\max} = \lceil t/s \rceil - 1$. In other words, in round $\lceil t/s \rceil$, the partial S-boxes in the partial round map r_p inevitably reintroduce higher degree terms into the internal state. □

For the common Poseidon instances, partial rounds usually have $s = 1$. Then $r_{\max} = t - 1$. Thus, if one ignores attack degrees of freedom and asks only whether the relevant linear space is nontrivial, a finite linear subspace trail typically lasts for at most $t - 1$ rounds. At round t, the partial S-box is activated.

For the Poseidon setting considered here, this gives a linear algebraic answer to the question on finite trail length raised in [15]. Once infinitely long subspace trails are excluded, the finite trail length is controlled by the dimension chain of intersections. For a random Cauchy linear layer in general position, the $\lceil t/s \rceil$-th internal partial round is already enough to activate at least one partial S-box.

If one considers the practical degrees of freedom in an attack and uses Ec extra constraints, then the realizable linear subspace trail length discussed in Section 3 is $\lfloor Ec/s \rfloor$. In the preimage problem, the extra constraint budget satisfies $Ec \leq t - 2$, since both c and d are at least 1. In the usual case $s = 1$, this gives $\lfloor Ec/s \rfloor \leq t-2 < t-1 = \lceil t/s \rceil - 1$. Thus, in the usual Poseidon setting, the number of realizable linearized partial rounds is limited by the available extra constraints rather than by the maximal finite trail length allowed by the MDS matrix.

For Poseidon2 and Neptune, the internal linear layers are weaker than the MDS matrix used in Poseidon from the viewpoint of the rank argument above. Therefore, the MDS-based proof does not directly apply to these layers. However, the next section uses the special structure of their non-MDS internal linear layers to give a direct analysis, and shows that the conclusions of this section extend to Poseidon2 and Neptune.

4 Linear Subspace Trails in Poseidon2 and Neptune

In the previous sections, we studied linear subspace trails for the internal linear layer of Poseidon. In this section we specialize the general algebraic criteria to the internal linear layers used in Poseidon2 and Neptune. As recalled in Section 2.1, both designs use efficient internal linear layers of the form $M_I = \Lambda + \mathbf{1}\mathbf{1}^\mathsf{T}$, namely a diagonal matrix plus the all-ones matrix. This special form allows the general conditions of Proposition 1 to be written in a more explicit way.

Let

$$M_I = \begin{pmatrix} \mu_0 & 1 & \cdots & 1 \\ 1 & \mu_1 & \cdots & 1 \\ \vdots & \vdots & \ddots & \vdots \\ 1 & 1 & \cdots & \mu_{t-1} \end{pmatrix} \in \mathbb{F}_p^{t\times t}.$$

Equivalently,

$$M_I = \Lambda + \mathbf{1}\mathbf{1}^\mathsf{T}, \qquad \Lambda = \mathrm{diag}(\lambda_0, \ldots, \lambda_{t-1}), \qquad \lambda_i = \mu_i - 1,$$

where $\mathbf{1}$ denotes the all-one column vector. For Poseidon2, this form is used for the generated internal matrices with $t = 2, 3, 4t'$ and $t' \in \mathbb{N}$. For Neptune, the same type of internal matrix is used for even state widths.

We first explore the determinant and characteristic polynomial of this matrix family. The following formula is a direct consequence of the matrix determinant lemma,

$$\det(A + uv^{\mathsf{T}}) = \det(A) + v^{\mathsf{T}} \operatorname{adj}(A)u,$$

where $\operatorname{adj}(A)$ denotes the adjugate matrix.

Lemma 4. *Let*

$$M = \operatorname{diag}(\lambda_0, \ldots, \lambda_{n-1}) + \mathbf{1}\mathbf{1}^{\mathsf{T}} \in \mathbb{F}_p^{n \times n}.$$

Then

$$\det(M) = \prod_{i=0}^{n-1} \lambda_i + \sum_{i=0}^{n-1} \prod_{j \neq i} \lambda_j.$$

Moreover, if $g(x) = \prod_{i=0}^{n-1}(x - \lambda_i)$, *then the characteristic polynomial of* M *is*

$$\chi_M(x) = \det(xI_n - M) = g(x) - g'(x).$$

Proof. The determinant formula follows from the matrix determinant lemma by taking $A = \operatorname{diag}(\lambda_0, \ldots, \lambda_{n-1})$ and $u = v = \mathbf{1}$.

Since $\operatorname{adj}(A) = \operatorname{diag}\left(\prod_{j \neq 0} \lambda_j, \ldots, \prod_{j \neq n-1} \lambda_j\right)$, we obtain

$$\det(A + \mathbf{1}\mathbf{1}^{\mathsf{T}}) = \prod_{i=0}^{n-1} \lambda_i + \sum_{i=0}^{n-1} \prod_{j \neq i} \lambda_j.$$

For the characteristic polynomial, apply the same formula to

$$xI_n - M = \operatorname{diag}(x - \lambda_0, \ldots, x - \lambda_{n-1}) - \mathbf{1}\mathbf{1}^{\mathsf{T}}.$$

This gives

$$\chi_M(x) = \prod_{i=0}^{n-1}(x - \lambda_i) - \sum_{i=0}^{n-1} \prod_{j \neq i}(x - \lambda_j).$$

The second term is exactly $g'(x)$, and hence

$$\chi_M(x) = g(x) - g'(x).$$

□

The expression $\chi_M(x) = g(x) - g'(x)$ gives a simple necessary condition for irreducibility. If two diagonal parameters coincide, say $\lambda_a = \lambda_b$ with $a \neq b$, then $g(x)$ has a repeated root. Equivalently, $g(x)$ and $g'(x)$ have a nontrivial common factor and the common factor also divides $\chi_M(x)$. Hence $\chi_M(x)$ is reducible over the base field. Therefore, any matrix in this family whose characteristic polynomial is irreducible over $\mathbb{F}_p$ must have pairwise distinct diagonal parameters λ_i, or equivalently pairwise distinct μ_i. The designers' check algorithm in [17] excludes reducible characteristic polynomials. Thus any internal linear layer satisfying the designers' conditions must have pairwise distinct diagonal entries.

We now specialize the linear trail criterion to Poseidon2 and Neptune. Splitting the internal matrix with respect to the active coordinate gives $M_I = \left(\begin{smallmatrix} A & B \\ C & D \end{smallmatrix}\right)$, where

$$B = \mathbf{1}^{\mathsf{T}} \in \mathbb{F}_p^{1 \times t-1},$$

and

$$D = \mathrm{diag}(\lambda_1, \ldots, \lambda_{t-1}) + \mathbf{1}\mathbf{1}^{\mathsf{T}} \in \mathbb{F}_p^{t-1\times t-1}.$$

Here $\mathbf{1}$ denotes the all-one vector of length $t-1$. By Proposition 1, the existence of a linear subspace trail of the prescribed length is equivalent to the existence of a nonzero vector X' satisfying

$$X' \in \ker\Big(\big(B^{\mathsf{T}}, (BD)^{\mathsf{T}}, (BD^2)^{\mathsf{T}}, \ldots, (BD^{\ell-1})^{\mathsf{T}}\big)^{\mathsf{T}}\Big).$$

Therefore, the maximum possible length of such a linear subspace trail is controlled by the rank growth of the matrix

$$\big(B^{\mathsf{T}}, (BD)^{\mathsf{T}}, (BD^2)^{\mathsf{T}}, \ldots, (BD^{\ell-1})^{\mathsf{T}}\big)^{\mathsf{T}}.$$

The special form of B and D reduces this matrix to a Vandermonde matrix. Let $\Lambda' = \mathrm{diag}(\lambda_1, \ldots, \lambda_{t-1})$. Since $D = \Lambda' + \mathbf{1}\mathbf{1}^{\mathsf{T}}$, we have, for any row vector v, $vD = v\Lambda' + (v\mathbf{1})\mathbf{1}^{\mathsf{T}}$.

The second term is a multiple of $B = \mathbf{1}^{\mathsf{T}}$. Therefore the row space generated by $B, BD, \ldots, BD^{\ell-1}$ is the same as the row space generated by

$$\mathbf{1}^{\mathsf{T}}, \mathbf{1}^{\mathsf{T}}\Lambda', \ldots, \mathbf{1}^{\mathsf{T}}(\Lambda')^{\ell-1}.$$

Equivalently, it is the row space of the Vandermonde matrix

$$V_\ell = \begin{pmatrix} 1 & 1 & \cdots & 1 \\ \lambda_1 & \lambda_2 & \cdots & \lambda_{t-1} \\ \lambda_1^2 & \lambda_2^2 & \cdots & \lambda_{t-1}^2 \\ \vdots & \vdots & \ddots & \vdots \\ \lambda_1^{\ell-1} & \lambda_2^{\ell-1} & \cdots & \lambda_{t-1}^{\ell-1} \end{pmatrix}.$$

Hence

$$\ker \begin{pmatrix} B \\ BD \\ \vdots \\ BD^{\ell-1} \end{pmatrix} = \ker(V_\ell).$$

Proposition 3. *Assume that $\lambda_1, \ldots, \lambda_{t-1}$ are pairwise distinct. Then, for every $1 \le \ell \le t-1$, we have* $\mathrm{rank}(V_\ell) = \ell$. *Consequently,*

$$\dim \ker(V_\ell) = t - 1 - \ell.$$

In particular, whenever the linear subspace trail is extended by one more internal round, the dimension of the corresponding subspace decreases by one. Thus the maximum length of a nontrivial linear subspace trail is $t-1$, in agreement with Proposition 1.

Proof. The matrix V_ℓ is an $\ell \times (t-1)$ Vandermonde matrix evaluated at the pairwise distinct points $\lambda_1, \ldots, \lambda_{t-1}$. For $\ell \le t-1$, its rows are linearly independent. Indeed, a linear relation among the rows would give a polynomial of degree at most $\ell - 1$ vanishing at the $t-1$ distinct points $\lambda_1, \ldots, \lambda_{t-1}$. If $\ell \le t-1$, this is possible only for the zero polynomial. Hence $\mathrm{rank}(V_\ell) = \ell$, and the kernel dimension is $t - 1 - \ell$.

Taking $\ell = t-1$, the matrix V_{t-1} is square and has determinant

$$\det(V_{t-1}) = \prod_{1\le i<j\le t-1} (\lambda_j - \lambda_i) \neq 0.$$

Thus its kernel is zero. Therefore the last nontrivial case occurs at $\ell = t-2$, which corresponds to a trail of length $t-1$ in agreement with Proposition 1. □

We also explore what happens if the diagonal parameters are not pairwise distinct. Let $m = \#\{\lambda_1, \ldots, \lambda_{t-1}\}$ be the number of distinct values among the diagonal parameters. Then $\operatorname{rank}(V_\ell) = \min(\ell, m)$, and therefore

$$\dim \ker(V_\ell) = t - 1 - \min(\ell, m).$$

In particular, if some of the λ_i coincide, then the constraint chain does not continue to decrease until zero. Instead, it stabilizes at dimension $(t-1) - m$.

Repeated diagonal parameters give explicit persistent directions. Let e_a and e_b denote the standard basis column vectors, whose a-th and b-th coordinates are 1, respectively, and whose other coordinates are 0. If $\lambda_a = \lambda_b$, then

$$e_a - e_b \in \ker(\mathbf{1}^\mathsf{T}), \qquad D(e_a - e_b) = \lambda_a(e_a - e_b).$$

Thus $e_a - e_b$ belongs to the kernel of the stacked matrix

$$\left(B^\mathsf{T}, (BD)^\mathsf{T}, (BD^2)^\mathsf{T}, \ldots\right)^\mathsf{T},$$

and therefore gives an infinitely long linear subspace trail. This is consistent with the invariant subspace criterion discussed above.

Therefore, repeated diagonal parameters μ_i among the inactive coordinates, equivalently repeated λ_i in the diagonal block D, necessarily lead to an infinite linear subspace trail. The irreducibility checks in [17] exclude this situation for the generated instances. When diagonal parameters are pairwise distinct, the maximum length of a finite linear subspace trail is exactly $t-1$ internal rounds. Thus, as in the case of Poseidon, the effective number of linearized partial rounds for Poseidon2 and Neptune is bounded by Ec, even though the structural finite trail length is $t-1$.

5 Conclusion

This paper revisits linear subspace trails in Poseidon-like hash functions, with the analysis based on the MDS matrix and centered on Poseidon. For infinitely long trails, we relate their existence to invariant subspace conditions of the linear layer and give a quantitative heuristic explanation for why such conditions are rare. For finite trails, we give a linear algebraic characterization of the constraint chain and show that, when each internal partial round activates s S-box coordinates, the typical maximal trail length is bounded by $\lceil t/s \rceil - 1$. We also relate this structural bound to preimage attacks by comparing it with the CICO constraint budget Ec, which determines how many linearized rounds can actually be realized in the attack model. For Poseidon2 and Neptune, the same proof based on the MDS matrix does not directly apply, but a direct analysis of their special internal linear layers gives more structured conclusions. Together, these results clarify the range in which linear subspace trails can affect algebraic attacks on Poseidon/Poseidon2 and Neptune.

References

[1] Tomer Ashur, Thomas Buschman, and Mohammad Mahzoun. Algebraic cryptanalysis of the HADES design strategy: Application to Poseidon and Poseidon2. In *ACISP*, LNCS 14896, pages 225–244. Springer, 2024. doi:10.1007/978-981-97-5028-3_12.

[2] Antoine Bak, Augustin Bariant, Aurélien Bœuf, Pierre Briaud, Morten Øygarden, and Atharva Phanse. The algebraic CheapLunch: Extending FreeLunch attacks on arithmetization-oriented primitives beyond CICO-1. Cryptology ePrint Archive, Paper 2025/2040, 2025. https://eprint.iacr.org/2025/2040.

[3] Antoine Bak, Augustin Bariant, Maël Hostettler, and Vincent Neiger. Resultants meet resultant: Improving CICO-1 and CICO-2 attacks on ZK-friendly permutations. *IACR Cryptology ePrint Archive*, Paper 2026/1281, 2026. https://eprint.iacr.org/2026/1281.

[4] Abdelhamid Bakhta, Eli Ben Sasson, Avihu Levy, and David Levit Gurevich. EIP-5988: Add Poseidon hash function precompile. Ethereum Improvement Proposals, no. 5988, 2022. https://eips.ethereum.org/EIPS/eip-5988.

[5] Magali Bardet, Jean-Charles Faugère, and Bruno Salvy. On the complexity of Gröbner basis computation of semi-regular overdetermined algebraic equations. In *International Conference on Polynomial System Solving*, pages 71–74, 2004.

[6] Augustin Bariant, Clémence Bouvier, Gaëtan Leurent, and Léo Perrin. Algebraic attacks against some arithmetization-oriented primitives. *IACR Transactions on Symmetric Cryptology*, 2022(3):73–101, 2022. doi:10.46586/tosc.v2022.i3.73-101.

[7] Augustin Bariant, Aurélien Boeuf, Axel Lemoine, Irati Manterola Ayala, Morten Øygarden, Léo Perrin, and Håvard Raddum. The algebraic FreeLunch: Efficient Gröbner basis attacks against arithmetization-oriented primitives. In *CRYPTO*, LNCS 14923, pages 139–173. Springer, 2024. doi:10.1007/978-3-031-68385-5_5.

[8] Augustin Bariant, Aurélien Boeuf, Pierre Briaud, Maël Hostettler, Morten Øygarden, and Håvard Raddum. Improved resultant attack against arithmetization-oriented primitives. In *CRYPTO 2025*, LNCS 16004, pages 335–367. Springer, 2025. doi:10.1007/978-3-032-01901-1_11.

[9] Tim Beyne, Anne Canteaut, Itai Dinur, Maria Eichlseder, Gregor Leander, Gaëtan Leurent, María Naya-Plasencia, Léo Perrin, Yu Sasaki, Yosuke Todo, and Friedrich Wiemer. Out of oddity: New cryptanalytic techniques against symmetric primitives optimized for integrity proof systems. In *CRYPTO*, LNCS 12172, pages 299–328. Springer, 2020. doi:10.1007/978-3-030-56877-1_11.

[10] Guido Bertoni, Joan Daemen, Michaël Peeters, and Gilles Van Assche. On the indifferentiability of the sponge construction. In *EUROCRYPT*, LNCS 4965, pages 181–197. Springer, 2008. doi:10.1007/978-3-540-78967-3_11.

[11] Guido Bertoni, Joan Daemen, Michaël Peeters, and Gilles Van Assche. Cryptographic sponge functions. Version 0.1, 2011. https://keccak.team/files/CSF-0.1.pdf.

[12] Jean-Charles Faugère and Ludovic Perret. Algebraic attacks against STARK-friendly ciphers. Cryptology ePrint Archive, Paper 2020/948, 2020. https://eprint.iacr.org/2020/948.

[13] Lorenzo Grassi, Reinhard Lüftenegger, Christian Rechberger, Dragos Rotaru, and Markus Schofnegger. On a generalization of substitution-permutation networks: The HADES design strategy. In *EUROCRYPT*, LNCS 12106, pages 674–704. Springer, 2020. doi:10.1007/978-3-030-45724-2_23.

[14] Lorenzo Grassi, Dmitry Khovratovich, Christian Rechberger, Arnab Roy, and Markus Schofnegger. Poseidon: A new hash function for zero-knowledge proof systems. In *30th USENIX Security Symposium*, pages 519–535. USENIX Association, 2021. https://www.usenix.org/conference/usenixsecurity21/presentation/grassi.

[15] Lorenzo Grassi, Christian Rechberger, and Markus Schofnegger. Proving resistance against infinitely long subspace trails: How to choose the linear layer. *IACR Transactions on Symmetric Cryptology*, 2021(2):314–352, 2021. doi:10.46586/tosc.v2021.i2.314-352.

[16] Lorenzo Grassi, Silvia Onofri, Marco Pedicini, and Luca Sozzi. Invertible quadratic non-linear layers for MPC/FHE/ZK-friendly schemes over $\mathbb{F}_{p^n}$: Application to Poseidon. *IACR Transactions on Symmetric Cryptology*, 2022(3):20–72, 2022. doi:10.46586/tosc.v2022.i3.20-72.

[17] Lorenzo Grassi, Dmitry Khovratovich, and Markus Schofnegger. Poseidon2: A faster version of the Poseidon hash function. In *AFRICACRYPT*, LNCS 14064, pages 177–203. Springer, 2023. doi:10.1007/978-3-031-37679-5_8.

[18] Lorenzo Grassi, Katharina Koschatko, and Christian Rechberger. Poseidon and Neptune: Gröbner basis cryptanalysis exploiting subspace trails. *IACR Transactions on Symmetric Cryptology*, 2025(2):34–86, 2025. doi:10.46586/tosc.v2025.i2.34-86.

[19] Kishan Chand Gupta, Sumit Kumar Pandey, Indranil Ghosh Ray, and Susanta Samanta. Cryptographically significant MDS matrices over finite fields: A brief survey and some generalized results. *Advances in Mathematics of Communications*, 13(4):779–843, 2019. doi:10.3934/amc.2019045.

[20] Thomas Jakobsen and Lars R. Knudsen. The interpolation attack on block ciphers. In *FSE*, LNCS 1267, pages 28–40. Springer, 1997.

[21] Nathan Keller and Asaf Rosemarin. Mind the middle layer: The HADES design strategy revisited. In *EUROCRYPT*, LNCS 12697, pages 35–63. Springer, 2021. doi:10.1007/978-3-030-77886-6_2.

[22] Katharina Koschatko, Lorenzo Grassi, Dmitry Khovratovich, Christian Rechberger, Markus Schofnegger, Verena Schröppel, and Zhuo Wu. Poseidon(2)b: Binary field versions of Poseidon/Poseidon2. *IACR Communications in Cryptology*, 2(4), 2026. doi:10.62056/a66ce0zn4.

[23] Katharina Koschatko, Reinhard Lüftenegger, and Christian Rechberger. Exploring the six worlds of Gröbner basis cryptanalysis: Application to Anemoi. *IACR Transactions on Symmetric Cryptology*, 2024(4):138–190, 2024. doi:10.46586/tosc.v2024.i4.138-190.

[24] Xuejia Lai and James L. Massey. A proposal for a new block encryption standard. In *EUROCRYPT*, LNCS 473, pages 389–404. Springer, 1991. doi:10.1007/3-540-46877-3_35.

[25] Simon-Philipp Merz and Àlex Rodríguez García. Skipping class: Algebraic attacks exploiting weak matrices and operation modes of Poseidon2(b). *IACR Cryptology ePrint Archive*, Paper 2026/306, 2026. https://eprint.iacr.org/2026/306.

[26] Hong-Sen Yang, Qun-Xiong Zheng, Jing Yang, Quan-feng Liu, and Deng Tang. A new security evaluation method based on resultant for arithmetic-oriented algorithms. In *ASIACRYPT*, LNCS 15490, pages 457–489. Springer, 2024. doi:10.1007/978-981-96-0941-3_15.

[27] Ziyu Zhao and Jintai Ding. Breaking Poseidon challenges with Graeffe transforms and complexity analysis by FFT lower bounds. Cryptology ePrint Archive, Paper 2025/950, 2025. `https://eprint.iacr.org/2025/950`.