

Poseidon: A New Hash Function for Zero-Knowledge Proof Systems

(Updated Version)

Lorenzo Grassi^{1,2}, Dmitry Khovratovich³, Christian Rechberger¹, Arnab Roy⁴ and Markus Schofnegger¹

¹ IAIK, Graz University of Technology

² Radboud University Nijmegen

³ Ethereum Foundation and Dusk Network

⁴ University of Klagenfurt

`l.grassi@cs.ru.nl,firstname.lastname@iaik.tugraz.at, khovratovich@gmail.com,
arnab.roy@aau.at`

Abstract. The area of practical computational integrity proof systems, like SNARKs, STARKs, Bulletproofs, is seeing a very dynamic development with several constructions having appeared recently with improved properties and relaxed setup requirements. Many use cases of such systems involve, often as their most expensive part, proving the knowledge of a preimage under a certain cryptographic hash function, which is expressed as a circuit over a large prime field. A zero-knowledge proof of coin ownership in the Zcash cryptocurrency is a notable example, where the inadequacy of the SHA-256 hash function for such a circuit caused a huge computational penalty. In this paper, we present a modular framework and concrete instances of cryptographic hash functions which work natively with $\text{GF}(p)$ objects. Our hash function POSEIDON uses up to 8x fewer constraints per message bit than Pedersen Hash.

Keywords: Zero-Knowledge Proof Systems · SNARK · STARK · Hash Function

Previous Version. *This paper is an updated version of [GKK⁺19, Version: 20200205:104144].*
The main differences are as follows.

- Here we propose only POSEIDON over prime fields, since it turned out that all applications we considered are more efficient when working with a prime field rather than a binary field.
- We briefly recall STARKAD, the previous version over binary fields, in Supplementary Material J.
- The security analysis has been updated in order to take care of recent attacks (e.g., [KR20, BCD⁺20]). In particular, an additional assumption on the linear layer has been added.
- The applications part and the comparison w.r.t. other schemes have been updated.

Contents

1	Introduction	2
2	The Poseidon Hash Function	4
2.1	Sponge Construction for POSEIDON $^\pi$	5
2.2	The HADES Design Strategy for Hashing	5
2.3	The Permutation Family POSEIDON $^\pi$	6
3	Applications	8
4	Concrete Instantiations of Poseidon$^\pi$	9
4.1	Main Instances	9
4.2	Domain Separation for POSEIDON	9
5	Cryptanalysis Summary of Poseidon	10
5.1	Security Claims	11
5.2	Summary of Attacks	11
5.3	Security Margin	12
5.4	Attack details	12
6	Poseidon in Zero-Knowledge Proof Systems	15
6.1	State of the Art	15
6.2	SNARKs with POSEIDON $^\pi$	16
6.3	STARKs with POSEIDON $^\pi$	20
A	Auxiliary Files	25
B	Efficient Implementation	25
C	Security Analysis: x^α-Poseidon$^\pi$	26
C.1	Statistical Attacks	27
C.2	Algebraic Attacks	30
D	Security Analysis: x^{-1}-Poseidon$^\pi$	36
D.1	Statistical Attacks	36
D.2	Algebraic Attacks	36
D.3	Gröbner Basis Attack	38
E	Selecting Number of Rounds in General Case	40
F	Initialization of the Grain LFSR Used for Parameter Generation	41
G	Concrete Instances with Security Margin	41
H	Compact Constraints for STARKs and SNARKs	41
I	Merkle Tree Instances of Poseidon	44
J	About Starkad	45

1 Introduction

The recent advances in computational integrity proof systems made a number of computational tasks verifiable in short time and/or in zero knowledge. Several protocols appeared that require one party to prove the knowledge of a seed-derived secret, of an element being part of a large set, or their combination. Whereas accumulator-based solutions [CL02, CKS09] and algebraic Schnorr proofs exist in the area, they are quite involving and thus error-prone, require a trusted setup, are limited in statement language, and are often slow. An alternative is to express secret derivation using cryptographic hash functions, and to prove set membership by presenting an opening in a properly chosen Merkle tree, also built on a cryptographic hash function. Such hash-based protocols require a computational integrity proof system, which can be applied to an arbitrary arithmetic circuit. However, for the protocol to be efficient, proofs must be generated and verified in reasonable time, which in turn requires the hash function to be cheap in a certain metric depending on the proof system.

In the middle of 2020, the most popular proof systems are ZK-SNARKs (Pinocchio [PHGR13], Groth16 [Gro16], PLONK [GWC19], Marlin [CHM⁺19] to name a few), Bulletproofs [BBB⁺18], ZK-STARKs [BBHR18], MPC-in-the-head systems [GMO16, AHIV17]. The former two groups have already been applied to a number of real-world protocols, whereas the latter ones are the most promising from the perspective of post-quantum security. These systems use two quite different circuit descriptions so that the proof size and generation time are computed differently:

- The R1CS format (rank-1 quadratic constraints) describes the circuit as a set of special quadratic polynomials of the form $L_1(X) \cdot L_2(X) = L_3(X)$, where X is the tuple of internal and input variables, L_i are affine forms and \cdot is the field multiplication, and (possibly in an affine-equivalent form) is used in almost all SNARKs and Bulletproofs. The circuit multiplication and addition gates are defined over a prime field $\text{GF}(p)$. The proof generation complexity is directly proportional to the number T of constraints, which often corresponds to the number of multiplication gates. The prime field $\text{GF}(p)$ is the scalar field of an elliptic curve, where for ZK-SNARKs the curve should be pairing-friendly and for Bulletproofs it should just be a secure curve.
- The AET metric is used in ZK-STARKs and (to some extent) in the PLONK proof system. The computation is expressed as a set of internal program states related to each other by polynomial equations of degree d . The state consists of w field elements and undergoes T transformations. The proof generation is roughly proportional to the product $w \cdot d \cdot T$. The number and sparsity of polynomial constraints do not play a major role.

Our goal was to design a family of hash functions that are optimal in the R1CS (as the most widespread) and good in the AET metric, while also supporting different finite field sizes. It turned out that the substitution-permutation network (SPN) design, well-known in symmetric cryptography, allows for a generic hash function framework where the only security-critical parameter that has to be changed for each instance is the number of rounds, and we provide an efficient and transparent strategy for its choice. The S-Box is chosen as power map $x \mapsto x^d$, where $d \geq 3$ is usually chosen as the smallest integer that guarantees invertibility and provides non-linearity. In particular, the cube function x^5 is almost universally chosen, apart from cases of fields where this function is not a bijection. Instead, we suggest other S-Boxes such as x^3 or $1/x$ for these cases. Thanks to a succinct representation of the functions and a low S-Box degree, we are able to optimize the circuit significantly for PLONK and RedShift proof systems, with performance improvements by

Table 1: Our primary proposals and their competitors. ‘Tree’ means Merkle tree arity and is equal to the rate/capacity ratio. Curve denotes the curve (BLS12-381, BN254, Ed25519) whose (subgroup) scalar field determines the prime size. The R1CS/bit costs are obtained by dividing the R1CS prover costs by message rate. Timings are from a third-party implementation of *Rescue* and POSEIDON on i9-8950 CPU @2.9 Ghz, 32 GB RAM.

Name	S-Box	Rate bits/perm.	SB size ($\log_2 p$)	Tree arity	R_F	R_P	Curve Scalar field	R1CS /perm.	R1CS /bit	Time /perm.
POSEIDON-80	x^5	510	255	2:1	8	33	BLS/BN/Ed	171	0.34	
	x^5	1020	255	4:1	8	35		225	0.22	
POSEIDON-128	x^5	510	255	2:1	8	57	BLS/BN/Ed	243	0.47	0.022 ms
	x^5	1020	255	4:1	8	60		300	0.29	
	x^5	2040	255	8:1	8	63		405	0.2	
POSEIDON-256	x^5	1020	255	2:1	8	120	BLS/BN/Ed	600	0.6	
	x^5	2040	255	4:1	8	120		600	0.3	
Pedersen Hash	-	516	-	2:1	-	-	BLS12-381	869	1.68	
<i>Rescue</i>	x^5 & $x^{1/5}$	510	255	2:1	16	-	BLS/BN/Ed	268	0.52	0.337 ms
		1020	255	4:1	10	-		300	0.29	
		2040	255	8:1	10	-		450	0.22	

a factor of up to 40.

Our Contributions. We design and analyze a family of hash functions over $\text{GF}(p)$ named POSEIDON. The internal permutation is called POSEIDON^π and is based on the HADES design strategy [GLR⁺20], which is essentially a strategy based on substitution-permutation networks with t cells, but including the use of so-called *partial* rounds, which use non-linear functions only for part of the state. In our specific construction, only one S-Box is used in these partial rounds, while full non-linear layers (i.e., t S-Boxes) are used in all other rounds. This is done to reduce the R1CS or AET cost.

We aim to support security levels of 80, 128, and 256 bits, where the security is the same for collision and preimage resistance. For each pair (basic field, security level) we suggest a concrete instance of POSEIDON. In our hash function, a few S-Box elements are reserved for the capacity (roughly double the security level in bits), and the rest for the rate. The permutation width is determined by the application: It is set close to 1280 bits for long-message hashing, whereas for Merkle trees we support various widths to enable 2:1, 4:1, and other arities and thus higher ZK performance.

We provide an extensive cryptanalysis of POSEIDON with an accent on algebraic methods as these prove to be the most effective. We explore different variants of interpolation, Gröbner basis, and higher-order differential attacks. As our permutations are quite wide, we do not aim for them behaving like randomly chosen permutations. Instead, for a security level of M bits we require that no attack could exhibit a non-random property of a permutation faster than in 2^M queries. We then calculate the maximum number of rounds for each field, security level, and fixed permutation width that can be attacked. Then we select the number of rounds for concrete instances together with a security margin.

We have evaluated the number of constraints in POSEIDON instances for the R1CS metric and the AET metric. Our primary proposals POSEIDON-80/128/256 are listed in Table 1 (BLS being BLS12-381¹, BN being BN254 [W⁺14], Ed being the Ristretto group²) and are compared to similar-purpose designs. As supplementary material we provide reference implementations of various instances of designs, scripts to choose and create those instances.

¹<https://electriccoin.co/blog/new-snark-curve/>

²<https://ristretto.group>

We also have third-party benchmarks of POSEIDON for regular hashing³ (Table 1) and in ZK proof systems: PLONK (Table 6), Groth16 (Table 3), and Bulletproofs (Table 5).

Comparison to [GLR⁺20]. Since the design of POSEIDON follows the same strategy as block ciphers in [GLR⁺20], we provide an explicit list of new material crafted for this paper:

- Hash-function specific (CICO, keyless, preimage) algebraic attacks, their analysis, and fixes against recent hash-only attacks.
- Orientation towards various zero knowledge proof systems. Tips how to improve prover performance in these systems.
- Concrete instances for Merkle trees and variable-length hashing.
- Concrete benchmarks for zero-knowledge proofs of accumulated values in Merkle trees, demonstration that it can be done in 1 second for billion-size tree.

Related Work. The Zcash designers introduced a new 256-bit hash function called Pedersen hash [HBHW19, p.134], which is effectively a vectorized Pedersen commitment in elliptic curve groups with short vector elements. For the claimed 128-bit security level, it utilizes 869 constraints per 516-bit message chunks, thus having 1.7 constraints per bit, whereas our POSEIDON instances use from 0.2 to 0.45 constraints per bit, depending on the underlying prime field.

For the binary field case, Ashur and Dhooghe [AD18] have recently introduced the STARK-friendly block cipher JARVIS and its derivative hash function FRIDAY with several instances and security levels. They use a key-alternating structure with a single inverse S-Box, followed by an affine transformation (with low degree in the extension field). However, both JARVIS and FRIDAY were successfully attacked shortly after their publication [ACG⁺19]. In the response, the authors created a new family of SNARK/STARK-friendly hash functions with *Vision* (binary fields) and *Rescue* (prime fields) being main instances [AABS⁺19]. The latter two share some similarity with our design with two important differences: First, all S-Box layers are full (there are no partial rounds). Moreover, every second layer has S-Boxes of the form $x^{1/d}$ for small d . This approach prevents some algebraic attacks but is also more expensive in software as the resulting power functions have high Hamming weight and thus require many squarings.

Structure of the Paper. We introduce POSEIDON as a HADES-based hash in Section 2 and follow up with real-world applications in Section 3. Concrete instances with round numbers and domain constants are given in Section 4. We summarize the cryptanalysis results in Section 5 with the details in Appendix. Finally, we estimate the performance of of POSEIDON instances in zero-knowledge proof systems in Section 6 by computing RICS (SNARK) and AET (STARK) costs.

2 The Poseidon Hash Function

In the following we propose the hash function POSEIDON⁴ over \mathbb{F}_p^t (for a prime $p \approx 2^n$), which is constructed by instantiating a sponge construction [BDPA08] with the POSEIDON ^{π} permutation. POSEIDON ^{π} is a variant of HADESMiMC proposed in [GLR⁺20], albeit instantiated with a fixed key (e.g., 0^κ , where κ is the security level of HADESMiMC).

³https://github.com/shamatar/poseidon_hash and https://github.com/shamatar/rescue_hash

⁴In Greek mythology, Poseidon was the brother of Zeus and Hades, and he was god of the Sea and other waters, of earthquakes, and of horses.

In the following, we sometimes use the notation $p \approx 2^n$ and $N = n \cdot t \approx \log_2(p) \cdot t$ to denote the approximate size of the texts in bits.

2.1 Sponge Construction for Poseidon $^\pi$

We recall that⁵ when the internal permutation \mathcal{P} of an N -bit sponge function (composed of a c -bit capacity and an r -word rate: $N = c + r$) is modeled as a randomly chosen permutation, it has been shown in [BDPA08] to be indifferentiable from a random oracle up to $2^{c/2}$ calls ($|\mathbb{F}|^{c/2}$ if capacity is counted in field elements) to \mathcal{P} . In other words, a sponge with a capacity of c provides $2^{c/2}$ bits of collision and $2^{c/2}$ bits of (second) preimage resistance. Given a permutation of size N and a desired security level s , we can hash $r = N - 2s$ bits per call to the permutation. Following this design strategy, we choose the number of rounds of the inner permutation POSEIDON $^\pi$ in order to ensure that such a permutation does not exhibit non-generic properties up to 2^M queries⁶, where M is the desired security level. For this we set the capacity to $2M$ and denote by POSEIDON- M a hash function that provides M bits of security against collision and preimage attacks.

2.2 The Hades Design Strategy for Hashing

Cryptographic permutations usually consist of an efficient round function which is applied sufficiently many times in order to make the permutation behave like a randomly drawn one. In general, the same round function is used throughout the permutation, in order to destroy all possible symmetries and structural properties of the round function.

In [GLR⁺20], the authors consider different round functions within the same construction. More precisely, they mix rounds with *full S-Box layers* and rounds with *partial S-Box layers*. The motivation to have different types of rounds is that full S-Box layers are expensive in software and ZK proof systems but are a good protection against statistical attacks, whereas partial layers are relatively cheap but are, in some cases, similarly good as full ones against algebraic attacks.

Details on the Hades Strategy. The HADES design strategy consists of R_f rounds in the beginning, in which S-Boxes are applied to the full state. After these rounds, R_P rounds in the middle contain only a single S-Box in each round, and the rest of the state goes through the non-linear layer unchanged (i.e., identity functions are used instead of the missing S-Boxes). Finally, R_f rounds at the end are applied by again using S-Boxes for the full state.

The idea of this approach is to provide arguments for the security against statistical attacks using the $R_F = 2R_f$ rounds with full S-Box layers in the beginning and in the end together with the *wide trail strategy* [DR01], which is also used in, e.g., the AES [DR02]. On the other hand, the R_P rounds with partial S-Box layers are a more efficient way to increase the degree of the overall function, and are mainly used for arguments against algebraic attacks.

A detailed overview of this approach is shown in Fig. 1.

⁵We present the Sponge construction over a binary field in order to follow the presentation made in [BDPA08]. It can easily be generalized for a prime field \mathbb{F}_p^t by replacing each (N/t) -bit word by a $(\lceil \log_2(p) \rceil)$ -bit one.

⁶In other words, such a permutation cannot be distinguished from a randomly drawn permutation.




Figure 1: Construction of the HADES-based POSEIDON^π permutation.

The Round Function. Each round function of our POSEIDON permutation consists of the following three components.

1. *AddRoundConstants*, denoted by $\text{ARC}(\cdot)$
2. *SubWords*, denoted by $\text{S-Box}(\cdot)$ or by $\text{SB}(\cdot)$
3. *MixLayer*, denoted by $\text{M}(\cdot)$

The *MixLayer* operation is the linear layer of our construction, and it consists in multiplying the state with a $t \times t$ MDS matrix in order to apply the wide trail strategy.

In total we get:

$$\underbrace{\text{ARC} \rightarrow \text{SB} \rightarrow \text{M}}_{\text{First round}} \rightarrow \cdots \rightarrow \underbrace{\text{ARC} \rightarrow \text{SB} \rightarrow \text{M}}_{(R-1)\text{-th round}} \rightarrow \underbrace{\text{ARC} \rightarrow \text{SB} \rightarrow \text{M}}_{R\text{-th round}}$$

While $\text{ARC}(\cdot)$ and $\text{M}(\cdot)$ are the same in each round, the number of S-Boxes is not the same, namely

- $R_f + R_f = R_F$ rounds have *full* S-Box layers, i.e., t S-Box functions, and
- R_P rounds have *partial* S-Box layers, i.e., 1 S-Box and $(t - 1)$ identity functions.

We refer to [GLR⁺20] for more details about the HADES design strategy.

2.3 The Permutation Family Poseidon^π

The primary application of our design is hashing in large prime fields, so POSEIDON^π takes inputs of $t \geq 2$ words in \mathbb{F}_p , where p is a prime of size $p \approx 2^n$ (i.e., $\lceil \log_2(p) \rceil = n$).

The linear layer of POSEIDON^π consists of the multiplication with a fixed $t \times t$ MDS matrix⁷. The number of rounds $R = 2 \cdot R_f + R_P$ depends on the choice of the S-Box and on the parameters p and t .

⁷The schemes can easily be modified to also allow near-MDS matrices.

The S-Box Layer. For the applications we have in mind, we focus on two S-Boxes.

- First, we consider the α -power S-Box, defined by $\text{S-Box}(x) = x^\alpha$, where α is the smallest positive integer s.t. $\gcd(\alpha, p-1) = 1$. In the following, these permutations are called “ x^α -POSEIDON $^\pi$ ”. Examples are given by $\alpha = 3$ (x^3 -POSEIDON $^\pi$) if $p \not\equiv 1 \pmod 3$ or $\alpha = 5$ (x^5 -POSEIDON $^\pi$) if $p \not\equiv 1 \pmod 5$.
- Secondly, we consider the inverse S-Box x^{-1} (under the assumption $\text{S-Box}(0) = 0$). In the following, these permutations are called “ x^{-1} -POSEIDON $^\pi$ ”.

It turns out that the S-Box x^5 is suitable for two of the most popular prime fields in ZK applications, concretely the prime subfields of the scalar field of the BLS12-381 and BN254 curves, so we mainly consider this S-Box, but try to present generic cryptanalytic results whenever possible for other cases.

The Linear Layer. A $t \times t$ MDS matrix⁸ with elements in \mathbb{F}_p exists if the condition (see [MS78] for details)

$$2t + 1 \leq p$$

is satisfied.

Given p and t , there are several ways to construct an MDS matrix. One of them is using a Cauchy matrix [YMT97], which we recall here briefly. For $x_i, y_i \in \mathbb{F}_p$, where $i \in [1, t]$, the entries of the matrix \mathcal{M} are defined by

$$\mathcal{M}_{i,j} = \frac{1}{x_i + y_j},$$

where the entries of $\{x_i\}_{1 \leq i \leq t}$ and $\{y_i\}_{1 \leq i \leq t}$ are pairwise distinct and $x_i + y_j \neq 0$, where $i \in \{1, \dots, t\}$ and $j \in \{1, \dots, t\}$.

Avoiding Insecure Matrices. We emphasize that not every MDS matrix provides the same level of security. In particular, the matrix \mathcal{M} must prevent the possibility to set up

- (1) invariant (or iterative) subspace trails (or equivalently, truncated differentials) with prob. 1 with inactive S-Boxes over more than $t - 1$ rounds⁹ (more details are given in the following), or
- (2) invariant (or iterative) subspace trails with prob. 1 and with active S-Boxes for any number of rounds.

Regarding the first point, let \mathcal{S}^i be the subspace s.t. no S-Box is active in the first i consecutive rounds, that is,

$$\mathcal{S}^{(i)} := \{v \in \mathbb{F}^t \mid [\mathcal{M}^j \cdot v]_0 = 0 \in \mathbb{F}, j < i\}, \quad (1)$$

where $[x]_0$ denotes the first word of $x \in \mathbb{F}^t$, $\mathcal{S}^{(0)} = \mathbb{F}^t$, and $\dim(\mathcal{S}^{(i)}) \geq t - i$. For each pair of texts (x, y) in the same coset of \mathcal{S}^i , no S-Boxes are active in the first i consecutive rounds. Hence, a truncated differential with prob. 1 (or equivalently, a subspace trail) can be set up for the first $i \leq t - 1$ rounds.

⁸A matrix $M \in \mathbb{F}^{t \times t}$ is called *maximum distance separable* (MDS) iff it has a branch number $\mathcal{B}(M)$ equal to $\mathcal{B}(M) = t + 1$. The branch number of M is defined as $\mathcal{B}(M) = \min_{x \in \mathbb{F}^t} \{\text{wt}(x) + \text{wt}(M(x))\}$, where wt is the Hamming weight in wide trail terminology. Equivalently, a matrix M is MDS iff every submatrix of M is non-singular.

⁹This assumption is made to fix a weakness in the previous version of POSEIDON, where one can exploit the choice of the matrix \mathcal{M} over $(\mathbb{F}_{2^n})^t$ in order to set up an infinitely long subspace trail with no active S-Boxes for any number R_P of partial rounds. We refer to [KR20, BCD⁺20] for more details about this weakness.

The matrix \mathcal{M} must be chosen s.t. no subspace trail with inactive/active S-Boxes can be set up for more than $t - 1$ rounds. We refer to [GRS20] for more details about this approach and how to determine if a matrix \mathcal{M} satisfies this property.

Hence, we suggest the following method to generate matrices:

1. Randomly select pairwise distinct $\{x_i\}_{1 \leq i \leq t}$ and $\{y_i\}_{1 \leq i \leq t}$, where $x_i + y_j \neq 0$ and where $i \in \{1, \dots, t\}$ and $j \in \{1, \dots, t\}$.¹⁰
2. Determine if the matrix is secure using the tools provided in [GRS20]. For a secure matrix, no infinitely long (invariant and/or iterative) subspace trail (with or without active S-Boxes) can be set up for the rounds with partial S-Box layers.
3. Repeat this procedure until a secure matrix is found.

We used this method to generate the matrices for the instantiations given in Section 4. For [GRS20, Algorithm 3], we used a search period of $l = 4t$, and we additionally made sure that no invariant subspace trails with active S-Boxes exist for M, M^2, \dots, M^l . In our experiments, we observed that only a few trials are needed in order to find a secure matrix for sufficiently large fields.

Efficient Implementation. We refer to Supplementary Material B for details about efficient POSEIDON $^\pi$ implementations. The main advantage of these strategies consists of reducing the number of constant multiplications in each round with a partial S-Box layer from t^2 to $2t$, which is particularly useful for large t and R_P . For example, we implemented x^3 -POSEIDON $^\pi$ with $(n, t, R_F, R_P) = (64, 24, 8, 42)$ in `Sage`, and we could observe that the performance improves by a factor of about 5, with the average computation time being 4 ms for the optimized version.

3 Applications

We suggest POSEIDON for all applications of zero-knowledge friendly hashing, concretely:

- Using POSEIDON for commitments in various protocols, where the knowledge of the committed value is proven in zero knowledge. For this we suggest a single-call permutation-based hashing with POSEIDON-128 and width from 2 to 5 field elements.
- Hashing multi-element objects with certain fields encoded as field elements, so that statements about these fields are proven in zero-knowledge. We suggest variable-length sponge-based hashing with POSEIDON-128 or POSEIDON-80 with width 5 (and rate 4).
- Using POSEIDON in Merkle trees to enable zero-knowledge proofs of knowledge of a leaf in the tree with optional statements about the leaf content. We recommend Merkle trees of arity 4 (so width 5) with POSEIDON-128 as the most performant, but trees of more conventional arities can be used as well.

There exist several third-party protocols that already use POSEIDON in these usecases:

- Filecoin¹¹ employs POSEIDON for Merkle tree proofs with different arities and for two-value commitments.
- Dusk Network uses¹² POSEIDON to build a Zcash-like protocol for securities trading.

¹⁰Possible choices of x_i, y_j are also given in [KR20].

¹¹<https://github.com/filecoin-project/neptune>

¹²<https://github.com/dusk-network/Poseidon252>

- Sovrin uses POSEIDON for Merkle-tree based revocation [sov19].
- Loopring¹³ uses POSEIDON for private trading on Ethereum.

4 Concrete Instantiations of Poseidon π

As of mid-2020, many protocols that employ zero-knowledge proofs use (or plan to use) pairing-based proof systems [PHGR13, Gro16, GWC19, CHM⁺19] or Bulletproofs [BBB⁺18]. The elliptic curves used by these systems are predominantly BLS12-381, BN254, and Ed25519. A hash function friendly for such a system would operate in the scalar prime field of the curve, and they all have a size of around 2^{255} .

4.1 Main Instances

We present POSEIDON π permutations for such prime fields, and leave the other cases to the appendix. The S-Box function can be chosen as x^5 in all cases, i.e., we use x^5 -POSEIDON for hashing in all such protocols, though concrete fields are slightly different (this affects only constants and matrices, but not the number of rounds).

We further set the security level M to 128 bits as this is the maximum possible level for these curves. Smaller security levels are possible too. This corresponds to a 255-bit capacity, i.e., one field element. We focus on two possible widths, namely $t = 3$ and $t = 5$, as they correspond to popular cases of 2-to-1 and 4-to-1 compression functions. In the Merkle tree case, this corresponds to trees of arity 2 and 4, respectively. The round numbers for 80- and 128-bit security levels are given in Table 2, and a more extensive set is given in the appendix.

All our MDS matrices are Cauchy matrices, and the method to construct them is further described in Section 2.3. We use sequences of integers for the construction.

The round constants are generated using the Grain LFSR [HJMM08] in a self-shrinking mode, and the detailed initialization and generation are described in Supplementary Material F. Using this method, the generation of round constants and matrices depends on the specific instance, and thus different round constants are used even if some of the chosen parameters (e.g., n and t) are the same.

We provide the round constants, the matrices, and test vectors in auxiliary files for four primary instantiations. We also make reference implementations for these instantiations and scripts to calculate the round numbers, the round constants, and the MDS matrices available online.¹⁴ We refer to Supplementary Material A for a more detailed overview of the auxiliary files.

4.2 Domain Separation for Poseidon

We suggest the following strategy for the domain separation when using POSEIDON. A constant that specifies the domain goes to the capacity element, and due to its size (around 256 bits) we do not expect any security degradation. Concretely, we propose:

- **Merkle Tree** (all leaves are present, up to arity 32). The capacity is $2^{\text{arity}} - 1$. A generic case is considered in Supplementary Material I. We use no padding here.

¹³<https://tinyurl.com/y7tl537o>

¹⁴<https://extgit.iaik.tugraz.at/krypto/hadhash>

Table 2: Concrete instantiations of POSEIDON $^\pi$ (with security margin) over BLS12-381, BN254, Ed25519 scalar fields.

POSEIDON-128				
S-Box	$\log_2(p)$	t	R_F	R_P
x^5	255, 254	3	8	57
		5	8	60
POSEIDON-80				
S-Box	$\log_2(p)$	t	R_F	R_P
x^5	255, 254	3	8	33
		5	8	35
POSEIDON-256				
S-Box	$\log_2(p)$	t	R_F	R_P
x^5	255, 254	6	8	120
		10	8	120

- **Merkle Tree** (some leafs may be empty). The capacity value equals the bitmask of which leafs are present. We use no padding here.
- **Variable-Length Hashing.** The capacity value is 2^{64} . The padding consists of one field element being 1, and the remaining elements being 0.
- **Constant-Length Hashing.** The capacity value is $length \cdot (2^{64})$. The padding consists of the field elements being 0.
- **Encryption.** The capacity value is 2^{32} . The padding consists of the field elements being 0.
- **Future Uses.** The capacity value is $identifier \cdot (2^{32})$. The padding depends on the application.

5 Cryptanalysis Summary of Poseidon

As for any new design, it is paramount to present a concrete security analysis. In the following, we provide an in-depth analysis of the security of our construction. Due to a lack of any method to ensure that a hash function based on a sponge construction is secure against all possible attacks, we base our argumentation on the following consideration. As we just recalled in the previous section, when the internal permutation \mathcal{P} of an $(N = c + r)$ -bit sponge function is modeled as a randomly chosen permutation, the sponge hash function is indistinguishable from a random oracle up to $2^{c/2}$ calls to \mathcal{P} . Thus, we choose the number of rounds of the inner permutation case in order to provide security against distinguishers relevant to collision/preimage attacks. Equivalently, this means that such a number of rounds guarantees that \mathcal{P} does not exhibit any relevant non-random/structural properties (among the ones known in the literature).

5.1 Security Claims

In terms of concrete security, we expect it to be infeasible for all our hash functions to find collision attacks or preimage attacks with a cost substantially lower than 2^M evaluations of the permutation.

To help increase confidence in our design and simplify external cryptanalysis, we also explicitly state another claim about our internal permutation about the difficulty of the so-called constrained-input constrained-output problem of the permutation: *We expect it to be infeasible for all our permutations to solve the CICO problem by fixing m_1 bits of the input and m_2 bits of the output of the permutation with a cost substantially lower than $2^{\min(M, m_1, m_2)}$ evaluations of the permutation.*

Even though an attack below this threshold may not affect any concrete applications of our hash functions, we would still consider it an important cryptanalytic result.

5.2 Summary of Attacks

Here we list the main points of our cryptanalysis results. The number of rounds $R = R_P + R_F$ we can break depends on the security level M and the number of S-Boxes t , which we specify for each concrete hash function instance in the next section.

Before going on, we point out that for all attacks that are in common to the ones proposed for the cipher HadesMiMC [GLR⁺20], here we limit ourselves to report the main idea and result. For all other cases (namely, higher-order differentials, zero-sum partitions, Gröbner basis attacks, and preimage attacks), we present here more details. In any case, *all details are provided as supplementary material.*

We highlight that the following cryptanalysis is not equivalent to the one presented for the block cipher HADESMiMC. Indeed, the scenarios are different (in one case the goal is to guarantee the impossibility to find the secret key, while here there is no secret key material and the goal is to guarantee that the internal permutation looks like a pseudo-random permutation). This means that certain attacks that we consider here are not valid in the case of a block cipher and vice-versa. Just to give some examples, the rebound attack [LMR⁺09, MRST09] holds only in the context studied here, while a MitM scenario (crucial in the case of an SPN cipher) does not work in the context of a sponge function, since the attacker does not know the full output. More details are given in the following.

Proposition 1. (Informal) *The following number of rounds for x^5 -POSEIDON-128 over \mathbb{F}_p with ≈ 256 -bit p protects against statistical and algebraic attacks:*

$$R_F = 6, \quad R = R_F + R_P = 56 + \lceil \log_5(t) \rceil.$$

Proof. We substitute $\alpha = 5, M = 128$ and $\log_2(p) = 255$ to Equations (2),(3),(5) and see that no one is satisfied, i.e., the attacks do not work. \square

Proposition 2 (Informal). *The following number of rounds for x^5 -POSEIDON-80 over \mathbb{F}_p with ≈ 256 -bit p protects against statistical and algebraic attacks:*

$$R_F = 6, \quad R = R_F + R_P = 35 + \lceil \log_5(t) \rceil.$$

Proposition 3 (Informal). *The following number of rounds for x^5 -POSEIDON-256 over \mathbb{F}_p with ≈ 256 -bit p protects against statistical and algebraic attacks:*

$$R_F = 6, \quad R = R_F + R_P = 111 + \lceil \log_5(t) \rceil.$$

5.3 Security Margin

Given the *minimum* number of rounds necessary to provide security against all attacks known in the literature, we *arbitrarily* decided to add

- (1) two more rounds with full S-Box layers (+2 R_F), and
- (2) 7.5% more rounds with partial S-Box layers (+7.5% R_P).

We further round up the R_P value to the next multiple of w in order to allow for more efficient representation in circuits for ZK proofs.

The resulting number of rounds for our primary instances is given in Table 2.

5.4 Attack details

5.4.1 Statistical Attacks

Differential/Linear Distinguishers. As shown in the appendix, at least 6 rounds with full S-Box layers are necessary to provide security against the statistical attacks we consider. In more details, for

$$R_F < \begin{cases} 6 & \text{if } M \leq (\lfloor \log_2 p \rfloor - \mathcal{C}) \cdot (t + 1) \\ 10 & \text{otherwise} \end{cases} \quad (2)$$

linear [Mat93] and differential [BS91, BS93] attacks may be possible, where $\mathcal{C} = 2$ for $S(x) = 1/x$ and $\mathcal{C} = \log_2(\alpha - 1)$ for $S(x) = x^\alpha$ (where remember that α is an odd integer number), e.g. $\mathcal{C} = 1$ for $S(x) = x^3$ and $\mathcal{C} = 2$ for $S(x) = x^5$.

Before going on, we highlight that we exploit only rounds with full S-Box layers in order to prevent statistical attacks (as done in [GLR⁺20]). As explained in [KR20], under the assumption made for the linear layer in Section 2.3, it is actually possible to exploit both the rounds with partial and full S-Box layers in order to guarantee security against some statistical attacks, like differential and linear attacks. Our decision to consider only rounds with full S-Box layers has been made since a similar condition on the rounds with full S-Box layers (e.g., $R_F \geq 6$) is necessary for the security against some algebraic attacks (e.g., Gröbner basis attacks – see in the following) and in order to provide simple security arguments for all statistical attacks (including e.g. the rebound one).

(“Invariant”) Subspace Trails. We emphasize that the choice of the matrix that defines the linear layer made in Section 2.3 prevents the existence of subspaces \mathcal{S} that generate infinitely long subspace trails, namely a finite collection of subspaces $\{\mathcal{S}_0, \dots, \mathcal{S}_{r-1}\}$ s.t. each coset of \mathcal{S}_i is mapped into a coset of \mathcal{S}_{i+1} with probability 1 (where the index is taken modulo r) an arbitrary number of times. This allows to fix the weakness of the previous version of POSEIDON.

Other Attacks. Finally, we briefly mention that the same number of rounds given before for the case of differential/linear attacks guarantees security against other attacks as truncated differentials [Knu94], impossible differentials [BBS99], rebound attacks [LMR⁺09, MRST09], and so on. More details are given in the appendix.

5.4.2 Algebraic Attacks

In order to estimate the security against algebraic attacks, we evaluate the degree of the reduced-round permutations and their inverses. Roughly speaking, our results can be summarized as follows, where $n \simeq \log_2(p)$.

Interpolation Attack. The interpolation attack [JK97] depends on the number of different monomials in the interpolation polynomial, where (an upper/lower bound of) the number of different monomials can be estimated given the degree of the function. The idea of such an attack is to construct an interpolation polynomial that describes the function. If the number of monomials is too large, then such a polynomial cannot be constructed faster than via a brute-force attack.

For a security level of M bits, the number of rounds that can be attacked is

- for $S(x) = x^\alpha$:

$$R \leq \lceil \log_\alpha(2) \cdot \min\{M, \log_2(p)\} \rceil + \lceil \log_\alpha t \rceil \quad (3)$$

- for $S(x) = 1/x$:

$$\lceil R_F \log_2(t) \rceil + R_P \leq \lceil \log_2(t) \rceil + \lceil 0.5 \cdot \min\{M, \log_2(p)\} \rceil \quad (4)$$

Gröbner Basis Attack. In a Gröbner basis attack [CLO97], one tries to solve a system of non-linear equations that describe the function. The cost of such an attack depends on the degree of the equations, but also on the number of equations and on the number of variables. Since there are several ways in which it is possible to describe the studied permutation, there are several ways to set up such a system of equations and so the attack. Here, we focus on two extreme cases:

1. In the first case, the attacker derives equations, one by word, for the entire r -round permutation. Assuming $S(x) = x^\alpha$ (analogous for the others), we show that the attack complexity is about α^{2t} (see below), therefore for a security level of M bits the attack works at most on $\log_\alpha 2^{\min\{n/2, M/2\}}$ rounds.
2. In the second case, since a partial S-Box layer is used, it may be more efficient to consider degree- α equations for single S-Boxes. In this case, more rounds can be necessary to guarantee security against this attack.

In both cases, it is possible to make use of the existence of the subspace $\mathcal{S}^{(r)}$ defined as in Eq. (1) in order to improve the attack. As shown in [BCD⁺20], such a subspace can be exploited in order to replace some non-linear equations of the system that we are trying to solve with linear equations. Indeed, given a text in a coset of the subspace $\mathcal{S}^{(r)}$, the output of such a text after r rounds with partial S-Box layers is simply the result of an affine map applied to the input (i.e., no S-Box is involved). As explained in detail in Supplementary Material C.2.2, this issue can easily be fixed both by a careful choice of the matrix that defines the linear layer (see Section 2.3 for details) and, if necessary, by adjusting the number of rounds with partial S-Box layers.¹⁵

¹⁵In a previous alternative version of POSEIDON ^{π} over $\mathbb{F}_{2^n}^t$, some of the matrices M that defined the linear layer had a low multiplicative order when t was a power of 2, namely $M^2 = \mu \cdot I$ for a certain $\mu \in \mathbb{F}_{2^n}$ and for $t = 2^r$ (where I is the identity matrix). This fact can be exploited in order to set up a preimage attack using the strategy just described. In particular, it is easy to find a subspace $\mathcal{S} \subseteq \mathbb{F}_{2^n}^t$ of dimension $t - 2$ s.t. for each pair of texts x, y in the same coset of \mathcal{S} :

$$R_P^{(1)}(x) - R_P^{(1)}(y) \in \mathcal{M}(\mathcal{S}) \quad \text{and} \quad R_P^{(2)}(x) - R_P^{(2)}(y) \in \mathcal{S},$$

With optimistic (for the adversary) complexity of the Gaussian elimination, we obtain for each S-Box two attacks which are faster than 2^M if either condition is satisfied:

- if $S(x) = x^\alpha$:

$$\begin{cases} R \leq \log_\alpha(2) \cdot \min \left\{ \frac{M}{3}, \frac{\log_2(p)}{2} \right\}, \\ R \leq t - 1 + \min \left\{ \frac{\log_\alpha(2) \cdot M}{t+1}, \frac{\log_\alpha(2) \cdot \log_2(p)}{2} \right\} \end{cases} \quad (5)$$

- if $S(x) = 1/x$:

$$\begin{cases} \lfloor R_F \log_2(t) \rfloor + R_P \leq \lceil 0.5 \cdot \min\{M, \log_2(p)\} \rceil + \lceil \log_2(t) \rceil \\ \lfloor R_F \log_2(t) \rfloor + R_P \leq t - 1 + \lceil \log_2(t) \rceil + \min \left\{ \left\lceil \frac{M}{t+1} \right\rceil, \lceil 0.5 \cdot \log_2(p) \rceil \right\} \end{cases} \quad (6)$$

Higher-Order Differential Attack. Working over $\mathbb{F}_{2^n}^t \equiv \mathbb{F}_2^{n \cdot t}$, the higher-order differential attack [Knu94] depends on the *algebraic degree* of the polynomial function that defines the permutation, where the algebraic degree δ of a function $f(x) = x^d$ of degree d over \mathbb{F}_{2^n} is defined as $\delta = \text{hw}(d)$ (where $\text{hw}(\cdot)$ is the hamming weight). The idea of such an attack is based on the property that given a function $f(\cdot)$ of algebraic degree δ , $\bigoplus_{x \in \mathcal{V} \oplus \phi} f(x) = 0$ if the dimension of the subspace \mathcal{V} satisfies $\dim(\mathcal{V}) \geq \delta + 1$. If the algebraic degree is sufficiently high, the attack does not work.

At first thought, one may think that this attack does not apply (or is much less powerful) in \mathbb{F}_p^t (due to the fact that the only subspaces of \mathbb{F}_p are $\{0\}$ and \mathbb{F}_p itself). Recently, it has been shown in [BCD⁺20] how to set up an higher-order differential over \mathbb{F}_p^t . Given f over \mathbb{F}_p of degree $d \leq p - 2$, $\sum_{x \in \mathbb{F}_p} f(x) = 0$. Since this result is related to the degree of the polynomial that describes the permutation, we claim that the number of rounds necessary to provide security against the interpolation attack provides security against this attack as well.

Zero-Sum Partitions. Another possible distinguisher that can be set up for some inner primitive in a hash function (with a relatively low degree) is based on the *zero-sum partition*. This direction has been investigated e.g. in [BCD11] for two SHA-3 candidates, *Luffa* and *KECCAK*. More generally, a zero-sum structure for a function $f(\cdot)$ is defined as a set Z of inputs z_i that sum to zero, and for which the corresponding outputs $f(z_i)$ also sum to zero, i.e., $\sum_i z_i = \sum_i f(z_i) = 0$. For an iterated function, the existence of zero sums is usually due to the particular structure of the round function or to a low degree. Since it is expected that a randomly chosen function does not have many zero sums, the existence of several such sets can be seen as a distinguishing property of the internal function.

Definition 1 (Zero-sum Partition [BCD11]). Let P be a permutation over \mathbb{F}_q^t for a prime $q \geq 2$. A zero-sum partition for P of size $K < t$ is a collection of K disjoint sets $\{X_1, \dots, X_K\}$ with the following properties:

- $X_i \subset \mathbb{F}^t$ for each $i = 1, \dots, k$ and $\bigcup_{i=1}^k X_i = \mathbb{F}^t$,
- $\forall i = 1, \dots, K$: the set X_i satisfies the zero-sum property $\sum_{x \in X_i} x = \sum_{x \in X_i} P(x) = 0$.

Here we explicitly state that we do not make claims about the security of POSEIDON^π against zero-sum partitions. This choice is motivated by the gap present in the literature

where $R_P^{(r)}$ denotes r partial rounds and where $\mathcal{M}(\mathcal{S}) = \{\mathcal{M} \cdot x \mid \forall x \in \mathcal{S}\}$ (that is, given an input in a coset of \mathcal{S} , the output is simply the result of an affine map applied to the input). This weakness has been fixed in this new version of POSEIDON^π , both due to a clever choice of the matrix that defines the linear layer and by adjusting the number of rounds.

between the number of rounds of the internal permutation that can be covered by a zero-sum partition and by the number of rounds in the corresponding sponge hash function that can be broken e.g. via a preimage or a collision attack. As a concrete example, consider the case of KECCAK: While 24 rounds of KECCAK- f can be distinguished from a random permutation using a zero-sum partition [BCD11] (that is, *full* KECCAK- f), attacks on KECCAK can only be set up for up to 6 rounds of KECCAK- f [GLL⁺20].

For completeness, we mention that a zero-sum partition on (a previous version of) round-reduced POSEIDON ^{π} has been proposed in [BCD⁺20]. Such an attack can cover up to $R_F = 6$ rounds (i.e., 2 rounds at the beginning and 4 rounds at the end) by exploiting the inside-out approach and by choosing a subspace of texts after the first R_f rounds with full S-Box layers and before the R_P rounds with partial S-Box layers. Since the number of rounds of this new version is not smaller than the number of rounds of the previous one, and since $R_F \geq 8$ (see Section 5.3), it seems that a zero-sum partition cannot be set up for full POSEIDON ^{π} .

6 Poseidon in Zero-Knowledge Proof Systems

Our hash functions have been designed to be friendly to zero-knowledge applications. Specifically, we aim to minimize proof generation time, proof size, and verification time (when it varies). Before presenting concrete results, we give a small overview of ZK proof systems to date.

6.1 State of the Art

Let \mathcal{P} be a circuit over some finite field \mathbb{F} where gates are some (low-degree) polynomials over \mathbb{F} with I and O being input and output variables, respectively: $\mathcal{P}(I) = O$. The *computational integrity problem* consists of proving that some given O_0 is the result of the execution of \mathcal{P} over some I_0 : $\mathcal{P}(I_0) = O_0$. It is not difficult to show that any limited-time program on a modern CPU can be converted to such a circuit [BCTV14], and making the proof zero-knowledge is often possible with little overhead.

The seminal PCP series of papers states that for any program \mathcal{P} it is possible to construct a proof of computational integrity, which can be verified in time sublinear in the size of \mathcal{P} . However, for a long time the prover algorithms were so inefficient that this result remained merely theoretical. Only recently, proof systems where the prover costs are polynomial in $|\mathcal{P}|$ were constructed, but they required a trusted setup: a verifier or someone else (not the prover) must process the circuit with some secret s and output a reference string S , used both by the prover and the verifier. In this setting, the prover’s work can even be made linear in $|\mathcal{P}|$, and the verifier’s costs are constant. These systems were called SNARKs for proof succinctness. The first generation of SNARKs, known as Pinocchio and Groth16 [PHGR13, Gro16], require a separate trusted setup for each circuit. The next generation, which includes Sonic [MBKM19], PLONK [GWC19], and Marlin [CHM⁺19], can use one reference string of size d for all circuits with at most d gates, thus simplifying the setup and its reuse. Later on, proof systems without trusted setups appeared, of which we consider Bulletproofs [BBB⁺18], STARKs [BBHR18], and RedShift [KPV19] the most interesting, though all of them come with deficiencies: Bulletproofs have linear verifier times (but rather short proofs), STARKs work with iterative programs, and RedShift has big proofs (up to 1 MB for millions of gates).

Current benchmarks demonstrate that programs with millions of gates can be processed within a few seconds with the fastest proof systems, which solves the computational integrity

problem for some practical programs. Among them, privacy-preserving cryptocurrencies, mixers, and private voting are prominent examples. In short, such applications work as follows:

1. Various users add publicly hashes of some secret and public values to some set V , which is implemented as a Merkle tree. Hashes can be currency transaction digests, public keys, or other credentials.
2. Only those who know a secret behind some hash are declared eligible for an action (e.g., to vote or to spend money).
3. A user who wants to perform the action proves that they know a tree leaf L and a secret K such that L is both the hash of K and a leaf in V . If the proof passes, the user is allowed to perform an action (e.g., to vote). If an action must be done only once, a deterministic hash of the secret and leaf position can be computed and published.

This paradigm is behind the cryptocurrency Zcash and the Ethereum mixer Tornado¹⁶.

The bottleneck of such a system is usually the proof creation time, which took 42 seconds in the early version of Zcash, and sometimes the verifier's time. Both are determined by the size of the circuit that describes a Merkle proof and are thus dependent on the hash function that constitutes the tree.

Unfortunately, a single hash function cannot be optimal for all ZK proof systems, because they use different arithmetizations: STARKs can use prime and binary fields, Bulletproofs uses any prime field, whereas most SNARKs use a prime field based on a scalar field of a pairing-friendly elliptic curve. Therefore, for each proof system a new instance of POSEIDON may be needed. In the future text we describe how this is done and how to optimize a circuit for some proof systems.

6.2 SNARKs with Poseidon ^{π}

In SNARKs, the prime field is typically the scalar field of some pairing-friendly elliptic curve. The primitive POSEIDON ^{π} can be represented as such a circuit with reasonably few gates, but the parameters of POSEIDON ^{π} must have been determined first by p . Concretely, after p is fixed, we first check if x^α is invertible in $\text{GF}(p)$, which is true if $p \bmod \alpha \neq 1$. If this inequality is not satisfied for a small α , we use the inverse S-Box or consider another prime power for the S-Box.

6.2.1 Groth16

Groth16 [Gro16] is an optimization of the Pinocchio proof system and currently the fastest SNARK with the smallest proofs. The Groth16 prover complexity is $\mathcal{O}(s)$, where s is the number of rank-1 constraints – quadratic equations of the form $(\sum_i u_i X_i)(\sum_i v_i X_i) = \sum_i w_i X_i$, where u_i, v_i, w_i are field elements and X_i are program variables. It is easy to see that the S-Box x^3 is represented by 2 constraints, the S-Box x^5 by 3 constraints, and the S-Box $1/x$ by 3 constraints (1 for the non-zero case, and two more for the zero case).

¹⁶<https://tinyurl.com/yd7j3w5h>

Thus, in total we have

$$\begin{aligned} &2tR_F + 2R_P \text{ constraints for } x^3\text{-POSEIDON}^\pi, \\ &3tR_F + 3R_P \text{ constraints for } x^5\text{-POSEIDON}^\pi, \\ &3tR_F + 3R_P \text{ constraints for } x^{-1}\text{-POSEIDON}^\pi. \end{aligned}$$

It requires a bit more effort to see that we do not need more constraints as the linear layers and round constants can be incorporated into these ones. However, it is necessary to do some preprocessing. For example, in the POSEIDON^π setting, the full S-Box layers are followed by a linear transformation M . Each round with a full S-Box layer can be represented by the following constraints in the SNARK setting:

$$\begin{aligned} \left(\sum_j M_{i,j}x_{i,j}\right) \cdot \left(\sum_j M_{i,j}x_{i,j}\right) &= y_i \quad 1 \leq i \leq t, \\ y_i \cdot \left(\sum_j M_{i,j}z_{i,j}\right) &= z_i, \end{aligned}$$

where $M = I_{t \times t}$ for the first round. However, in a round with a partial S-Box layer, we will have only one such constraint for $j = 1$. For the rest of the $t - 1$ variables we will have linear constraints of the form

$$\sum_j M_{i,j}x_{i,j} = u_i, \text{ where } 2 \leq i \leq t.$$

Since the linear constraints have little complexity effect in Groth16, in the partial S-Box rounds they can be composed with the ones from the previous round(s) using

$$\sum_k M_{i,k} \left(\sum_j M_{i,j}x_{i,j}\right) = v_k \quad 2 \leq k \leq t.$$

We can now calculate the number of constraints for the sponge mode of operation and for Merkle trees. In sponges, the $2M$ bits are reserved for the capacity, so $N - 2M$ bits are fed with the message. Therefore, we get

- $\frac{2tR_F + 2R_P}{N - 2M}$ constraints per bit for $x^3\text{-POSEIDON}^\pi$,
- $\frac{3tR_F + 3R_P}{N - 2M}$ constraints per bit for $x^5\text{-POSEIDON}^\pi$,
- $\frac{3tR_F + 3R_P}{N - 2M}$ constraints per bit for $x^{-1}\text{-POSEIDON}^\pi$.

For the Merkle tree, we suggest a 1-call sponge where all branches must fit into the rate. Then a Merkle tree has arity $\frac{N}{2M} - 1$. Based on that we can calculate how many constraints we need to prove the opening in a Merkle tree of, for example, 2^{32} elements (the recent ZCash setting). The tree will have $\frac{32}{\log_2[\frac{N}{2M} - 1]}$ levels with the number of constraints in each according to the above. The libsnark performance of POSEIDON is given in Table 3.

As an example, we calculate the concrete number of constraints for a Merkle tree proof, where the tree has 2^{30} elements, assuming a security level of 128 bits and a prime field of size close to 2^{256} . We take the S-Box equal to x^5 as it fits many prime fields: Ristretto (the prime group based on the scalar field of Ed25519), BN254, and BLS12-381 scalar fields. The results are in Table 4.

Table 3: `libsnark`[lib] performance of x^5 -POSEIDON over a prime field and with 127-bit collision resistance.

Arity	Width	libsnark ZK proof time for one hash		R1CS constraints
		Prove	Verify	
2:1	3	57.5ms	1.2ms	225
3:1	4	84.0ms	1.2ms	243
8:1	9	120.2ms	1.3ms	335

Table 4: Number of constraints for a circuit proving a Merkle tree entry in the set of 2^{30} elements, with the S-Box x^5 over a suitable prime field with 128-bit collision resistance.

POSEIDON				
Arity	Width	R_F	R_P	Total constraints
2:1	3	8	57	7290
4:1	5	8	60	4680
8:1	9	8	63	4050
Pedersen hash				
510	171	-	-	43936
<i>Rescue</i>				
2:1	3	16	-	7920
4:1	5	10	-	4500
8:1	9	10	-	5400

6.2.2 Bulletproofs

Bulletproofs [BBB⁺18] is a proof system that does not require a trusted setup. It is notable for short proofs which are logarithmic in the program size, and also for the shortest range proofs that do not require a trusted setup. However, its verifier is linear in the program size. For the use cases where the trusted setup is not an option, the Bulletproofs library ‘dalek’ is among the most popular ZK primitives. We have implemented¹⁷ a Merkle tree prover for POSEIDON in Bulletproofs using the same constraint system as for Groth16 with results outlined in Table 5. The performance varies since the underlying curves are based on prime fields of different size and weight: BN254 uses a 254-bit prime whereas BLS12-381 uses a 381-bit one (the reason for that is the recent reevaluation of discrete logarithm algorithms specific to pairing-friendly curves).

6.2.3 PLONK

PLONK [GWC19] is a novel but popular SNARK using a universal trusted setup, where a structured reference string of size d can be used for any circuit of d gates or less. The setup is pretty simple as for the secret k the values $k^d \cdot B$ are stored, where B is an elliptic curve point and \cdot denotes scalar multiplication. The setup is used for short polynomial commitments and proofs by Kate et al. [KZG10].

¹⁷https://github.com/lovesh/bulletproofs-r1cs-gadgets/blob/master/src/gadget_poseidon.rs

Table 5: Bulletproofs performance to prove 1 out of 2^{30} -Merkle tree.

Field	Arity	Merkle 2^{30} -tree ZK proof		R1CS constraints
		Prove	Verify	
POSEIDON-128				
BLS12-381	2:1	16.8s	1.5s	7200
	4:1	13.8s	1.65s	4320
	8:1	11s	1.4s	3870
BN254	2:1	11.2s	1.1s	7200
	4:1	9.6s	1.15s	4320
	8:1	7.4s	1s	3870
Ristretto	2:1	8.4s	0.78s	7200
	4:1	6.45s	0.72s	4320
	8:1	5.25s	0.76s	3870

Table 6: PLONK performance to prove a 1-out-of- 2^n -Merkle tree of arity 4.

Field	Set size	Merkle 2^n -tree ZK proof		R1CS constraints
		Prove	Verify	
POSEIDON-128				
BLS12-381	2^{16}	3.59s	0.7ms	1920
	2^{34}	13.8s	1.65s	4080
	2^{68}	9.9s	2.7ms	8160

The standard version of PLONK works with the same constraint system as we have described, plus it uses special machinery to lay out wires in the circuit. A prover first crafts three polynomials of degree equal to the number of gates, which are responsible for left input, right input, and output, respectively. Then they allocate several supplementary polynomials to describe the wire layout. The prover complexity for a POSEIDON^π permutation with the S-Box x^5 of width w and R rounds is $11(w(w+6)+3)R$ point multiplications, and the proof has 7 group elements and 7 field elements. A third-party non-optimized implementation of a PLONK prover in Rust (Intel(R) Core(TM) i5-7300HQ CPU @ 2.50GHz) gives us benchmarks, which we provide in Table 6.

As we have almost identical rounds, the PLONK compiler can be heavily optimized. Concretely, we suggest the following.

- Define a separate polynomial for each S-Box line.
- Get rid of wire layout polynomials.
- Express round transitions as a system of affine equations over polynomial values at adjacent points.

As a result, our optimized PLONK compiler makes only $(w+11)R$ point multiplications for a single permutation call, whereas the proof consists of $((w+3)$ group elements and $2w$ field elements. This might bring a 25-40x increase in performance depending on w .

6.2.4 RedShift

RedShift [KPV19] is a STARK-inspired proof system which works with an arbitrary set of constraints. It can be viewed as PLONK with pairing-based polynomial commitments with

the trusted setup being replaced by Reed-Solomon trustless commitments. The RedShift proof is $c_\lambda \log d^2$ KB large, where d is the degree of circuit polynomials and $c_\lambda \approx 2.5$ for a 120-bit security. Due to similarity, we can make the same optimizations as in PLONK, so that the entire Merkle tree proof requires polynomials of degree 4800 for width 5, resulting in the entire proof being around 12 KB in size. Unfortunately, no RedShift library is publicly available so far, and hence we could not measure the actual performance.

6.3 STARKs with Poseidon $^\pi$

ZK-STARKs [BBHR18] is a proof system for the computational integrity, which is not vulnerable to quantum computers and does not use a trusted setup. STARKs operate with programs whose internal state can be represented as a set of w registers, each belonging to a binary field $\mathbb{GF}(2^n)$ or to a 2^n -subgroup \mathbb{G} of a prime-order group (this is our primary case, as the scalar fields of BLS12-381 and BN254 have such a big subgroup).

The program execution is then represented as a set of T internal states. The computational integrity is defined as the set of all wT registers satisfying certain s polynomial equations (constraints) of degree d .

STARK Costs

According to [Sta18], the number of constraints does not play a major role in the prover, verifier, or communication complexity, which are estimated as follows:

- $8w \cdot T \cdot d \cdot \log(wT)$ operations in \mathbb{G} for the prover,
- a prover memory in $\Omega(w \cdot T \cdot n)$, and
- a communication (verifier time) of $n \cdot (m + \log^2(8Td))$,

where m is the maximum number of variables in a constraint polynomial.

The primitive POSEIDON $^\pi$ can be represented as such a program with few registers, a small number of steps, and low degree. Following the same approach as for SNARKs in Section 6.2, we keep in registers only S-Box inputs and the permutation outputs. Setting $w = t$, we get $T = R_F + \lceil R_P/t \rceil$ and $wT = tR_F + R_P$. Thus, the complexity is as follows:

- $24(tR_F + R_P) \cdot \log_2(tR_F + R_P)$ operations in \mathbb{G} for the prover,
- a prover memory in $\Omega(63 \cdot (tR_F + R_P))$, and
- a communication (verifier time) of $63 \cdot (t + \log_2^2(24(tR_F + R_P)))$.

We suggest $t \in \{3, 5\}$ in order to support the same Merkle tree cases as before. Thus, for our primary instance POSEIDON-128, we get an AET cost of 20540 for each permutation call for a width of 3. As we process 510 bits per call, we obtain a prover complexity of 40 operations per bit. For a width of 5 we get an AET cost of 38214, which translates to 38 operations per bit in \mathbb{G} .

References

- [AÅBL12] Mohamed Ahmed Abdelraheem, Martin Ågren, Peter Beelen, and Gregor Leander. On the Distribution of Linear Biases: Three Instructive Examples. In *CRYPTO 2012*, volume 7417 of *LNCS*, pages 50–67, 2012.

- [AABS⁺19] Abdelrahman Aly, Tomer Ashur, Eli Ben-Sasson, Siemen Dhooghe, and Alan Szepieniec. Design of symmetric-key primitives for advanced cryptographic protocols. Cryptology ePrint Archive, Report 2019/426, 2019. <https://eprint.iacr.org/2019/426>.
- [ACG⁺19] Martin R. Albrecht, Carlos Cid, Lorenzo Grassi, Dmitry Khovratovich, Reinhard Lüftenegger, Christian Rechberger, and Markus Schofnegger. Algebraic Cryptanalysis of STARK-Friendly Designs: Application to MARVELLous and MiMC. In *ASIACRYPT 2019*, volume 11923 of *LNCS*, pages 371–397, 2019.
- [AD18] Tomer Ashur and Siemen Dhooghe. Marvellous: a stark-friendly family of cryptographic primitives. Cryptology ePrint Archive, Report 2018/1098, 2018. <https://eprint.iacr.org/2018/1098>.
- [AHIV17] Scott Ames, Carmit Hazay, Yuval Ishai, and Muthuramakrishnan Venkatasubramanian. Liger: Lightweight sublinear arguments without a trusted setup. In *CCS*, pages 2087–2104. ACM, 2017.
- [BBB⁺18] Benedikt Bünz, Jonathan Bootle, Dan Boneh, Andrew Poelstra, Pieter Wuille, and Gregory Maxwell. Bulletproofs: Short proofs for confidential transactions and more. In *IEEE Symposium on Security and Privacy*, pages 315–334. IEEE Computer Society, 2018.
- [BBHR18] Eli Ben-Sasson, Iddo Bentov, Yinon Horesh, and Michael Riabzev. Scalable, transparent, and post-quantum secure computational integrity. *IACR Cryptology ePrint Archive*, 2018:46, 2018.
- [BBS99] Eli Biham, Alex Biryukov, and Adi Shamir. Cryptanalysis of Skipjack Reduced to 31 Rounds Using Impossible Differentials. In *EUROCRYPT 1999*, volume 1592 of *LNCS*, pages 12–23, 1999.
- [BCD11] Christina Boura, Anne Canteaut, and Christophe De Cannière. Higher-Order Differential Properties of Keccak and *Luffa*. In *FSE 2011*, volume 6733 of *LNCS*, pages 252–269, 2011.
- [BCD⁺20] Tim Beyne, Anne Canteaut, Itai Dinur, Maria Eichlseder, Gregor Leander, Gaëtan Leurent, María Naya-Plasencia, Léo Perrin, Yu Sasaki, Yosuke Todo, and Friedrich Wiemer. Out of Oddity – New Cryptanalytic Techniques against Symmetric Primitives Optimized for Integrity Proof Systems. Cryptology ePrint Archive, Report 2020/188, to appear at CRYPTO’2020, 2020. <https://eprint.iacr.org/2020/188>.
- [BCLR17] Christof Beierle, Anne Canteaut, Gregor Leander, and Yann Rotella. Proving Resistance Against Invariant Attacks: How to Choose the Round Constants. In *CRYPTO 2017*, volume 10402 of *LNCS*, pages 647–678, 2017.
- [BCTV14] Eli Ben-Sasson, Alessandro Chiesa, Eran Tromer, and Madars Virza. Succinct non-interactive zero knowledge for a von neumann architecture. In *USENIX Security Symposium*, pages 781–796. USENIX Association, 2014.
- [BDPA08] Guido Bertoni, Joan Daemen, Michaël Peeters, and Gilles Van Assche. On the Indifferentiability of the Sponge Construction. In *EUROCRYPT 2008*, volume 4965 of *LNCS*, pages 181–197, 2008.
- [BFSY05] M Bardet, JC Faugere, B Salvy, and BY Yang. Asymptotic behaviour of the index of regularity of quadratic semi-regular polynomial systems. In *The Effective Methods in Algebraic Geometry Conference (MEGA)*, pages 1–14, 2005.

- [BS91] Eli Biham and Adi Shamir. Differential Cryptanalysis of DES-like Cryptosystems. *Journal of Cryptology*, 4(1):3–72, 1991.
- [BS93] Eli Biham and Adi Shamir. *Differential Cryptanalysis of the Data Encryption Standard*. Springer, 1993.
- [CGL⁺20] Carlos Cid, Lorenzo Grassi, Reinhard Lüftenegger, Christian Rechberger, and Markus Schofnegger. Higher-Order Differentials of Word-Oriented SPN Schemes with Low-Degree S-Boxes. Cryptology ePrint Archive, Report 2020/536, 2020. <https://eprint.iacr.org/2020/536>.
- [CHM⁺19] Alessandro Chiesa, Yuncong Hu, Mary Maller, Pratyush Mishra, Noah Vesely, and Nicholas Ward. Marlin: Preprocessing zkSNARKs with universal and updatable SRS. Cryptology ePrint Archive, Report 2019/1047, 2019. <https://eprint.iacr.org/2019/1047>.
- [CKS09] Jan Camenisch, Markulf Kohlweiss, and Claudio Soriente. An accumulator based on bilinear maps and efficient revocation for anonymous credentials. In *Public Key Cryptography*, volume 5443 of *LNCS*, pages 481–500. Springer, 2009.
- [CL02] Jan Camenisch and Anna Lysyanskaya. Dynamic Accumulators and Application to Efficient Revocation of Anonymous Credentials. In *CRYPTO 2002*, volume 2442 of *LNCS*, pages 61–76. Springer, 2002.
- [CLO97] David A. Cox, John Little, and Donal O’Shea. *Ideals, varieties, and algorithms - an introduction to computational algebraic geometry and commutative algebra (2. ed.)*. Undergraduate texts in mathematics. Springer, 1997.
- [DKP⁺19] Itai Dinur, Daniel Kales, Angela Promitzer, Sebastian Ramacher, and Christian Rechberger. Linear Equivalence of Block Ciphers with Partial Non-Linear Layers: Application to LowMC. In *EUROCRYPT 2019*, volume 11476 of *LNCS*, pages 343–372, 2019.
- [DKR97] Joan Daemen, Lars R. Knudsen, and Vincent Rijmen. The Block Cipher Square. In *FSE 1997*, volume 1267 of *LNCS*, pages 149–165, 1997.
- [DR01] Joan Daemen and Vincent Rijmen. The wide trail design strategy. In *IMACC*, volume 2260 of *Lecture Notes in Computer Science*, pages 222–238. Springer, 2001.
- [DR02] Joan Daemen and Vincent Rijmen. *The Design of Rijndael: AES - The Advanced Encryption Standard*. Information Security and Cryptography. Springer, 2002.
- [EGL⁺20] Maria Eichlseder, Lorenzo Grassi, Reinhard Lüftenegger, Morten Øygaard, Christian Rechberger, Markus Schofnegger, and Qingju Wang. An Algebraic Attack on Ciphers with Low-Degree Round Functions: Application to Full MiMC. Cryptology ePrint Archive, Report 2020/182, 2020. <https://eprint.iacr.org/2020/182>.
- [Gen07] Giulio Genovese. Improving the algorithms of berlekamp and niederreiter for factoring polynomials over finite fields. *J. Symb. Comput.*, 42(1-2):159–177, 2007.
- [GKK⁺19] Lorenzo Grassi, Daniel Kales, Dmitry Khovratovich, Arnab Roy, Christian Rechberger, and Markus Schofnegger. Starkad and Poseidon: New Hash Functions for Zero Knowledge Proof Systems. IACR Cryptology ePrint Archive, Report 2019/458, 2019. Version: 20200205:104144.

- [GLL⁺20] Jian Guo, Guohong Liao, Guozhen Liu, Meicheng Liu, Kexin Qiao, and Ling Song. Practical Collision Attacks against Round-Reduced SHA-3. *Journal of Cryptology*, 33(1):228–270, 2020.
- [GLR⁺20] Lorenzo Grassi, Reinhard Lüftenegger, Christian Rechberger, Dragos Rotaru, and Markus Schofnegger. On a Generalization of Substitution-Permutation Networks: The HADES Design Strategy. In *EUROCRYPT 2020*, volume 12106 of *LNCS*, pages 674–704, 2020.
- [GMO16] Irene Giacomelli, Jesper Madsen, and Claudio Orlandi. ZKBoo: Faster Zero-Knowledge for Boolean Circuits. In *USENIX Security Symposium*, pages 1069–1083. USENIX Association, 2016.
- [Gra18] Lorenzo Grassi. Mixture Differential Cryptanalysis: a New Approach to Distinguishers and Attacks on round-reduced AES. *IACR Trans. Symmetric Cryptol.*, 2018(2):133–160, 2018.
- [Gro16] Jens Groth. On the size of pairing-based non-interactive arguments. In *EUROCRYPT 2016*, volume 9666 of *LNCS*, pages 305–326. Springer, 2016.
- [GRR17] Lorenzo Grassi, Christian Rechberger, and Sondre Rønjom. A New Structural-Differential Property of 5-Round AES. In *EUROCRYPT 2017*, volume 10211 of *LNCS*, pages 289–317, 2017.
- [GRS20] Lorenzo Grassi, Christian Rechberger, and Markus Schofnegger. Weak Linear Layers in Word-Oriented Partial SPN and HADES-Like Ciphers. Cryptology ePrint Archive, Report 2020/500, 2020. <https://eprint.iacr.org/2020/500>.
- [GWC19] Ariel Gabizon, Zachary J. Williamson, and Oana Ciobotaru. PLONK: permutations over lagrange-bases for oecumenical noninteractive arguments of knowledge. *IACR Cryptology ePrint Archive*, 2019:953, 2019.
- [HBHW19] Daira Hopwood, Sean Bowe, Taylor Hornby, and Nathan Wilcox. Zcash protocol specification: Version 2019.0-beta-37 [overwinter+sapling]. Technical report, Zerocoin Electric Coin Company, 2019. available at <https://github.com/zcash/zips/blob/master/protocol/protocol.pdf>.
- [HJMM08] Martin Hell, Thomas Johansson, Alexander Maximov, and Willi Meier. The Grain Family of Stream Ciphers. In *The eSTREAM Finalists*, volume 4986 of *LNCS*, pages 179–190. Springer, 2008.
- [JK97] Thomas Jakobsen and Lars R. Knudsen. The Interpolation Attack on Block Ciphers. In *FSE 1997*, volume 1267 of *LNCS*, pages 28–40, 1997.
- [JNP13] Jérémy Jean, María Naya-Plasencia, and Thomas Peyrin. Multiple Limited-Birthday Distinguishers and Applications. In *SAC 2013*, volume 8282 of *LNCS*, pages 533–550, 2013.
- [Knu94] Lars R. Knudsen. Truncated and Higher Order Differentials. In *FSE 1994*, volume 1008 of *LNCS*, pages 196–211, 1994.
- [KPV19] Assimakis Kattis, Konstantin Panarin, and Alexander Vlasov. Redshift: Transparent snarks from list polynomial commitment iops. Cryptology ePrint Archive, Report 2019/1400, 2019. <https://eprint.iacr.org/2019/1400>.
- [KR20] Nathan Keller and Asaf Rosemarin. Mind the Middle Layer: The HADES Design Strategy Revisited. Cryptology ePrint Archive, Report 2020/179, 2020. <https://eprint.iacr.org/2020/179>.

- [KZG10] Aniket Kate, Gregory M. Zaverucha, and Ian Goldberg. Constant-size commitments to polynomials and their applications. In *ASIACRYPT*, volume 6477 of *LNCS*, pages 177–194. Springer, 2010.
- [LAAZ11] Gregor Leander, Mohamed Ahmed Abdelraheem, Hoda AlKhazimi, and Erik Zenner. A Cryptanalysis of PRINTcipher: The Invariant Subspace Attack. In *CRYPTO 2011*, volume 6841 of *LNCS*, pages 206–221, 2011.
- [lib] C++ library for zkSNARK.
- [LMR⁺09] Mario Lamberger, Florian Mendel, Christian Rechberger, Vincent Rijmen, and Martin Schl affer. Rebound Distinguishers: Results on the Full Whirlpool Compression Function. In *ASIACRYPT 2009*, volume 5912 of *LNCS*, pages 126–143, 2009.
- [Mat93] Mitsuru Matsui. Linear Cryptanalysis Method for DES Cipher. In *EUROCRYPT 1993*, volume 765 of *LNCS*, pages 386–397, 1993.
- [MBKM19] Mary Maller, Sean Bowe, Markulf Kohlweiss, and Sarah Meiklejohn. Sonic: Zero-knowledge snarks from linear-size universal and updateable structured reference strings. Cryptology ePrint Archive, Report 2019/099, 2019. <https://eprint.iacr.org/2019/099>.
- [MRST09] Florian Mendel, Christian Rechberger, Martin Schl affer, and S oren S. Thomsen. The Rebound Attack: Cryptanalysis of Reduced Whirlpool and Gr ostl. In *FSE 2009*, volume 5665 of *LNCS*, pages 260–276, 2009.
- [MS78] F. J. MacWilliams and N. J. A. Sloane. *The Theory of Error-Correcting Codes*. North-holland Publishing Company, 1978.
- [NK92] Kaisa Nyberg and Lars R. Knudsen. Provable Security Against Differential Cryptanalysis. In *CRYPTO 1992*, volume 740 of *LNCS*, pages 566–574, 1992.
- [Nyb94] Kaisa Nyberg. Differentially uniform mappings for cryptography. In *EUROCRYPT 1993*, volume 765 of *LNCS*, pages 55–64, 1994.
- [PHGR13] Bryan Parno, Jon Howell, Craig Gentry, and Mariana Raykova. Pinocchio: Nearly practical verifiable computation. In *IEEE Symposium on Security and Privacy*, pages 238–252. IEEE Computer Society, 2013.
- [sov19] 2019. Mike Lodder, Sovrin’s principal cryptographer www.sovrin.org, private communication.
- [Sta18] StarkWare Industries Ltd. The complexity of STARK-friendly cryptographic primitives. Private communication, 2018.
- [W⁺14] Gavin Wood et al. Ethereum: A secure decentralised generalised transaction ledger. ethereum project yellow paper.(2014), 2014.
- [YMT97] A. M. Youssef, S. Mister, and S. E. Tavares. On the Design of Linear Transformations for Substitution Permutation Encryption Networks. In *School of Computer Science, Carleton University*, pages 40–48, 1997.

SUPPLEMENTARY MATERIAL

A Auxiliary Files

We provide the following files as supplementary material:

- `calc_round_numbers.py`
- `generate_parameters_grain.sage`
- `poseidonperm_x5_254_3.sage`
- `poseidonperm_x5_254_5.sage`
- `poseidonperm_x5_255_3.sage`
- `poseidonperm_x5_255_5.sage`
- `test_vectors.txt`
- `poseidonperm_x3_64_24_optimized.sage`

The first two scripts can be used to calculate round numbers and generate round constants and matrices, while the next five files contain the reference implementations and test vectors. Finally, we provide an optimized implementation of x^3 -POSEIDON $^\pi$, using the modifications shown in Supplementary Material B.

B Efficient Implementation

Like for LowMC, the fact that the non-linear layer is partial in R_P rounds can be used to reduce the size of the round constants required in each partial round. Referring to [DKP⁺19], we recall here an equivalent representation of an SPN with partial non-linear layers for an efficient implementation.

Round Constants. In the description of an SPN, it is possible to swap the order of the linear layer and the round constant addition as both operations are linear. The round constant then needs to be exchanged with an equivalent one. For round constant $c^{(i)}$, the equivalent one can be written as $\hat{c}^{(i)} = MC^{-1}(c^{(i)})$, where MC is the linear layer in the i -th round.

If one works with partial non-linear layers, it is possible to use this property to move parts of the original round constants from the last round all the way through the permutation to the beginning. In order to do this, it is sufficient to split the round constants in two parts, one that applies to the S-Box part of the non-linear layer and one that applies to the identity part of the non-linear layer. The constant part that only applies to the non-linear layer part can now move further up. Working in this way for all round constants, it is possible to finally end up with an equivalent representation in which round constants are only added to the output of the S-Boxes apart from one constant which is applied to the entire state after the first R_f full rounds.

Linear Layer. The situation for our design is simpler than for LowMC, since it is always possible to guarantee the existence of invertible sub-matrices. Hence, a similar trick proposed for LowMC in [DKP⁺19] works here as well for the matrix multiplication.

Focusing on the rounds with a single S-Box, let \mathcal{M} be the $t \times t$ MDS matrix of the linear layer:

$$\mathcal{M} = \left[\begin{array}{c|cccc} \mathcal{M}_{0,0} & \mathcal{M}_{0,1} & \mathcal{M}_{0,2} & \cdots & \mathcal{M}_{0,t-1} & \mathcal{M}_{0,t} \\ \mathcal{M}_{1,0} & & & & & \\ \mathcal{M}_{2,0} & & & & & \\ \vdots & & & & & \\ \mathcal{M}_{t-1,0} & & & & & \\ \mathcal{M}_{t,0} & & & & & \end{array} \right] = \left[\begin{array}{c|c} \mathcal{M}_{0,0} & v \\ w & \hat{\mathcal{M}} \end{array} \right],$$

where $\hat{\mathcal{M}}$ is a $(t-1) \times (t-1)$ MDS matrix (note that since \mathcal{M} is MDS, every submatrix of \mathcal{M} is also MDS), v is a $1 \times (t-1)$ matrix and w is a $(t-1) \times 1$ vector. By simple computation, the following equivalence holds:

$$\mathcal{M} = \underbrace{\left[\begin{array}{c|c} 1 & 0 \\ 0 & \hat{\mathcal{M}} \end{array} \right]}_{\mathcal{M}'} \times \underbrace{\left[\begin{array}{c|c} \mathcal{M}_{0,0} & v \\ \hat{w} & I \end{array} \right]}_{\mathcal{M}''}, \quad (7)$$

where

$$\hat{w} = \hat{\mathcal{M}}^{-1} \times w$$

and I is the $(t-1) \times (t-1)$ identity matrix. Note that both \mathcal{M}' and \mathcal{M}'' are two invertible matrices¹⁸.

As for the round constants discussed previously, it is possible to use Eq. (7) in order to swap the S-Box layer (formed by a single S-Box and $t-1$ identity functions) and the matrix multiplication with the matrix \mathcal{M}' . As a result, each linear part in the R_P partial rounds is defined only by a multiplication with a matrix of the form \mathcal{M}'' , which is a *sparse matrix*, since $(t-1)^2 - (t-1) = t^2 - 3t + 2$ coefficients of \mathcal{M}'' are equal to zero (moreover, $t-1$ coefficients of \mathcal{M}'' are equal to one). It follows that this optimized representation potentially greatly reduces the number of operations needed to compute the linear layer multiplication.

Computing Round Constants and Matrices. We provide an algorithmically optimized implementation of x^3 -POSEIDON $^\pi$ in the supplementary material (see Supplementary Material A). This file also contains functions to calculate the equivalent round constants and the equivalent matrices for the partial rounds. These functions can be used to generate these values for other instantiations of POSEIDON $^\pi$.

C Security Analysis: x^α -Poseidon $^\pi$

In the following, we assume that $x \mapsto x^\alpha$ is invertible (that is, $\gcd(\alpha, p-1) = 1$).

Differences to the Analysis for HadesMiMC. Although some of our results are similar to those provided for HADESMiMC in [GLR⁺20], we emphasize that our analysis is more tailored towards our specific construction, which is a family of *unkeyed* permutations used

¹⁸First of all, $\det(\mathcal{M}') = \det(\hat{\mathcal{M}}) \neq 0$ since $\hat{\mathcal{M}}$ is an MDS matrix, and so it is invertible. Secondly, $\det(\mathcal{M}) = \det(\mathcal{M}') \cdot \det(\mathcal{M}'')$. Since $\det(\mathcal{M}) \neq 0$ and $\det(\mathcal{M}') \neq 0$, it follows that $\det(\mathcal{M}'') \neq 0$.

in the context of a sponge function. This is why, for example, some MitM approaches do not work, as opposed to a (keyed) block cipher.

We make further distinctions also in the case of algebraic attacks. For example, we assume settings with $\chi < t$ input and output words, since this reflects the situation an attacker has to deal with in a sponge setting. This distinction affects our analysis regarding interpolation attacks, and most prominently our analysis regarding Gröbner basis attacks.

Moreover, we consider the case of a combination of certain algebraic attacks with invariant subspace trails, which is not directly applicable in the keyed setting.

C.1 Statistical Attacks

Since the results w.r.t. impossible differentials [BBS99], the multiple-of-8 property [GRR17], mixture differentials [Gra18] and integral distinguishers [DKR97] are equivalent to the ones proposed for HADES MiMC, we omit their analysis here and we refer to [GLR⁺20] for all details.

C.1.1 Differential Cryptanalysis

Here we focus only on the S-Boxes used in the paper, namely $x \mapsto x^3$ and $x \mapsto x^5$.

$S(x) = x^3$. Differential cryptanalysis [BS91, BS93] and its variations are the most widely used techniques to analyze symmetric-key primitives. The differential probability of any function over the finite field \mathbb{F}_p is defined as

$$\text{Prob}[\alpha \rightarrow \beta] := |\{x : f(x + \alpha) - f(x) = \beta\}|/p.$$

Since the cube function $f(x) = x^3$ is an almost perfect non-linear permutation (APN) [NK92, Nyb94], it has an optimal differential probability over a prime field. In other words, for this function the probability is bounded above by $2/|\mathbb{F}_p|$.

As largely done in the literature, we claim that POSEIDON ^{π} is secure against differential cryptanalysis if each differential has probability at most p^{-t} . Since it is in general hard to compute the probability of a differential, we assume that this fact is satisfied if each characteristic has probability at most $p^{-2 \cdot t}$. In order to compute the minimum number of rounds to guarantee this, we work only with the rounds with full S-Box layers. In other words, we limit ourselves to work with a “weaker” version of the permutation defined as

$$R^{R_f} \circ L \circ R^{R_f}(\cdot), \tag{8}$$

where

- L is an *invertible linear layer* (which is the “weakest” possible assumption),
- $R(\cdot) = M \circ \text{S-Box} \circ \text{ARK}(\cdot)$ where $\text{S-Box}(\cdot)$ is a full S-Box layer (remember that M is an MDS matrix).

We are going to show that this “weaker” permutation is secure against differential cryptanalysis for $R_F = 2R_f = 10$. As a result, it follows that also POSEIDON ^{π} (instantiated with R_F rounds with full S-Box layers) is secure against such an attack. Indeed, if the linear layer L (which we only assume to be invertible) is replaced by R_P rounds of POSEIDON ^{π} , its security cannot decrease. *The same strategy is exploited in the following in order to prove security against all attacks in this subsection.*

In order to prove the result just given, we need a lower bound for the number of active S-Boxes. Observe that the minimum number of active S-Boxes in the permutation

$$R^s \circ L \circ R^r(\cdot) \equiv SB \circ \underbrace{M \circ SB}_{s-1 \text{ times}} \circ \underbrace{L'}_{\equiv L \circ M(\cdot)} \circ SB \circ \underbrace{M \circ SB}_{r-1 \text{ times}}(\cdot)$$

(where $s, r \geq 1$, $R(\cdot)$ is a round with a full S-Box layer and where L' is an invertible linear layer) is at least¹⁹

$$\text{number active S-Boxes} \geq \underbrace{(\lfloor s/2 \rfloor + \lfloor r/2 \rfloor) \cdot (t+1)}_{\text{due to final/initial rounds}} + (s \bmod 2) + (r \bmod 2).$$

We emphasize that the (middle) linear layer $L'(\cdot) \equiv L \circ M(\cdot)$ plays *no* role in the computation of the previous number. Since at least $2 \cdot (t+1)$ S-Boxes are active in the 4 middle rounds of $R^2 \circ L \circ R^2(\cdot)$, and since the maximum differential probability of the cube S-Box is $DP_{max} = 2/p$, each characteristic has a probability of at most $(2/p)^{2 \cdot (t+1)}$, where

$$\begin{cases} p^{-2t} \cdot (p^{-2} \cdot 2^{2t+2}) \leq p^{-2t} & \text{if } \log_2 p \geq t+1, \\ p^{-1.25 \cdot t} \cdot (p^{-0.75 \cdot t - 2} \cdot 2^{2t+2}) < p^{-1.25 \cdot t} & \text{otherwise} \end{cases}$$

(where the second inequality holds since $\log_2(p) \cdot [0.75 \cdot t + 2] \geq 2t + 2$) for each $t \geq 2$ and $\log_2(p) \geq 3$. By doubling this number of rounds, we get that each characteristic has a probability of at most $p^{-2.5 \cdot t}$. Finally, 1 more round guarantees that no differential attack can be set up.

Security up to $2^M \leq p^t$. For completeness, we present the number of rounds necessary to provide security up to 2^M (that is, the data and the computational cost of the attacker are upper-bounded by 2^M). Using the same analysis as before, it turns out that

$$R_F = \begin{cases} 6 & \text{if } M \leq (t+1) \cdot (\log_2 p - 1) \\ 10 & \text{otherwise} \end{cases}$$

guarantees that no differential attack can be set up.

$\mathbf{S}(x) = x^5$. As before, x^5 -POSEIDON $^\pi$ instantiated by $\text{S-Box}(x) = x^5$ is secure against statistical attacks if and only if

$$R_F^{stat} \geq 6.$$

The main difference here is due to differential and linear attacks. In particular, since²⁰ $DP_{max}(\text{S-Box}(x) = x^5) = 4/p$, it follows that the minimum number of rounds necessary to guarantee security against linear and differential attacks is given by

$$R_F = \begin{cases} 6 & \text{if } 2t + 2 < N + \lceil \log_2(p) \rceil - M \\ 10 & \text{if } 2t + 2 \geq N + \lceil \log_2(p) \rceil - M \end{cases}$$

for a security level up to $2^M \leq p^t$ (that is, the data and the computational cost of the attacker are upper-bounded by 2^M).

¹⁹If $s = 2 \cdot s'$ is even, the minimum number of active S-Boxes over $R^s(\cdot)$ is $\lfloor s/2 \rfloor \cdot (t+1)$. Instead, if $s = 2 \cdot s' + 1$ is odd, the minimum number of active S-Boxes over $R^s(\cdot)$ is $\lfloor s/2 \rfloor \cdot (t+1) + 1$.

²⁰Note that $(x + \Delta_I)^5 - x^5 = \Delta_O$ is an equation of degree 4, hence there are at most 4 different solutions.

C.1.2 Linear Cryptanalysis

Similar to differential attacks, linear attacks [Mat93] pose no threat to the POSEIDON $^\pi$ family of permutations instantiated with the same number of rounds previously defined for classical differential cryptanalysis. In more detail, this follows from the fact that the maximum square correlation of the cube function is limited to $2/p$ (see [AÅBL12] for details). As a result, it offers the best possible resistance against linear cryptanalysis much like an APN function provides optimal resistance against differential cryptanalysis.

C.1.3 Truncated Differential

A variant of classical differential cryptanalysis is the truncated differential one [Knu94], in which the attacker can specify only part of the difference between pairs of texts.

We consider the “weaker” permutation described in Eq. (8) again. Focusing only on active/passive bytes (and not on the actual differences), there exist several differentials with probability 1 for a maximum of 1 round of POSEIDON $^\pi$, e.g.,

$$[\alpha, 0, \dots, 0]^T \xrightarrow{R(\cdot)} \mathcal{M} \times [\beta, 0, \dots, 0]^T,$$

where α, β denote non-zero differences. Due to the next S-Box layer, the linear relations given by $\mathcal{M} \times (\beta, 0, \dots, 0)^T$ are destroyed in the next round. As a result, no probability-one truncated differential covers more than a single round.

Since no linear relation survives the S-Box layer, it seems hard to set up a truncated differential for more than 2 rounds. As a result, it turns out that 4 rounds with full S-Box layers make POSEIDON $^\pi$ secure against this attack.

C.1.4 Rebound Attacks

The rebound attacks [LMR⁺09, MRST09] have much improved the best known attacks on many hash functions, especially for AES-based schemes. The goal of this attack is to find two (input, output) pairs (p^1, c^1) and (p^2, c^2) such that the two inputs satisfy a certain (truncated) input difference and the corresponding outputs satisfy a certain (truncated) output difference.

The rebound attack consists of two phases, called *inbound* and *outbound* phase. According to these phases, the internal permutation of the hash function is split into three sub-parts. Let f be the permutation, then we get $f = f_{fw} \circ f_{in} \circ f_{bw}$. The part of the inbound phase is placed in the middle of the permutation and the two parts of the outbound phase are placed next to the inbound part. In the outbound phase, two high-probability (truncated) differential trails are constructed, which are then connected in the inbound phase. Since the rebound attack is a differential attack, first the attacker needs to construct a “good” (truncated) differential trail. A good trail used for a rebound attack should have a high probability in the outbound phases and can have a rather low probability in the inbound phase. In the first phase, the attacker uses the knowledge of the key to find pairs of texts that satisfy the middle rounds of the truncated differential trail. In the second one, they propagate the solutions found in the first phase in the forward and in the backward directions, and check if at least one of them satisfies the entire differential trail.

The best rebound attack on AES proposed in [JNP13] covers 8 rounds. Here we claim that 6 rounds with full S-Box layers are sufficient to protect POSEIDON $^\pi$ from this attack. To support it, note that (1st) 1 round of POSEIDON $^\pi$ provides full diffusion while 2 rounds of AES are necessary to provide it and (2nd) the best truncated differential covers 1 round

of POSEIDON $^\pi$ vs 3 rounds of AES²¹. Since the best results on AES in the literature cover at most 8 rounds, due to the similarity between AES and POSEIDON $^\pi$ and due to the previous observations, we argue that it is not possible to mount a rebound attack on more than 5 rounds with full S-Box layers of POSEIDON $^\pi$. Hence, 6 rounds of POSEIDON $^\pi$ with full S-Box layers are sufficient to guarantee security against this attack.

C.1.5 Invariant Subspace Attack

The invariant subspace attack [LAZ11] makes use of affine subspaces that are invariant under the round function.

Definition 2 (Invariant Subspace Trail). Let K_{weak} be a set of keys and $k \in K_{\text{weak}}$, with $k \equiv (k^0, k^1, \dots, k^r)$ where k^j is the j -th round key. For each $k \in K_{\text{weak}}$, the subspace U generates an *invariant subspace trail* of length r for the function $F_k(\cdot) \equiv F(\cdot) \oplus k$ if for each $i = 1, \dots, r$ there exists a non-empty set $A_i \subseteq U^c$ for which the following property holds:

$$\forall a_i \in A_i : \exists a_{i+1} \in A_{i+1} \text{ s.t. } F_k^i(U \oplus a_i) = U \oplus a_{i+1}.$$

All keys in the set K_{weak} are *weak keys*.

Definition 3 (Subspace Trails). Let $(U_1, U_2, \dots, U_{r+1})$ denote a set of $r + 1$ subspaces with $\dim(U_i) \leq \dim(U_{i+1})$. If for each $i = 1, \dots, r$ and for each a_i , there exists (unique) $a_{i+1} \in U_{i+1}^c$ such that

$$F(U_i \oplus a_i) \subseteq U_{i+1} \oplus a_{i+1},$$

then $(U_1, U_2, \dots, U_{r+1})$ is a *subspace trail* of length r for the function F . If all the previous relations hold with equality, the trail is called a *constant-dimensional subspace trail*.

Here we consider separately the two cases, namely rounds with full S-Box layers and rounds with partial S-Box layers.

- In the first case, as the round constant addition translates this invariant subspace [BCLR17], random round constants provide a good protection against these attacks.
- In the second case, it is always possible to construct a subspace trail for (at least) $t - 1$ rounds without activating any S-Box and for any choice of the linear layer. Due to the assumption made in Section 2.3, $t - 1$ is also an upper bound for the number of rounds for which such a trail can be constructed.

In other words, due to the choice of the matrix that defines the linear layer (described in details in Section 2.3), it is not possible to cover more than $t - 1$ rounds with a subspace trail (or equivalently, a truncated differential with prob. 1) with active/inactive S-Boxes. We refer to [GRS20] for more details about the possibility to set up infinitely long subspace trails for partial SPN schemes (equivalently, for the rounds with partial S-Box layers in HADES-like schemes).

C.2 Algebraic Attacks

First, we introduce a simple lemma which follows from the iterative structure of the HADES permutation.

Lemma 1. *The algebraic degree $D_\alpha(r)$ of r -round POSEIDON $^\pi$ with the S-Box x^α as a function of input variables is at most α^r , and we expect it to reach this degree no matter if partial or full rounds are used.*

²¹The best truncated differential distinguisher with prob. 1 covers 2 rounds of AES.

C.2.1 Interpolation Attack

One of the most powerful attacks is the interpolation attack, introduced by Jakobsen and Knudsen [JK97] in 1997. In the case of a keyed function $E_k : \mathbb{F} \rightarrow \mathbb{F}$, the strategy of the attack is to construct a polynomial representation of the function without knowledge of the secret key. If an adversary can construct such a polynomial then they can compute any output without knowing the key, thus enabling forgeries (for MAC settings) and other attacks. The interpolation polynomial $P(x)$ representing $E_k(x)$ can be constructed using e.g. the Vandermonde matrix – cost approximately of $\mathcal{O}(t^2)$ – or the Lagrange’s theorem – cost approximately of $\mathcal{O}(t \cdot \log t)$, where x is the indeterminate corresponding to the input.

Such an attack can be opportunely modified for the case of an unkeyed permutation $E(\cdot)$. In such a case, assume it is possible to construct the interpolation polynomial without using the full code book. In this case, such a polynomial can be exploited to set up a forgery attack on the permutation E , which, in general, is not possible for a (pseudo-)random permutation.

In more detail, each output word of an SPN permutation can be represented as a multivariate polynomial where the variables are the inputs to each S-Box. Consider a permutation input where χ input words are unknown to us, and the other $t - \chi$ words are known. A (rough) estimation of the number of monomials of the interpolation polynomial (and so of the complexity of the attack) is given by

$$(D_\alpha(r) + 1)^\chi.$$

As a result, by requiring the number of monomials to be close to the number of possible input values p^χ , the number of rounds must be at least $r \simeq \log_\alpha(p)$.

However, just reaching the full degree is not sufficient to prevent the interpolation attack. First, the polynomial should be dense to guarantee that most monomials occur in it. As shown in details in [GLR⁺20], one needs at least $1 + \lceil \log_\alpha(p) \rceil + \lceil \log_\alpha(t) \rceil$ rounds in order to guarantee that the interpolation polynomial E_k is *dense*.

Since $\text{S-Box}^{-1}(x) = x^{1/\alpha}$ has a higher degree than $\text{S-Box}(x) = x^\alpha$ (e.g., $x^{1/3} = x^{(2p-1)/3}$), we do not expect the attack to perform better when considering the backward direction instead of the forward one. In particular, we expect that 2 rounds are sufficient to reach the maximum degree in the backward direction. Moreover, we emphasize that we do not claim security w.r.t. the MitM variant of the interpolation attack. Indeed, since the permutation is used in a sponge construction, the attacker can only see part of the output of the permutation, which naturally prevents the MitM variant attack.

Secondly, we consider the algebraic degree not at round r but at round $r - 1$ to account for the partial S-Box case where the degree increase is delayed for $t - 1$ words by 1 round. As a result, the total number of rounds R must satisfy ²²

$$R \geq 1 + \lceil \log_\alpha(p) \rceil + \log_\alpha(t)$$

to thwart the interpolation attack.

Security up to $2^M \leq p^t$. For completeness, we present the number of rounds necessary to provide security up to 2^M (that is, the data and the computational cost of the attacker are upper-bounded by 2^M).

²²We emphasize that in this analysis we do not take into account the cost to construct the interpolation polynomial, which is (in general) non-negligible.

Using the argumentation given before, the number of rounds must satisfy

$$\left(\alpha^{r-\log_\alpha(t)-1} + 1\right)^\chi \approx 2^{\min\{M, \log_2(p)\chi\}},$$

that is, $r \geq 1 + \log_\alpha(t) + \min\{M, \log_2(p)\chi\}/\chi \cdot \log_\alpha(2)$. The maximum number of attacked rounds is achieved for $\chi = 1$. As a result, we have $R_P + R_F \geq (1 + \lceil \log_\alpha(2) \cdot \min\{M, \log_2(p)\} \rceil) + \log_\alpha(t)$.

C.2.2 Gröbner Basis Attack

We consider the Gröbner basis attack in the same setting: Some permutation inputs are unknown and the rest are known to the attacker. Given some words of the permutation output, they have to find the unknowns.

A Gröbner basis attack usually consists of three steps, namely,

- (1) computing the Gröbner basis in *degrevlex* order,
- (2) converting the result to the *lexicographic* order, and
- (3) factorizing the univariate polynomial, and back-substituting its roots.

As largely done in the literature, we assume that the security of permutations against Gröbner basis attacks follows from the infeasible complexity of computing the Gröbner basis in *degrevlex* order. For generic systems, the complexity of this step for a system of \mathcal{N} polynomials f_i in \mathcal{V} variables is $\mathcal{O}\left(\binom{\mathcal{V}+D_{\text{reg}}}{D_{\text{reg}}}\omega\right)$ operations over the base field \mathbb{F} [CLO97], where D_{reg} is the *degree of regularity* and $2 \leq \omega < 3$ is the linear algebra constant (note that the memory requirement of these algorithms is of the same order as the running time). The degree of regularity depends on the degrees of the polynomials d and the number of polynomials \mathcal{N} . When $\mathcal{V} = \mathcal{N}$, we have the simple closed-form approximation

$$D_{\text{reg}} := 1 + \sum_{i=0}^{\mathcal{N}-1} (d_i - 1), \quad (9)$$

where d_i is the degree of the i -th polynomial f_i in the polynomial system we are trying to solve (see [BFSY05] for details). In the over-determined case, i.e., $\mathcal{V} < \mathcal{N}$, the degree of regularity can be estimated by developing the Hilbert series of an ideal generated by generic polynomials $\langle f_0, \dots, f_{\mathcal{N}-1} \rangle$ of degrees d_i (under the assumption that the polynomials behave like generic systems). Closed-form formulas for D_{reg} are known for some special cases, but not in general.

Full-Permutation Equations. In the first case, we derive equations (one by word) for the entire r -round permutation. We consider the case in which the number χ of unknown input variables equals the number of known output variables. Then we get χ equations of degree $D_\alpha(r) = \alpha^r$ in χ variables. We now estimate the complexity of solving algorithms based on different values of χ .

If $\chi = 1$ (i.e., we have one equation in one variable), the resulting system of equations is already a Gröbner basis consisting of a single equation of degree α^r . This means that we can omit the step of computing a Gröbner basis and we focus on the final step, which consists of finding the roots of the corresponding polynomial. This can be done by using e.g. a fast Berlekamp factorization algorithm [Gen07], where the complexity of factoring a

degree- D polynomial over \mathbb{F}_p for a prime number p is an element in $\mathcal{O}(D^\omega)$ for $2 \leq \omega \leq 3$. Hence, we want that

$$\log_2(\alpha^{\omega \cdot r}) \geq \log_2(\alpha^{2r}) \geq \min\{M, \log_2(p)\},$$

which implies

$$r \geq \left\lceil \frac{\min\{M, \log_2(p)\}}{2 \log_2(\alpha)} \right\rceil,$$

where $r = R_F + R_P$.

If $\chi \geq 2$, we consider multiple equations in multiple variables. Computing a Gröbner basis for the resulting system of equations can be estimated by

$$\begin{pmatrix} \mathcal{V} + D_{\text{reg}} \\ D_{\text{reg}} \end{pmatrix}^2 \approx \binom{\chi \alpha^r}{\chi}^2 \leq \left[\frac{(\chi \cdot \alpha^r)^\chi}{\chi!} \right]^2 \leq \alpha^{2r\chi},$$

since $\frac{(\chi \alpha^r)!}{[\chi(\alpha^r - 1)]!} = (\chi \alpha^r) \cdot \dots \cdot [\chi(\alpha^r - 1) + 1] \leq (\chi \alpha^r)^\chi$ and since $x! \leq x^x$ for each $x \geq 1$. Focusing on the complexity of computing a Gröbner basis and on an M -bit security level, the number of rounds to be attacked can be estimated as

$$\alpha^{2r\chi} \leq 2^{\min\{M, \log_2(p)\chi\}}.$$

Since this value is maximized for $\chi = 2$ (in this setting), we conjecture that security is provided by choosing

$$r = R_F + R_P \geq \min \left\{ \left\lceil \frac{M}{4 \log_2(\alpha)} \right\rceil, \left\lceil \frac{\log_2(p)}{2 \log_2 \alpha} \right\rceil \right\},$$

where $\frac{1}{4 \log_2(3)} \approx 0.16$ and $\frac{1}{4 \log_2(5)} \approx 0.11$ for the cases we focus on in this paper.

By combining these two cases ($\chi = 1$ and $\chi \geq 2$), we finally get

$$r = R_F + R_P \geq \min \left\{ \left\lceil \frac{M}{2 \log_2(\alpha)} \right\rceil, \left\lceil \frac{\log_2(p)}{2 \log_2 \alpha} \right\rceil \right\}.$$

Exploit the Subspace $\mathcal{S}^{(t)}$ to Speed-Up the Attack. As shown in [BCD⁺20], the previous attack can be improved by exploiting the existence of the subspace $\mathcal{S}^{(r)}$ defined as in Eq. (1). In particular, given a text in a coset of $\mathcal{S}^{(r)}$, the output of this text after r rounds with partial S-Box layers is just the result of an affine map applied to the input. As explained in details in [BCD⁺20], this fact can be exploited in order to replace some non-linear equations of the system that we are trying to solve with linear equations. In particular, by starting in the “middle” with texts in a coset of $\mathcal{S}^{(r)}$, one faces the following scenario

$$\begin{array}{ccc} \text{input} \xleftarrow{R_F^{-R_f}(\cdot)} \xleftarrow{R_P^{-\frac{R_P-r}{2}}(\cdot)} & \text{text in coset of } \mathcal{S}^{(r)} & \xrightarrow{R_P^r(\cdot)} \\ & & \text{“linear” equations} \end{array} \quad (10)$$

$$\text{text in coset of } \mathcal{M}'(\mathcal{S}^{(r)}) \xrightarrow{R_P^{\frac{R_P-r}{2}}(\cdot)} \xrightarrow{R_F^{R_f}(\cdot)} \text{output},$$

where \mathcal{M}' is a $t \times t$ matrix that depends on the details of the cipher and on the subspace $\mathcal{S}^{(r)}$ (and where $\mathcal{M}'(\mathcal{S}^{(r)}) = \{\mathcal{M}' \times x \mid x \in \mathcal{S}^{(r)}\}$).

The number of rounds r described by linear equations is related to the dimension of the subspace $\mathcal{S}^{(r)}$. This fact has an impact on the chance to find a solution of the system

of equations we are trying to solve (indeed, the number of variables of such a system of equations is related to the dimension of the subspace). In any case, due to the assumption made in Section 2.3 on the matrix that defines the linear layer, this strategy does not allow the attacker to cover more than $t - 1$ rounds with partial S-Box layers.

The scenario Eq. (10) can be described by

- (*first part*) t equations of degree $\alpha^{R_f + \frac{R_P - r}{2}}$,
- (*middle part*) d linear equations and $t - d$ non-linear equations of degree α^r , and
- (*final part*) t equations of degree $\alpha^{R_f + \frac{R_P - r}{2}}$.

Let us focus on the middle part (r rounds). Let $\mathcal{S} = \langle s_1, \dots, s_d \rangle$ be the subspace defined before of dimension d , and let $\mathcal{S}^c = \langle w_1, \dots, w_{t-d} \rangle$ where \mathcal{S}^c denotes the complementary space of \mathcal{S} (obviously, $\mathbb{F}^t = \langle s_1, \dots, s_d, w_1, \dots, w_{t-d} \rangle$). Let x and y be resp. the input and the corresponding output of such a middle part, where $y = F(x)$ for a specific function $F(\cdot)$ of degree α^r . Given x in a coset of \mathcal{S} and by definition of \mathcal{S} , the output y is in a subspace of $\mathcal{M}'(\mathcal{S}^{(r)})$ for a particular matrix \mathcal{M}' . Hence, the middle part can be described by a system of d linear equations of the form

$$s_i^T \cdot y = s_i^T \cdot (\mathcal{M}' \times x) \quad \forall i \in \{1, \dots, d\}$$

and $t - d$ non-linear equations of the form

$$w_i^T \cdot y = w_i^T \cdot F(x) \quad \forall i \in \{1, \dots, t - d\},$$

where $a^T \cdot b$ denotes the scalar product between $a, b \in \mathbb{F}^t$.

Working as before, this means that the attacker has to face a system of $t + (t - d) + \chi$ equations in $t + (t - d) + \chi$ variables. By adding constraints on the previous $t - d$ non-linear equations (which corresponds on “fixing” the coset of the subspace \mathcal{S}), these number of equations/variables can be reduced to $t + \chi$.

By using the previous computation and since $D_{\text{reg}} = 1 + (t + \chi) \cdot (\alpha^{R_f + \frac{R_P - r}{2}} - 1) \approx (t + \chi) \cdot \alpha^{\frac{R_f + R_P - r}{2}}$, the cost can be approximated by

$$\begin{aligned} \binom{(t + \chi) \cdot \alpha^{\frac{R_f + R_P - r}{2}}}{t + \chi}^\omega &\leq \left[\frac{\left((t + \chi) \cdot \alpha^{\frac{R_f + R_P - r}{2}} \right)^{t + \chi}}{(t + \chi)!} \right]^2 \leq \\ &\leq \alpha^{(R_f + R_P - r) \cdot (t + \chi)} \end{aligned}$$

assuming $\omega \approx 2$. If we target a security level of M bits, the number of rounds to be attacked is calculated as

$$\begin{aligned} \alpha^{(R_f + R_P - r) \cdot (t + \chi)} &\leq 2^{\min\{M, \log_2(p)\chi\}}, \\ \implies R_f + R_P &\geq r + \min \left\{ \frac{M \log_\alpha(2)}{t + \chi}, \frac{\log_2(p)\chi \log_\alpha(2)}{t + \chi} \right\}. \end{aligned}$$

Since the maximum number of rounds to be attacked is achieved resp. for $\chi = 1$ in the first term and for $\chi = t$ in the second one, the security is provided by choosing

$$R_f + R_P \geq t - 1 + \min \left\{ \frac{\log_\alpha(2) \cdot M}{t + 1}, \frac{\log_\alpha(2) \cdot \log_2(p)}{2} \right\},$$

where remember that $r \leq t - 1$.

Working at Round Level: Equations for Each S-Box. Here we consider equations of degree α for each S-Box, which relate its inputs and outputs. Given χ unknown permutation inputs and χ known outputs, we get $(t-1)R_F + R_P + \chi$ unknown S-Boxes, and for each we use 1 variable (for its input). In total, we get $(t-1)R_F + R_P$ equations for the S-Box inputs in all rounds, and χ equations for the last-round outputs. Denoting $q = (t-1)R_F + R_P + \chi$, the degree of regularity is estimated by $D_{\text{reg}} \approx 1 + 2q$. The attack complexity can then be estimated by

$$\begin{aligned} \binom{\mathcal{V} + D_{\text{reg}}}{D_{\text{reg}}} &\approx \binom{\alpha \cdot q}{q} = \left(\frac{(\alpha \cdot q)!}{q! \cdot [(\alpha - 1) \cdot q]!} \right)^2 \approx \\ &\approx \left(\frac{\alpha^\alpha}{(\alpha - 1)^{\alpha-1}} \right)^{2q} \cdot \left(\frac{2\pi \cdot (\alpha - 1) \cdot q}{\alpha} \right)^{-1}, \end{aligned}$$

where we used Stirling's approximation for the factorials. We introduce two constants:

$$\mathcal{C} := 2 \cdot \log_2 \left(\frac{\alpha^\alpha}{(\alpha - 1)^{\alpha-1}} \right) \quad \text{and} \quad \mathcal{C}' := \log_2 \left(\frac{2\pi \cdot (\alpha - 1) \cdot q}{\alpha} \right).$$

Since in our cases $q \geq 8$ (where q denotes the number of variables used in the attack, which is the same as the number of S-Boxes in our construction), in the following we use $\mathcal{C}' = 3 + \log_2 \left(\frac{\pi \cdot (\alpha - 1)}{\alpha} \right)$. It follows that the attack complexity can be estimated by

$$2^{\mathcal{C} \cdot q - \mathcal{C}'} \leq 2^{\mathcal{C} \cdot q}.$$

More concretely,

$$\mathcal{C}(\alpha = 3) = 5.5 \quad \text{and} \quad \mathcal{C}(\alpha = 5) = 7.2.$$

If we target a security level of M bits, the number of rounds to be attacked is calculated as

$$\begin{aligned} 2^{\mathcal{C} \cdot ((t-1)R_F + R_P + \chi)} &\leq 2^{\min\{M, \log_2(p)\chi\}} \\ \implies \mathcal{C} \cdot ((t-1)R_F + R_P + \chi) &\leq \min\{M, \log_2(p)\chi\} \\ \implies (t-1)R_F + R_P + \chi &\leq \mathcal{C}^{-1} \cdot \min\{M, \log_2(p)\chi\}. \end{aligned}$$

Since the maximum number of rounds to be attacked is achieved for $\chi = 1$, we conjecture security by choosing $(t-1)R_F + R_P$ greater than

$$\left[2 \cdot \log_2 \left(\frac{\alpha^\alpha}{(\alpha - 1)^{\alpha-1}} \right) \right]^{-1} \cdot \min\{M, \log_2(p)\} + t - 2,$$

where $t-1$ rounds are added to prevent the version of the attack just described that makes use of the subspace \mathcal{S}^i in order to skip rounds without activating any S-Box.

Conclusion. Combining the strategies together, we get the following conditions:

$$\begin{aligned} R_F + R_P &\geq \log_\alpha(2) \cdot \min \left\{ \frac{M}{3}, \frac{\log_2(p)}{2} \right\}, \\ R_F + R_P &\geq t - 1 + \min \left\{ \frac{\log_\alpha(2) \cdot M}{t + 1}, \frac{\log_\alpha(2) \cdot \log_2(p)}{2} \right\} \\ (t-1)R_F + R_P &\geq \\ &\left[2 \cdot \log_2 \left(\frac{\alpha^\alpha}{(\alpha - 1)^{\alpha-1}} \right) \right]^{-1} \cdot \min\{M, \log_2(p)\} + t - 2, \end{aligned} \tag{11}$$

where $n \approx \log_2 p$. Note that if the first inequality is satisfied, then the third one is satisfied as well: hence, it is sufficient that the first two conditions are satisfied.

D Security Analysis: x^{-1} -Poseidon $^\pi$

Here we propose the security analysis of POSEIDON $^\pi$ instantiated with S-Box(x) = x^{-1} in GF(p). In particular, we focus only on the attacks that *depend* on the details of the S-Box, like differential/linear attacks and the algebraic attacks.

D.1 Statistical Attacks

Differential Cryptanalysis. As before, x^{-1} -POSEIDON $^\pi$ instantiated by S-Box(x) = x^5 is secure against statistical attacks if and only if

$$R_F^{stat} \geq 6.$$

The main difference here is due to differential and linear attacks. In particular, since $DP_{max}(S\text{-Box}(x) = x^{-1}) = 4/p$, it follows that the minimum number of rounds necessary to guarantee security against linear and differential attacks is given by

$$R_F = \begin{cases} 6 & \text{if } 2t + 2 < N + \lceil \log_2(p) \rceil - M \\ 10 & \text{if } 2t + 2 \geq N + \lceil \log_2(p) \rceil - M \end{cases}$$

for a security level up to $2^M \leq p^t$ (that is, in the case in which the data and the computational costs of the attacker are upper-bounded by 2^M).

Rebound Attacks. Due to the argumentation for the security of x^α -POSEIDON $^\pi$ against the rebound attack, we argue that 6 rounds provide security also for x^{-1} -POSEIDON $^\pi$.

D.2 Algebraic Attacks

D.2.1 Interpolation Attack

As we have already seen, in an interpolation attack [JK97], the goal is to determine the polynomial representation of a state word. Since the inverse function has a high degree, one may think that the interpolation attack can cover only few rounds in this case. However, exploiting the original idea proposed by Jakobsen and Knudsen in [JK97], it is possible to show that

- for a full S-Box layer, the S-Box $f(x) = x^{-1}$ has the same behavior as a function of algebraic degree t (i.e., the number of words)²³ “from the point of view” of the interpolation attack, and
- for a partial S-Box layer (with a single S-Box), the S-Box $f(x) = x^{-1}$ has the same behavior as a function of algebraic degree 2 “from the point of view” of the interpolation attack.

Note that the two previous cases lead to two completely different results, while we emphasize that the two previous cases (full or partial S-Box layers) are equivalent for the cube S-Box. It follows that the choice of using partial or full S-Box layers in order to protect from algebraic attacks also depends on the details of the S-Box.

²³More precisely, the degree of S-Box(x) = $x^{-1} \equiv x^{p-2}$ “from the point of view of the interpolation attack” is $\min\{t, p-2\}$, where t is due to the fraction representation and $p-2$ is due to the “normal” representation. Since $2t \leq p+1$ in order to guarantee that a $t \times t$ MDS matrix with coefficients in \mathbb{F}_p exists, it follows that $\min\{t, p-2\} = t$.

Full S-Box Layer. First, consider $t = 1$. In this case, the function can be written as

$$f(x) = \frac{x + A}{B \cdot x + C}$$

for *any* number of rounds and for some constants A, B, C . This means that 4 texts are sufficient to break the permutation.

Consider the case $t = 2$. Let $f_i^r(\cdot) \equiv \frac{Nf_i^r(\cdot)}{Df_i^r(\cdot)}$ (for $i = 0, 1$) be the interpolation polynomial at round r of the i -th word. By simple computation, the i -th word of the function at round $r + 1$ (assuming a full S-Box layer) for $i = 0, 1$ can be written as

$$\begin{aligned} f_i^{r+1}(x \equiv [x_0, x_1]) &= \\ &= \frac{A}{f_0^r(x \equiv [x_0, x_1]) + k_0} + \frac{B}{f_1^r(x \equiv [x_0, x_1]) + k_1} = \\ &= \frac{A \cdot Df_0^r(x)}{Nf_0^r(x) + k_0 \cdot Df_0^r(x)} + \frac{B \cdot Df_1^r(x)}{Nf_1^r(x) + k_1 \cdot Df_1^r(x)} = \\ &= \frac{A \cdot [Nf_1^r(x) + k_1 \cdot Df_1^r(x)] \times Df_0^r(x)}{[Nf_0^r(x) + k_0 \cdot Df_0^r(x)] \times [Nf_1^r(x) + k_1 \cdot Df_1^r(x)]} + \\ &+ \frac{B \cdot [Nf_0^r(x) + k_0 \cdot Df_0^r(x)] \times Df_1^r(x)}{[Nf_0^r(x) + k_0 \cdot Df_0^r(x)] \times [Nf_1^r(x) + k_1 \cdot Df_1^r(x)]} = \\ &= \frac{Nf_i^{r+1}(x \equiv [x_0, x_1])}{Df_i^{r+1}(x \equiv [x_0, x_1])} \end{aligned}$$

for some constants A, B . It follows that the degree of the function increases at most by a factor of 2 (where the degree after the first round is 1). As a result, the number of unknown coefficients after r rounds is at most $2 \cdot (2^{r-1} + 1)^2$, where the degree of the numerator (and so the number of unknown coefficients) is always less than or equal to the degree of the denominator.

Since the number of unknown coefficients after r rounds for t words is approximately $2 \cdot (t^{r-1} + 1)^t$, the permutation can be considered secure if $2 \cdot (t^{r-1} + 1)^t \simeq 2^N$, that is, $t^{r-1} \simeq p$, which implies

$$r \geq \log_t(2) \cdot \log_2(p) + 1.$$

As a result, the total number of rounds (with full S-Box layers) must be

$$R_F \geq 2 + \frac{\log_2(p)}{\log_2(t)}.$$

Partial S-Box Layer. Referring to the expression of f_i^r given before, it is possible to note that all denominators at round r (for any r) are in general equal, while all numerators are in general different, that is,

$$\forall i, j \in [0, 1, \dots, t-1] : \quad Df_i^r = Df_j^r.$$

This observation seems to have no effect on the complexity of the previous attack. Indeed, since the S-Boxes are applied to each word and since the numerators are different, it turns out that the denominators of S-Box(f^r) (which correspond to the numerators of f^r) are all different.

However, this is not the case when working with a partial non-linear layer, e.g., with a non-linear layer composed of a single S-Box. Consider first the case $t = 2$ assuming the

S-Box is applied only to the first word (we use the same notation as before):

$$\begin{aligned}
f_i^{r+1}(x \equiv [x_0, x_1]) &= \\
&= \frac{A}{f_0^r(x \equiv [x_0, x_1]) + k_0} + B \cdot [f_1^r(x) + k_1] = \\
&= \frac{A \cdot Df_0^r(x)}{Nf_0^r(x) + k_0 \cdot Df_0^r(x)} + \frac{B \cdot [Nf_1^r(x) + k_1 \cdot Df_1^r(x)]}{Df_1^r(x)} = \\
&= \frac{A \cdot Df_0^r(x) \times Df_1^r(x)}{[Nf_0^r(x) + k_0 \cdot Df_0^r(x)] \times Df_1^r(x)} + \\
&+ \frac{B \cdot [Nf_1^r(x) + k_1 \cdot Df_1^r(x)]}{Df_1^r(x)} = \frac{Nf_i^{r+1}(x \equiv [x_0, x_1])}{Df_i^{r+1}(x \equiv [x_0, x_1])}.
\end{aligned}$$

In this case, there is no difference w.r.t. the previous case.

Consider now the case $t \geq 3$. By previous observation, it follows that $Df_i^r(x) = Df_j^r(x)$ for each $i, j \geq 1$, which implies that

$$Df_i^{r+1} = [Nf_0^r(x) + k_0 \cdot Df_0^r(x)] \times Df_1^r(x)$$

also for the case $t \geq 3$. This fact has a huge impact on the number of monomials of the corresponding polynomial at round r . Indeed, the number of unknown coefficients after r rounds for t words is approximately $2 \cdot (2^{r-1} + 1)^t$, which is much smaller than $2 \cdot (t^{r-1} + 1)^t$ for large t . The permutation can be considered secure if $2 \cdot (2^{r-1} + 1)^t \simeq 2^N$, that is, $2^{r-1} \simeq p$, which implies $r \geq \log_2(p) + 1$. As a result, the total number of rounds (with full S-Box layers) must be

$$R \equiv R_P + R_F \geq \log_2(p) + 1.$$

Full & Partial S-Box Layers: Security up to $2^M \leq 2^N$. The previous result can be improved. Since at least $R_F \geq 6$ rounds have a full S-Box layer, it follows that the number of unknown coefficients after $R = R_P + R_F$ rounds for t words is approximately

$$2 \cdot (2^{R_P} \cdot t^{R_F-1} + 1)^t \equiv 2 \cdot (2^{R_P+(R_F-1) \cdot \log_2(t)} + 1)^t.$$

The permutation can be considered secure if $2 \cdot (2^{R_P+(R_F-1) \cdot \log_2(t)} + 1)^t \simeq 2^N$, that is,

$$R_P + \log_2(t) \cdot R_F \geq 2 + \log_2(p) + \log_2(t).$$

If one aims to provide a security of up to M bits and using the argumentation given before, the number of rounds must satisfy

$$R_P + R_F \cdot \log_2(t) \geq 2 + \log_2(t) + \min\{M, \log_2(p)\}.$$

D.3 Gröbner Basis Attack

We use the same setting with χ unknown inputs and χ known outputs.

Case: Input-Output. After $R_F + R_P \geq 1$ rounds and using the “fraction representation” just proposed for the interpolation attack, the minimum degree of a variable in the output polynomials is $d = t^{R_F-1} 2^{R_P}$, using the equivalence

$$f(x) \equiv \frac{Nf(x)}{Df(x)} = C \quad \text{if and only if} \quad Nf(x) = C \cdot Df(x).$$

Therefore we get χ equations of degree d of χ variables, so the degree of regularity is $D_{\text{reg}} = 1 + \chi(d - 1)$. If we target the security level of M bits, the number of rounds to be attacked is calculated as

$$(t^{R_F-1} 2^{R_P})^{2\chi} \leq 2^{\min\{M, \log_2(p)\chi\}} \implies \\ R_F \log_2(t) - \log_2(t) + R_P \leq \frac{\min\{M, \log_2(p)\chi\}}{2\chi}.$$

Since the maximum number of attacked rounds is achieved for $\chi = 1$, the security is provided by choosing

$$R_F \log_2(t) + R_P \geq \log_2(t) + 0.5 \cdot \min\{M, \log_2(p)\}.$$

Similarly, in order to prevent the modified version of the attack described in Supplementary Material C.2.2, the number of rounds must satisfy

$$R_F \log_2(t) + R_P \geq t - 1 + \log_2(t) + \min\left\{\frac{M}{t+1}, 0.5 \cdot \log_2(p)\right\}.$$

Equations for Each S-Box. Here we consider equations of degree 2 for each S-Box, which relate its inputs and outputs. Given χ unknown permutation inputs and χ known outputs, we get $(t-1)R_F + R_P + \chi$ unknown S-Boxes, and for each we use 1 variable for its input. In total we get $(t-1)R_F + R_P$ equations for the S-Box inputs in all rounds, and χ equations for the last round outputs. Denoting $q = (t-1)R_F + R_P + \chi$, the degree of regularity is estimated by $D_{\text{reg}} = 1 + q$. The attack complexity is lower-bounded by

$$\binom{\mathcal{V} + D_{\text{reg}}}{D_{\text{reg}}}^2 \approx \binom{2q}{q}^2 \approx \frac{2^{4q}}{\pi q} \approx 2^{3.9q},$$

where we use the same approximations as in Supplementary Material C.2. If we target a security level of M bits, the number of rounds to be attacked is calculated as

$$2^{3.9((t-1)R_F + R_P + \chi)} \leq 2^{\min\{M, \log_2(p)\chi\}}.$$

Since the maximum number of rounds to be attacked is achieved for $\chi = 1$, the security is provided by chosen

$$(t-1)R_F + R_P \geq \frac{\min\{M, \log_2(p)\}}{3.9} + t - 2 \approx \\ \approx 0.25 \cdot \min\{M, \log_2(p)\} + t - 2,$$

where we add $t-1$ rounds in order to prevent the improved version of the attack described in Supplementary Material C.2.2.

Conclusion. Combining the two strategies together, the security is achieved if the following conditions are satisfied:

$$R_F \log_2(t) + R_P \geq \log_2(t) + 0.5 \cdot \min\{M, \log_2(p)\} + t - 1, \\ R_F \log_2(t) + R_P \geq t - 1 + \log_2(t) + \min\left\{\frac{M}{t+1}, 0.5 \cdot \log_2(p)\right\}, \\ (t-1)R_F + R_P \geq 0.25 \cdot \min\{M, \log_2(p)\} + t - 2.$$

As before, note that if the first inequality is satisfied, the third one is satisfied as well.

E Selecting Number of Rounds in General Case

The design goal is to offer a family of hash functions which minimize the R1CS costs. It turns out that for the fixed S-Box function the minimum costs are delivered by a primitive with the smallest number of S-Boxes, though the field size also plays a role. For each combination (security level M , prime/binary field type, S-Box size, S-Box function) we minimize the number of S-Boxes taking into account Eq. (3) and Eq. (5).

Minimizing the Number of S-Boxes. In our design strategy, we always exploit the wide trail strategy in order to provide security against statistical attacks. In other words, for this class of attacks, we only work with rounds with full S-Box layers. All our instances are secure against statistical attacks if

$$R_F^{stat} \geq \begin{cases} 6 & \text{if } M \leq (\log_2 p - \mathcal{C}) \cdot (t + 1) \\ 10 & \text{otherwise} \end{cases}$$

where $\mathcal{C} = 2$ for $S(x) = 1/x$ and $\mathcal{C} = \log_2(\alpha - 1)$ for $S(x) = x^\alpha$. In order to minimize the number of S-Boxes for given n and t , the goal is to find the best ratio between R_F and R_P that minimizes

$$\text{number of S-Boxes} = t \cdot R_F + R_P, \quad (12)$$

where $t \geq 2$ and where the number of non-linear operations is proportional to the number of S-Boxes.

Overall, the S-Box type and the number of rounds should be chosen as follows.

- If you plan to use a prime field \mathbb{F}_p and $\gcd(\alpha, p - 1) = 1$ for $\alpha = 3$ or $\alpha = 5$, then
 - (1) use the S-Box $S(x) = x^\alpha$,
 - (2) select the minimum R_F that provides security against statistical attacks (i.e., at least $R_F = 6$), and
 - (3) select R_P that minimizes $tR_F + R_P$ such that Eq. (3) and Eq. (5) are not satisfied.
- For all other cases, we suggest to
 - (1) use the S-Box $S(x) = x^{-1}$,
 - (2) select the minimum R_F that provides security against statistical attack (i.e., at least $R_F = 6$), and
 - (3) select R_P that minimizes $tR_F + R_P$ such that Eq. (3) and Eq. (5) are not satisfied.

Before going on, we mention that other S-Boxes can be used as well (e.g., $x \mapsto x^7$). We have set up a script that calculates the number of rounds accordingly, using the security margin further described at the end of Section 5. Our resulting instances are given in Table 7 and Table 8.

Results via Script. A complete analysis on how to set up the script – in order to guarantee security and to find the best ratio between R_P and R_F – for this case has been proposed in [GLR⁺20]. For this reason, we refer to [GLR⁺20], and here we only report the minimum number of rounds necessary to provide security.

For completeness, we mention that the simplest way to set up the script is to test (e.g., by brute force) all possible values R_P and R_F that provide security (equivalently, for which the previous inequalities are satisfied), and find the ones that minimize the cost.

F Initialization of the Grain LFSR Used for Parameter Generation

We use the Grain LFSR for the generation of pseudo-random numbers.

1. First, the state is initialized with 80 bits b_0, b_1, \dots, b_{79} as follows:
 - (a) b_0, b_1 describe the field and are fixed to the binary representation of 1 for \mathbb{F}_p ,
 - (b) b_i for $2 \leq i \leq 5$ describe the S-Box,
 - (c) b_i for $6 \leq i \leq 17$ are the binary representation of n ,
 - (d) b_i for $18 \leq i \leq 29$ are the binary representation of t ,
 - (e) b_i for $30 \leq i \leq 39$ are the binary representation of R_F ,
 - (f) b_i for $40 \leq i \leq 49$ are the binary representation of R_P ,
 - (g) b_i for $50 \leq i \leq 79$ are set to 1.
2. Update the bits using $b_{i+80} = b_{i+62} \oplus b_{i+51} \oplus b_{i+38} \oplus b_{i+23} \oplus b_{i+13} \oplus b_i$.
3. Discard the first 160 bits.
4. Evaluate bits in pairs: If the first bit is a 1, output the second bit. If it is a 0, discard the second bit.

If a randomly sampled integer is larger than or equal to p , we discard this value and take the next one. We generate numbers starting from the most significant bit (i.e., the first bit sampled represents the most significant bit of the first constant). However, starting from the least significant bit or the most significant one makes no difference regarding the security.

Note that cryptographically strong randomness is not needed for the round constants or the matrices, and other methods can also be used.

G Concrete Instances with Security Margin

Concrete instances for x^3 -POSEIDON $^\pi$, x^5 -POSEIDON $^\pi$, and x^{-1} -POSEIDON $^\pi$, including the security margin, are given in Table 7, Table 8, and Table 9, respectively.

H Compact Constraints for STARKs and SNARKs

In this section we show how to generate constraints for S-Boxes that depend only on a few variables. This is useful when t is relatively small compared to R_P .

Let us denote the outputs of the ARK transformation in round r by $A_r^1, A_r^2, \dots, A_r^t$. Let us also denote the inputs to the MDS matrix in round r by $B_r^1, B_r^2, \dots, B_r^t$. We obtain

Table 7: Concrete instances for x^3 -POSEIDON $^\pi$ over \mathbb{F}_p , where S-Box(x) = x^3 .

Security M	Text Size $N = \log_2(p) \times t$	S-Box Size $(\lceil \log_2 p \rceil)$	# S-Boxes (t)	R_F	R_P	Cost Eq. (12)
128	1536	768	2	8	84	100
128	1536	384	4	8	84	116
128	1536	256	6	8	84	132
128	1536	192	8	8	88	152
128	1536	96	16	8	64	192
256	1536	768	2	8	170	186
256	1536	384	4	8	172	204
256	1536	256	6	8	174	222
256	1536	192	8	8	128	192
256	1536	96	16	8	64	192

Table 8: Concrete instances for x^5 -POSEIDON $^\pi$ over \mathbb{F}_p , where S-Box(x) = x^5 .

Security M	Text Size $N = \log_2(p) \times t$	S-Box Size $(\lceil \log_2 p \rceil)$	# S-Boxes (t)	R_F	R_P	Cost Eq. (12)
128	1536	768	2	8	56	72
128	1536	384	4	8	56	88
128	1536	256	6	8	60	108
128	1536	192	8	8	64	128
128	1536	96	16	8	48	176
256	1536	768	2	8	116	132
256	1536	384	4	8	116	148
256	1536	256	6	8	120	168
256	1536	192	8	8	88	152
256	1536	96	16	8	48	176

that in full rounds

$$S(A_r^i) = B_r^i,$$

whereas in partial rounds

$$S(A_r^t) = B_r^t, \quad A_r^i = B_r^i, i < t.$$

The S-Box inputs will then be

$$\begin{aligned} A_r^i \text{ for } i < t : \quad r \in [1; R_F/2] \cup [R_F/2 + 1 + R_P; R_F + R_P] \\ A_r^t : \quad r \in [R_F/2 + 1; R_F/2 + R_P]. \end{aligned} \tag{13}$$

It is obvious that the equations above plus the MDS and ARK affine transformations

$$MB_r + \mathbf{K}_r = \mathbf{A}_{r+1} \tag{14}$$

fully determine the permutation. Therefore, using $2Rt$ variables $\{A_r^i, B_r^i\}$ we can describe the permutation using constraints of degree equal to the S-Box degree d .

As we have mentioned, one can do better by substituting Eq. (14) for all B_r^i thus getting equations in only Rt variables $\{A_r^i\}$. In the same way we can get rid of A_r^i that are not inputs to S-Boxes, thus leaving with equations in $R_F t + R_P$ variables. However, this procedure creates equations with too many variables as S-Box outputs in the first partial rounds now depend on the S-Box inputs of all future partial rounds. Our goal is to construct compact equations in the same variables. We are going to work with two

Table 9: Concrete instances for x^{-1} -POSEIDON $^\pi$ over \mathbb{F}_p , where $S\text{-Box}(x) = x^{-1}$.

Security M	Text Size $N = \log_2(p) \times t$	S-Box Size $(\lceil \log_2 p \rceil)$	# S-Boxes (t)	R_F	R_P	Cost Eq. (12)
128	1536	768	2	8	66	82
128	1536	384	4	8	60	92
128	1536	256	6	8	60	108
128	1536	192	8	8	56	120
128	1536	96	16	8	32	160
256	1536	768	2	8	134	150
256	1536	384	4	8	128	160
256	1536	256	6	8	126	174
256	1536	192	8	8	88	152
256	1536	96	16	8	32	160

consecutive segments of t partial rounds each, let them be rounds from r to $r + 2t - 1$. We proceed as follows:

1. Express A_j^t , $j \in [r + t; r + 2t - 1]$ as affine functions of

$$A_{r+t}^1, A_{r+t}^2, \dots, A_{r+t}^t, B_{r+t}^t, B_{r+t+1}^t, \dots, B_{r+2t-1}^t.$$

2. Using Gaussian elimination, express $(A_{r+t}^1, A_{r+t}^2, \dots, A_{r+t}^t)$ as affine functions of

$$A_{r+t}^t, A_{r+t+1}^t, \dots, A_{r+2t-1}^t, B_{r+t}^t, B_{r+t+1}^t, \dots, B_{r+2t-1}^t.$$

3. Express B_j^t , $j \in [r; r + t - 1]$ as affine functions of

$$A_{r+t}^1, A_{r+t}^2, \dots, A_{r+t}^t, A_{r+1}^t, A_{r+2}^t, \dots, A_{r+t}^t.$$

4. Using Gaussian elimination, express $(A_{r+t}^1, A_{r+t}^2, \dots, A_{r+t}^t)$ as affine functions of

$$A_{r+1}^t, A_{r+2}^t, \dots, A_{r+t}^t, B_r^t, B_{r+1}^t, \dots, B_{r+t-1}^t.$$

5. Combine items 2 and 4 and get a system of t affine equations that link

$$A_{r+t}^t, A_{r+t+1}^t, \dots, A_{r+2t-1}^t, B_{r+t}^t, B_{r+t+1}^t, \dots, B_{r+2t-1}^t, \\ A_{r+1}^t, A_{r+2}^t, \dots, A_{r+t}^t, B_r^t, B_{r+1}^t, \dots, B_{r+t-1}^t.$$

6. Substitute $B_j^t \leftarrow S(A_j^t)$ and get a system of t degree- d polynomial constraints on $(A_{r+t}^t, A_{r+t+1}^t, \dots, A_{r+2t-1}^t, A_{r+1}^t, A_{r+2}^t, \dots, A_{r+t}^t)$.

The resulting system P_1 of polynomial constraints does not depend on r except for the constant term, which is determined by round constants. We do not have a formal proof that the systems composed at steps 1 and 3 have rank t but experiments show that it is the case for all matrices we tried.

We thus get the following system of constraints for the entire permutation:

- For the first group of full rounds: $t(R_F/2 - 1)$ constraints of degree d that link $(A_r^1, A_r^2, \dots, A_r^t, A_{r+1}^1, A_{r+1}^2, \dots, A_{r+1}^t)$, $r \in [1; R_F/2 - 1]$. The constraints depend on r in the constant terms only.

- Bridging the last full and t first partial rounds: t constraints of degree d that link $(A_{R_F/2}^1, A_{R_F/2}^2, \dots, A_{R_F/2}^t, A_{R_F/2+1}^t, A_{R_F/2+2}^2, \dots, A_{R_F/2+t}^t)$.

To get them, we express $A_{R_F/2+1}^t, A_{R_F/2+2}^2, \dots, A_{R_F/2+t}^t$ as affine functions of $(A_{R_F/2}^1, A_{R_F/2}^2, \dots, A_{R_F/2}^t, B_{R_F/2}^t, B_{R_F/2+1}^2, \dots, B_{R_F/2+t-1}^t)$ and then substituting B as degree- d functions.

- For all R_P partial rounds: divide them into groups of t rounds and use the system P_1 (derived above) to link them consecutively, thus $R_P - t$ constraints in total.
- Bridging the last t partial rounds and round $(R_P + R_F/2 + 1)$ (full one): t constraints of degree d .
- For the last group of full rounds: $t(R_F/2 - 1)$ constraints of degree d .

This totals to $tR_F + R_P - t$ constraints of degree d .

I Merkle Tree Instances of Poseidon

As a hash function used in a Merkle tree of a fixed arity always gets a message input of bounded length, it makes sense to have a compact padding of input elements. Concretely, we suggest the following:

- POSEIDON instances with width t are used for Merkle trees with arity $t - c$ where c is the capacity (at word level, namely c elements in \mathbb{F}). As the hash output is equal in size to the capacity, we should have $c|(t - c)$. In the prime field setting with 128-bit security and a 256-bit field we will have $c = 1$ and arity $t - 1$.
- A tree node may have from 0 to $t - c$ child elements. The missing child element is denoted by \emptyset and we denote $\widehat{\mathbb{F}} = \mathbb{F} \cup \{\emptyset\}$.
- A node hash function \widehat{H} maps $\widehat{\mathbb{F}}^{t-c}$ to \mathbb{F}^c . Therefore, a missing subtree of depth 1 (a single node) is represented as \emptyset , a missing subtree of depth 2 has the hash $\widehat{H}_\emptyset^2 = \widehat{H}(\emptyset, \emptyset, \dots, \emptyset)$, and a missing subtree of depth d has the hash

$$\widehat{H}_\emptyset^d = \widehat{H}(\widehat{H}_\emptyset^{d-1}, \widehat{H}_\emptyset^{d-1}, \dots, \widehat{H}_\emptyset^{d-1}).$$

- The node hash function \widehat{H} , based on the permutation Π of width t , is defined as follows:

$$\widehat{H}(X_{c+1}, X_{c+2}, \dots, X_t) = \Pi_{c+1\dots 2c}(\widetilde{X}_1, \widetilde{X}_2, \widetilde{X}_3, \dots, \widetilde{X}_t),$$

where $\Pi_{c+1\dots 2c}$ are the first c non-capacity elements of the output of Π and

$$\widetilde{X}_1 = \dots = \widetilde{X}_{c-1} = 0, \quad \widetilde{X}_c = \sum_i 2^i [X_i \neq \emptyset];$$

$$\widetilde{X}_{i>c} = \begin{cases} X_i, & X_i \neq \emptyset; \\ 0, & X_i = \emptyset \end{cases}$$

where $[]$ is the Iverson bracket (1 if the input is true, 0 otherwise).

Sponge Padding. For a variable-length sponge instance of POSEIDON, we pad all message strings from \mathbb{F}^* with a single element $1 \in \mathbb{F}$ and then, if necessary, with as many zero elements as necessary to have the message length be a multiple of $t - c$.

J About Starkad

Besides POSEIDON, in the previous version of the paper [GKK⁺19, Version: 20200205:104144], we also proposed STARKAD, which is basically the counterpart of POSEIDON for the case of binary fields. As already mentioned in the introduction, for this new version of the paper, we decided to focus on POSEIDON only, because all applications we considered are more efficient when working with a prime field rather than a binary field.

For this reason, here we limit ourselves to recall few details of STARKAD. Before going on, we highlight that “weaker” versions of STARKAD have recently been broken in [KR20, BCD⁺20]. In the following, we explain how to fix the design in order to prevent the attack presented in these papers.

The Permutation Starkad^π

The hash function²⁴ STARKAD for the *binary case* is constructed by instantiating a sponge construction [BDPA08] with the permutation STARKAD^π. This permutation is almost equivalent to the permutation POSEIDON^π defined in Section 2.3, with only some minor differences, namely

- (1) it is defined over $\mathbb{F}_{2^n}^t$ instead of \mathbb{F}_p^t (it uses binary fields instead of prime fields), and
- (2) in order to define the S-Box as $S(x) = x^\alpha$, the positive integer $\alpha \geq 3$ must be chosen as the smallest integer s.t. $\gcd(\alpha, 2^n - 1) = 1$ (e.g., $\alpha = 3$ if and only if n is odd).

In order to prevent the attacks presented in [KR20, BCD⁺20], it is sufficient to choose the MDS matrix that defines the linear layer as described in Section 2.3. Namely, we assume that there exists no infinitely long subspace trail (either invariant or iterative and either with active or inactive S-Boxes) for the rounds with partial S-Box layers. In order to satisfy this assumption, it is sufficient to test the MDS matrix using the tools provided in [GRS20].

Cryptanalysis and (Old) Number of Rounds of Starkad^π

The security analysis of STARKAD^π is similar to the one proposed for POSEIDON^π. In particular, *the number of rounds necessary to provide security for STARKAD^π over $\mathbb{F}_{2^n}^t$ cannot be smaller than the number of rounds for POSEIDON^π over \mathbb{F}_p^t where $p > 2^n$ is the prime closest to 2^n .*

At the same time, some important differences arise:

- We expect that the higher-order differential attack is more competitive (or at least, not less competitive) in the case of a binary field than in the case of a prime field. This means that in principle, more rounds are necessary in order to provide security of STARKAD^π compared to POSEIDON^π.
- The MDS matrices proposed for the previous version of STARKAD^π were weak matrices, in the sense that their squares were multiples of the identity. This weakness has been showed and exploited in [KR20, BCD⁺20]. The new assumption on the MDS matrix prevents this weakness.

²⁴*About the name:* Starkad was a legendary hero in Norse mythology, who used to hash his enemies with 2^2 swords in 2^3 arms.

Regarding the first point, we decided not to present a complete analysis of security of STARKAD $^\pi$ w.r.t. higher-order differential attacks. We point out that several works have recently been published in the literature (see [EGL⁺20, CGL⁺20] for more details) whose goal is to better understand the growth of the degree in the case of schemes (over binary fields) with large S-Boxes. An open problem is to better understand how the combination of rounds with full and rounds with partial S-Box layers influences such growth of the algebraic degree.

In the following table, we limit ourselves to recall the number of rounds proposed in the previous version of the paper [GKK⁺19, Version: 20200205:104144]. *We stress that we do not claim that STARKAD $^\pi$ instantiated with this number of rounds is secure.*

Table 10: Old range of different parameter sets for STARKAD $^\pi$ instantiated by S-Box(x) = x^3 (with security margin). *We do not claim that the following round numbers provide security.*

Security M	Text Size $N = n \times t$	S-Box Size (n)	# S-Boxes (t)	R_F	R_P	Field
128	1512	63	24	8	45	\mathbb{F}_{2^n}
128	1551	33	47	8	25	\mathbb{F}_{2^n}
128	1581	31	51	8	24	\mathbb{F}_{2^n}
256	1512	63	24	8	45	\mathbb{F}_{2^n}
256	1551	33	47	8	25	\mathbb{F}_{2^n}
256	1581	31	51	8	24	\mathbb{F}_{2^n}