Cryptanalytic Audit of the XHash Sponge Function and its Components

Vincent Rijmen
vincent.rijmen@kuleuven.be

April 29, 2024

Executive Summary

In this audit we started from the security analysis provided in the design documentation [1]. We extended the analysis in several directions and confirmed the security claims that were made by the designers.

The simple algebraic description of the S-boxes made it possible to extend the original analysis of the security of the XHash permutations against differential cryptanalysis, resulting in some interesting new properties. These properties did not result in any attacks.

We added an analysis of the security of the XHash8 permutation against saturation attacks and made a few comments on linear cryptanalysis of the permutation.

Finally, we examined the XHash padding rule, the RPX representation of the XHash permutations and the XHash sponge function. They have all been defined according to the state of the art and maintain the security level of the underlying XHash8/XHash12 permutations.

Contents

1 Introduction

This report is the result of an audit of the XHash hash functions. In particular, the security of the following components was evaluated:

1. The XHash permutations;
2. The XHash padding rule;
3. The RPX representation of the XHash permutations
4. The XHash sponge function.

The audit is based exclusively on the textual specification [1]. The compliance of available software implementations to the specification was not part of this audit.

2 Differential attacks

Resistance against ordinary differential attacks is argued in the design documentation [1, Section 5.1]. In this section we comment on some aspects of that study and extend it. All arithemtic operations in this section are taken modulo p, where $p = 2^{64} - 2^{32} + 1$ is the 64-bit prime number used in the definition of XHash8/12.

2.1 Ordinary differential attacks

The designers state: "an adversary controls only the outer part of the sponge and therefore they can only create a difference in at most 8 field elements."

This statement refers to the use of XHash8/XHash12 in a sponge construction. However, apparently it does not take into account the padding rule. Since an adversary can modify the length of the input, they can also influence the ninth state element, which contains the domain separation identifier. Hence an adversary can create a difference in $8 + 1 = 9$ field elements on the input.

Luckily, this observation does not affect the bound on the probability of a characteristic given in [1, Section 5.1], since that bound is determined by the minimum number of active π_0 and π_1 S-boxes, which is given by $\beta_F - 4$, where β_F denotes the branch number of the linear layer of F (hence $\beta_F = 13$).

2.2 Differential uniformity of the S-box

The designers give a bound for the differential uniformity of a power map [1, Theorem 5.1]. The differential uniformity is determined by the number of roots for the following polynomial:

$$q(x) = (x + \alpha)^\gamma - x^\gamma - \beta \tag{1}$$

(The last term is missing in [1].) Clearly, the bound of Theorem 5.1 can be strengthened for invertible maps, since the differential uniformity of a map equals the differential uniformity of its inverse. Therefore, all S-boxes of XHash8/12 have the same differential uniformity, which is at most 6.

Working out (1) for $\gamma = 7$, we obtain:

$$q(x) = 7\alpha x^6 + 21\alpha^2 x^5 + 35\alpha^3 x^4 + 35\alpha^4 x^3 + 21\alpha^5 x^2 + 7\alpha^6 x + \alpha^7 - \beta$$

Inspired by the study of the differential uniformity of AES [3], we investigate the special case $\alpha^7 = \beta$. We obtain

$$q(x) = 7\alpha x(x + \alpha)(x^2 + \alpha x + \alpha^2)^2$$

Since

$$p - 3 = (2^{33} - 1)^2 \bmod p,$$

the quadratic polynomial $x^2 + \alpha x + \alpha^2$ has two roots and

$$x^2 + \alpha x + \alpha^2 = (x - \alpha(2^{32} - 1))(x + \alpha 2^{32})$$

We conclude that all the differentials (α, α^7) have exactly 4 right pairs. There might well exist other choices of α and β that result in a higher number of right pairs, but we did not find any.

2.3 Number of right pairs in XHash components

The simple algebraic description of the XHash8/XHash12 S-boxes leads to the following property.

Theorem 1. *All differentials (α, β) over an XHash8/XHash12 S-box with $\beta \neq \alpha^7/64$ have an even number of right pairs.*

Proof. As explained earlier, it suffices to give a proof for S-box π_0. Let $(u, u + \alpha)$ be a right pair for π_0. Then we have from (1)

$$(u + \alpha)^7 - u^7 - \beta = 0$$

Using

$$\begin{aligned}(u + \alpha)^7 &= -(-u - \alpha)^7 \\ u^7 &= -(-u)^7 = -((-u - \alpha) + \alpha)^7\end{aligned}$$

we obtain

$$((-u - \alpha) + \alpha)^7 - (-u - \alpha)^7 - \beta = 0$$

which means that $(-u - \alpha, -u)$ is a right pair for the differential (α, β). The pairs $(u, u + \alpha)$, $(-u - \alpha, -u)$ are two different pairs if and only if $u \neq -\alpha/2$. Using (1) again we obtain that the pair $(-\alpha/2, \alpha/2)$ can be a right pair only if $\beta = \alpha^7/64$. Hence for all other differentials, the number of right pairs must be even. □

The proof of Theorem 1 introduces a kind of mixed quartets (u_0, u_1, u_2, u_3) defined by $u_1 = u_0 + \alpha$, $u_2 = -u_1$ and $u_3 = u_2 + \alpha$. Each quartet defines 0 or 2 right pairs for a differential (α, β), except when $\beta = \alpha^7/64$. It can easily be verified that the application of an XHash8/XHash12 S-box to a mixed quartet (u_0, u_1, u_2, u_3) results in a new mixed quartet (v_0, v_1, v_2, v_3) with $v_1 = v_0 + \beta$, $v_2 = -v_1$ and $v_3 = v_2 + \beta$. It follows that a differential over a series of S-box applications will also have an even number of right pairs, except if we are unlucky and hit somewhere the special $\beta = \alpha^7/64$ value. Furthermore, the application of the MC layer to a mixed quartet results in a new mixed quartet. Hence also a differential over a series of S-box and MDS applications will typically have an even number of right pairs. Since addition with a constant does not transform a mixed quartet into a mixed quartet, the property does not hold over a whole step of XHash8/XHash12.

2.4 Security of XHash12 versus XHash8

The designers claim that the resistance of XHash12 against differential cryptanalysis is at least as high as the resistance of XHash8 against differential cryptanalysis, *because for any trail* pattern*, XHash12 activates the same number or more S-boxes than XHash8* [1]. While this claim might look plausible at first sight, we argue here that there is no guarantee for its correctness.

Let's measure security against differential cryptanalysis by studying (the maximum of) the EDP values of the characteristics. Denote by $\mathcal{N}_X$ the number of characteristics with nonzero EDP value over the transformation X. If X is the operation (MC), then the input difference determines uniquely the output difference. Hence, $\mathcal{N}_{(MC)} = p^{12}$. If X contains nonlinear elements, then $\mathcal{N}_X$ will increase because each input difference can result in several output differences. For example, we can use the bound on the differential uniformity of the S-boxes to obtain $\mathcal{N}_{\pi_0} \geq 1 + (p-1)^2/6$. For the step (F) we get

$$\mathcal{N}_{(F)} = 1 + 12\mathcal{N}_{\pi_0} + \binom{12}{2}(\mathcal{N}_{\pi_0})^2 + \cdots + \binom{12}{12}(\mathcal{N}_{\pi_0})^{12} = (1+\mathcal{N}_{\pi_0})^{12} \approx (\mathcal{N}_{\pi_0})^{12}$$

Symmetry implies that $\mathcal{N}_{(B)} = \mathcal{N}_{(F)}$. However,

$$\mathcal{N}_{(B')} \approx (\mathcal{N}_{\pi_0})^8 < \mathcal{N}_{(B)}$$

It follows that $\mathcal{N}_{\text{XHash8}} < \mathcal{N}_{\text{XHash12}}$. It seems hasty to conclude that the maximum of the EDP of the characteristics in the smaller set is guaranteed to be larger than the maximum of the EDP of the characteristics in the larger set. In particular, from existing work on AES characteristics we know that the characteristics with the largest EDP values show a typical pattern where rounds with many active S-boxes are followed by rounds with only few active S-boxes. The existence of such characteristics relies on the fact that the S-boxes have for each input difference many output differences possible, hence it is possible to produce a difference that will be converted by the MDS layer to differences with only a few active S-boxes in the next step. Since XHash8 has steps with fewer S-boxes, there are less possibilities to produce a suitable difference.

2.5 Higher-order differential attacks

A single round of XHash8 is vulnerable to a higher-order differential attack, that we describe here. Let $b = (F)(a)$, $c = (B')(b)$. Then

$$b_s = \mathbf{C}_{i|s} + \sum_{t=0}^{11} \mathbf{M}_{s,t}(a_t)^7, s = 0, \ldots, 11$$

$$c_s = \mathbf{C}_{i+1|s} + b_s = \mathbf{C}_{i+1|s} + \mathbf{C}_{i|s} + \sum_{t=0}^{11} \mathbf{M}_{s,t}(a_t)^7, s = 1, 4, 7, 10$$

It follows that c_1, c_4, c_7 and c_{10} can be expressed as functions of $a_0, \ldots, a_{11}$ with degree 7. Hence it is trivial to construct differentials of order 7 with probability 1. We see however no way to extend these differentials over more steps.

2.6 Plateau trails and related differentials

Reasonings on the security of cryptographic primitives against differential cryptanalysis are often implicitly based on the *Hypothesis of stochastic equivalence*. Reformulated in statistical terms, this hypothesis states that the differential probability of a differential/characteristic is very close to the *expected* differential probability, which is computed by assuming that all roundkeys are independent uniformly distributed random variables. In the case of block ciphers, there are no experimental results supporting the validity of this hypothesis. A fortiori, the validity can be questioned in the case of hash functions or permutation-based primitives. The recently proposed *quasi-differential cryptanalysis* tries to remedy this situation [2]. For many interesting cases, it is not known yet how to overcome the computational challenges posed by this framework.

Plateau trails can be see as a special case of quasi-differential trails. It has been observed that plateau trails are caused by a special property of the diffusion layer, called *related differentials*. Two differentials (α, β), $(\alpha^\diamond, \beta^\diamond)$ over an n-component map M are called related differentials if for $i = 0, 1, \ldots, n-1$:

$$\alpha_i = 0 \text{ OR } \alpha_i^\diamond = 0 \text{ OR } \alpha_i = \alpha_i^\diamond$$
$$\beta_i = 0 \text{ OR } \beta_i^\diamond = 0 \text{ OR } \beta_i = \beta_i^\diamond$$

For the MixColumns layer of AES, related differentials with Hamming weight 5 are known. (5 is the minimum weight of a nontrivial differential over MixColumns.)

For the MDS layer of XHash8/XHash12, no related differentials with Hamming weight 13 are known. We searched for related differentials by extrapolating from the related differentials of the MixColumns layer of AES, but did not find any. We conjecture that none exist.

3 Saturation attacks

The design documentation [1] does not mention security against saturation attacks. We introduce here a distinguisher for a single round of XHash8/XHash12 based on saturation properties.

3.1 Saturating F_p

Using straightforward techniques, a saturation distinguisher can be constructed for the sequence of two steps $(F)(B)$ or $(F)(B')$. A set of inputs with 11 coordinates fixed and one coordinate saturated will be transformed to a set with all 12 coordinates saturated. There seems to be no way to propagate this property through the $(P3)$ step, if we reason on F_p only.

3.2 Saturating F_{p^3}

The design documentation [1] uses F_p as the *base* field, in which most components of the hash functions are described. The S-boxes π_2 are then the exception. However, it is also possible to describe the hash functions using F_{p^3} as the base field. If we take this approach, then a state a can be described as a 3-tuple $(\mathbf{a}_0, \mathbf{a}_1, \mathbf{a}_2)$ with $\mathbf{a}_i \in F_{p^3}$. The S-boxes π_0 and π_1 become opaque invertible S-boxes. Also the description of the MDS layer becomes more complicated,

but it can easily be seen that the *MDS properties* are maintained. In particular, we have the following.

Theorem 2. *Let* $(\alpha, \beta) = (\alpha_0, \ldots, \alpha_{11}; \beta_0, \ldots, \beta_{11})$ *denote a differential over the MDS layer of XHash8/XHash12. If*

$$\alpha_3 = \alpha_4 = \cdots = \alpha_{11} = 0$$

then the probability of (α, β) *is nonzero only if*

$$(\beta_0, \beta_1, \beta_2) \neq (0,0,0), (\beta_3, \beta_4, \beta_5) \neq (0,0,0), (\beta_6, \beta_7, \beta_8) \neq (0,0,0) \text{ and } (\beta_9, \beta_{10}, \beta_{11}) \neq (0,0,0)$$

Proof. Let $\mathbf{M}_0$ denote the 3×3 submatrix of $\mathbf{M}$ consisting of columns $0, 1, 2$ and rows $0, 1, 2$. Since $\mathbf{M}$ is an MDS matrix, $\mathbf{M}_0$ is invertible. Hence an input difference $(\alpha_0, \alpha_1, \alpha_2, 0, \ldots, 0)$ with $(\alpha_0, \alpha_1, \alpha_2) \neq (0,0,0)$ is mapped to an output difference $\beta_0, \ldots, \beta_{11})$ with $(\beta_0, \beta_1, \beta_2) \neq (0,0,0)$. Applying a similar reasoning we obtain that also $(\beta_3, \beta_4, \beta_5) \neq (0,0,0)$, $(\beta_6, \beta_7, \beta_8) \neq (0,0,0)$ and $(\beta_9, \beta_{10}, \beta_{11}) \neq (0,0,0)$. □

Let R denote a single round, i.e. the sequence of the three steps $(F)(B)(P3)$ or $(F)(B')(P3)$. Let R^- denote a single round without the last MDS layer. Then we have the following.

Theorem 3. *Let* S *be a saturated input set*

$$S = \{a \in (F_{p^3})^3 | \mathbf{a}_1 = c_1, \mathbf{a}_2 = c_2\}$$

with c_1, c_2 *two arbitrary constants. Then*

1. $R^-(S)$ *is saturated in the following way:*

$$\forall a, b \in R(S) \text{ with } a \neq b : (a_0, a_1, a_2) \neq (b_0, b_1, b_2), (a_3, a_4, a_5) \neq (b_3, b_4, b_5),$$
$$(a_6, a_7, a_8) \neq (b_6, b_7, b_8), \text{ and } (a_9, a_{10}, a_{11}) \neq (b_9, b_{10}, b_{11})$$

2. *The set* $R(S)$ *has the* sum *property:*

$$\sum_{a \in S} R(a) = (0, 0, \ldots, 0)$$

Proof.

1. Two different elements of S have the property that their difference equals $(\alpha_0, \alpha_1, \alpha_2, 0, \ldots, 0)$ with $(\alpha_0, \alpha_1, \alpha_2) \neq (0,0,0)$. This property passes through the S-box layer π_0 with probability 1. Using Theorem 2, we see that the MDS-layer of the step (F) expands this property and produces a difference $\beta = (\beta_0, \ldots, \beta_{11})$ with $(\beta_{4i}, \beta_{4i+1}, \beta_{4i+2}) \neq (0,0,0), i = 0, 1, 2, 3$. Steps (B) or (B') transform this difference to an output difference with the same property, and the same is true for the S-box layer π_2.

2. For any field F_q with $q > 2$ we have

$$\sum_{x \in F_q} x = 0$$

Using part (1) we obtain that

$$\sum_{a \in S} R(a) = \left(\sum_{x \in F_{p^3}} x, \sum_{x \in F_{p^3}} x, \sum_{x \in F_{p^3}} x, \sum_{x \in F_{p^3}} x \right)$$
$$= (0,0,0;0,0,0;0,0,0;0,0,0)$$

□

It follows that we have a saturation distinguisher over 1 full round of the hash function. We see no way to extend this distinguisher over more steps.

4 Linear attacks

Linear cryptanalysis of non-binary ciphers is defined in [4] by considering the group characters χ_u, where

$$\chi_u(x) = e^{\frac{2\pi i}{p} ux}$$

The correlation of an approximation (u, v) of π_0 is given by:

$$\mathrm{cor}_{\pi_0}(u,v) = \frac{1}{p} \sum_{x \in F_p} e^{\frac{2\pi i}{p}(vx^7 - ux)} = \frac{1}{p} \sum_{z \in F_p} e^{\frac{2\pi i}{p}(vu^{-7}z^7 - z)}$$

It follows that $\mathrm{cor}_{\pi_0}(u,v) = \mathrm{cor}_{\pi_0}(wu, w^7 v), \forall w \in F_p$, which will cause some patterns in the correlation matrix of π_0. Observe that the correlation matrix of π_1 is the transpose of the correlation matrix of π_0. Hence, it exhibits the same patterns. Since π_2 is based on the same power mapping, but over a larger field, the correlation matrix of π_2 exhibits similar patterns.

The size of the components used by XHash8/XHash12 makes it impossible to compute correlation matrices in the naive way. We see no way to accelerate this computation. Hence it is difficult to obtain more results. Furthermore, even if one could find linear approximations with a relatively high correlation over the whole hash function, it is not known how to exploit them for an attack on a hash function.

References

[1] Tomer Ashur, Al Kindi, Mohammad Mahzoun. XHash8 and XHash12: efficient STARK-friendly hash functions.

[2] Tim Beyne, Vincent Rijmen. Differential cryptanalysis in the fixed-key model. CRYPTO 2022, LNCS Vol. 13509, pp. 687–716.

[3] Joan Daemen, Vincent Rijmen. Understanding two-round differentials in AES. SCN 2006, LNCS Vol. 4116, pp. 78–94.

[4] Thomas Baignères, Jacques Stern, Serge Vaudenay. Linear cryptanalysis of non binary ciphers. SAC 2007. LNCS Vol. 4876, pp. 184–211.