

Security Analysis of BLAKE2's Modes of Operation

Atul Luykx¹, Bart Mennink¹ and Samuel Neves²

¹ Dept. Electrical Engineering, ESAT/COSIC, KU Leuven, and iMinds, Belgium
atul.luykx@esat.kuleuven.be, bart.mennink@esat.kuleuven.be

² CISUC, Dept. of Informatics Engineering, University of Coimbra, Portugal
sneves@dei.uc.pt

Abstract. BLAKE2 is a hash function introduced at ACNS 2013, which has been adopted in many constructions and applications. It is a successor to the SHA-3 finalist BLAKE, which received a significant amount of security analysis. Nevertheless, BLAKE2 introduces sufficient changes so that not all results from BLAKE carry over, meaning new analysis is necessary. To date, all known cryptanalysis done on BLAKE2 has focused on its underlying building blocks, with little focus placed on understanding BLAKE2's generic security. We prove that BLAKE2's compression function is indifferentiable from a random function in a weakly ideal cipher model, which was not the case for BLAKE. This implies that there are no generic attacks against any of the modes that BLAKE2 uses.

Keywords: BLAKE · BLAKE2 · hash function · indifferentiability · PRF

1 Introduction

Widespread adoption of cryptographic algorithms in practice often occurs regardless of their scrutiny by the cryptographic community. Although competitions such as AES and SHA-3 popularize thoroughly analyzed algorithms, they are not the only means with which practitioners find new algorithms. Standards, textbooks, and social media are sometimes more effective than publications and competitions.

Nevertheless, analysis of algorithms is important regardless of how they were popularized, and can result in finding insecurities, but also new techniques. For example, the PLAID protocol avoided cryptographic scrutiny by being standardized via the *Cards and personal identification* subcommittee of ISO, instead of via the *Cryptography and security mechanisms* working group, and when properly analyzed, PLAID turned out to be significantly weaker than claimed [DFP⁺14]. Similarly the ANSI authenticated encryption algorithm EAX' modified Bellare, Rogaway, and Wagner's EAX algorithm, thereby introducing security vulnerabilities [MLM13]. In other cases modifications actually improve security, as with AMAC [BBT16], which processes the output of a hash function to construct a MAC.

BLAKE2. Since its introduction in 2013, the hash function BLAKE2 has seen quick adoption, despite the fact that it had not received as much analysis as the SHA-3 finalists. It is a modification of the SHA-3 finalist BLAKE, which has high software performance and withstood extensive cryptanalysis [CPB⁺12, Section 3.1]. BLAKE2 simplifies BLAKE, resulting in better efficiency and ultimately its use in numerous constructions [FLW14, BDK16, CJMS14, JAA⁺15, AABS14, HKR15] and applications [Per16, HBHW16, Ros13].

Although BLAKE2 is based on BLAKE, one cannot claim that BLAKE's cryptanalysis [BNR11, AGK⁺10a, DK11] directly carries over. Nevertheless, the cryptanalytic techniques used for BLAKE can generally be applied to BLAKE2, resulting in an increasing amount of novel cryptanalysis [GKN⁺14, Hao14, KNP⁺15, EFK15]. The generic security of BLAKE2's mode, however, has not yet been analyzed, and the indistinguishability analysis of BLAKE [ALM12, CNY11] does not carry over. The main reason for this is that BLAKE2 weakens its underlying compression function and uses it in many different modes: the plain HAIFA mode, a tree mode, a parallelized mode, and a keyed HAIFA mode. Additionally, these modes initialize the state using the salt, that can be freely chosen by the user.

Even slight modifications to modes or the underlying primitives might introduce vulnerabilities. Besides the EAX example given above, other examples include Dual Counter Mode [BS01, DGW01] versus IAPM [Jut01], and the masking used in OTR [Min14, BS16] versus the classical XEX-masking [Rog04]. Therefore, properly analyzing the security of the BLAKE2 modes of operation is important.

Results. Unlike BLAKE, the BLAKE2 compression function already achieves indistinguishability *at the compression function level*. Using a weakly ideal block cipher, we prove that the compression function is indistinguishable from a random function up to a query complexity of about $2^{n/2}$, where n is the state size of the compression function. The derivation in part relies on the fact that the BLAKE2 compression function can be seen as a $7n/2$ -to- n -bit compression function based on a $2n$ -bit block cipher.

Using the indistinguishability composition result from Maurer et al. [MRH04], the indistinguishability of the BLAKE2 hashing modes (based on an idealized underlying block cipher) directly follows from the already existing indistinguishability analyses on the modes (based on an ideal compression function) and the newly derived indistinguishability result of the BLAKE2 compression function. In other words, rather than deriving three tedious hash function indistinguishability proofs—as were the norm for the SHA-3 finalists [ALM12, CNY11, AMP10, BMN10, MPS16, BDPV08, BKL⁺09]¹—for BLAKE2 it suffices to derive a surprisingly short and simple compression function indistinguishability result and rely on the generic indistinguishability of the overlying modes. Note that our results also imply that the BLAKE2 compression function could be used in *any* hash function mode that has an indistinguishability proof if the underlying compression function is ideal.

We furthermore consider security of the keyed version of BLAKE2, and demonstrate that the newly obtained indistinguishability result on the BLAKE2 compression function immediately entails strong PRF-security of keyed BLAKE2 in the multi-key setting as long as the total query complexity is at most $2^{n/2}$.

2 BLAKE2

BLAKE2 consists of a compression function that internally uses a block cipher. This compression function is used to instantiate various keyless and keyed modes. We will discuss the block cipher and compression function design in this section, omitting technical details that are irrelevant for the generic analysis. We refer to the original publication [ANWW13] and the RFC [SA15] for details. The keyless hashing modes are discussed in Section 5 and the keyed hashing mode in Section 6.

Throughout, we adopt the following notation. For two bit strings x, y , their concatenation is interchangeably denoted by $x\|y$, (x, y) , and xy . For a bit string y of even length, we denote by $L(y)$ its left half and by $R(y)$ its right half, so that $y = L(y)\|R(y)$. We denote by $n \in \{256, 512\}$ the state size of the hash and compression function, and $w = n/8$ the word

¹As well as beyond the SHA-3 finalists, as almost all known compression functions are differentiable, the only known exceptions being the compression function of MD6 [DRRS09] and some double length compression functions [Men13].

size. The fixed initialization value is denoted by $IV \in \{0, 1\}^n$; we refer to [ANWW13] for the specification of the initialization value, where the only important property of IV is that $L(IV) \neq aaaa$ for any word $a \in \{0, 1\}^w$. A $2n$ -bit state $aaaabbbbccccdddd$ may be uniquely identified by its representation as a 4×4 matrix

$$\begin{pmatrix} a & a & a & a \\ b & b & b & b \\ c & c & c & c \\ d & d & d & d \end{pmatrix};$$

we use either notation interchangeably.

2.1 Block Cipher

BLAKE2 internally uses a block cipher $E : \{0, 1\}^{2n} \times \{0, 1\}^{2n} \rightarrow \{0, 1\}^{2n}$. In this work, we will focus on the generic security of BLAKE2, and consider E as an idealized block cipher. However, while BLAKE’s underlying block cipher had no known weaknesses and could reasonably be modeled as an ideal cipher, this is no longer the case in BLAKE2. In particular, the property

$$E \left(\begin{pmatrix} k & k & k & k \\ k & k & k & k \\ k & k & k & k \\ k & k & k & k \end{pmatrix}, \begin{pmatrix} a & a & a & a \\ b & b & b & b \\ c & c & c & c \\ d & d & d & d \end{pmatrix} \right) = \begin{pmatrix} a' & a' & a' & a' \\ b' & b' & b' & b' \\ c' & c' & c' & c' \\ d' & d' & d' & d' \end{pmatrix}, \quad (1)$$

for arbitrary words $a, b, c, d, k \in \{0, 1\}^w$ can be used to efficiently distinguish it from an ideal cipher. This property is a central part of the “chosen-IV” attacks of Guo et al. [GKN⁺14, Section 3 and 4], and is the generalization of a well-known property of permutations derived from ChaCha [BHH⁺15, Section 4]. As discussed in Section 3, we will deal with this caveat by modeling E as a weakly ideal cipher.

2.2 Compression Function

The BLAKE2 compression function $F : \{0, 1\}^n \times \{0, 1\}^{2n} \times \{0, 1\}^{n/4} \times \{0, 1\}^{n/4} \rightarrow \{0, 1\}^n$ gets as input a state value h , a message m , a counter t , and a flag f , and is defined as follows:

$$x \leftarrow h \parallel 0^{n/2} \parallel t \parallel f \oplus 0^n \parallel IV \quad (2a)$$

$$y \leftarrow E(m, x) \quad (2b)$$

$$h' \leftarrow L(y) \oplus R(y) \oplus h \quad (2c)$$

$$\mathbf{return } h', \quad (2d)$$

where we recall that IV is a fixed initialization value throughout this work. The BLAKE2 compression function F is depicted in Figure 1. Note that for the specification of the compression function, we have not put any restrictions on the input values h, t, f : an adversary has full freedom to select these values. In the BLAKE2 hashing and MAC modes (Sections 5 and 6), the counter t and flag f are subject to specific formats.

For further analysis, we will also require the “inverse” of equation (2a). Define by \mathbf{parse}_{IV} the mapping that gets as input an $x \in \{0, 1\}^{2n}$, and outputs the unique $(h, z, t, f) \in \{0, 1\}^n \times \{0, 1\}^{n/2} \times \{0, 1\}^{n/4} \times \{0, 1\}^{n/4}$ such that

$$h \parallel z \parallel t \parallel f = x \oplus 0^n \parallel IV.$$

Note that the z -value equals $0^{n/2}$ if and only if x could appear in the computation of (2a). We define $\mathbf{parse}_{IV,z} : \{0, 1\}^{2n} \rightarrow \{0, 1\}^{n/2}$ as the function \mathbf{parse}_{IV} restricted to the z -value, or more formally,

$$\mathbf{parse}_{IV,z}(x) = L(R(\mathbf{parse}_{IV}(x))).$$




Figure 1: BLAKE2 compression function

3 Security Model

For $l, m \in \mathbb{N}$ such that $l \geq m$, denote by $\text{Func}(l, m)$ the set of all functions $F : \{0, 1\}^l \rightarrow \{0, 1\}^m$, and by $\text{Block}(m)$ the set of all block ciphers $E : \{0, 1\}^m \times \{0, 1\}^m \rightarrow \{0, 1\}^m$.

3.1 Weakly Ideal Cipher Model

A naive way of modeling the block cipher E in BLAKE2 would be to consider it as an ideal cipher: $E \xleftarrow{\$} \text{Block}(2n)$. However, this solution would not properly capture the structural property (1) of the BLAKE2 block cipher as discussed in Section 2.1. Instead, we will generate E as a *weakly ideal cipher*, i.e., an ideal cipher with the restriction that it adheres to property (1). The approach shows similarities with the weakly ideal compression functions used by Liskov [Lis06] to prove security of the zipper hash function if the underlying compression function can be inverted. The weakly ideal cipher also resembles ideas of the indistinguishability analyses of the SHA-3 candidates Shabal if the underlying block cipher shows some non-random behavior [BCC⁺09] and SIMD if the underlying compression function is distinguishable from a random function [BFL10]. We remark that, unlike these works, our analysis should not be seen as a patch to an unexpected property of BLAKE2. It appears that the property was known to the designers in advance (see, e.g., [AGK⁺10b, Appendix C]) and simply accepted as being a harmless property. The weakly ideal cipher model could also be seen as a specific instance of the weakened models of Katz et al. [KLT15] and Mennink and Preneel [MP15], but these models are much more general and more involved.

In more detail, consider the partition $\{0, 1\}^{2n} = W \cup S$ (W for weak and S for strong), where:

$$W = \{aaaabbbbccccddd \in \{0, 1\}^{2n} \mid a, b, c, d \in \{0, 1\}^w\}, \quad (3)$$

$$S = \{0, 1\}^{2n} \setminus W. \quad (4)$$

We define by $\text{Block}^*(2n)$ to be the set of all block ciphers $E \in \text{Block}(2n)$ with the additional restriction that

$$E \left(\left(\begin{array}{cccc} k & k & k & k \\ k & k & k & k \\ k & k & k & k \\ k & k & k & k \end{array} \right), \cdot \right) \quad (5)$$

is W - and S -subspace invariant for all $k \in \{0, 1\}^w$, that is, inputs in W map to W , and

likewise for S . For notational simplicity, define the set of weak keys for E as

$$\text{WK} = \{kkkkkkkkkkkkkkkkkk \in \{0, 1\}^{2n} \mid k \in \{0, 1\}^w\}. \quad (6)$$

A random $E \stackrel{\$}{\leftarrow} \text{Block}^*(2n)$ can now be modeled as follows: on input of $(k, x) \in \text{WK} \times W$, it generates its response y randomly from W up to repetition; on input of $(k, x) \in \text{WK} \times S$, it generates its response y randomly from S up to repetition. For key values $k \in \{0, 1\}^{2n} \setminus \text{WK}$, it behaves like an ideal cipher. The case of inverse queries is analogous.

We remark that, by resorting to the weakly ideal cipher model, we do not make stronger assumptions than those used in previous results, and, despite the fact that we give distinguishers more power (by weakening the cipher), we are able to get similar results. Concerning the most up to date cryptanalysis on BLAKE2's block cipher, there is currently no reason to believe that it does not approximate a weakly ideal cipher as we define it.

3.2 Indifferentiability

One way to measure the extent to which a certain cryptographic function behaves like a random function is via the indistinguishability framework, where a distinguisher is given access to either the cryptographic function or the random function, with the goal to distinguish both worlds. The indistinguishability security model inherently relies on the existence of secret information in both worlds, either a secret key, or the random function. Therefore, for keyless cryptographic hash functions the indistinguishability framework is inadequate, and we will use the indifferentiability framework of Maurer et al. [MRH04]. At a high level, the indifferentiability framework measures the distance from a construction $\mathcal{C}^{\mathcal{P}}$ based on an ideal subcomponent \mathcal{P} , for instance a compression function based on an ideal cipher or a hash function based on a compression function, to an ideal functionality \mathcal{R} , and it guarantees that a construction has no structural design flaws. In this work, we employ the adaptation and simplification of the indifferentiability framework by Coron et al. [CDMP05]. We note that this indifferentiability framework only applies to single-stage games; cf., Ristenpart et al. [RSS11].

Definition 1. Let \mathcal{C} be a construction with oracle access to an ideal primitive \mathcal{P} . Let \mathcal{R} be an ideal primitive with the same domain and range as \mathcal{C} . Let \mathcal{S} be a simulator with the same domain and range as \mathcal{P} with oracle access to \mathcal{R} , and let \mathcal{D} be a distinguisher. The differentiability advantage of \mathcal{D} is defined as

$$\text{Indiff}_{\mathcal{C}^{\mathcal{P}}, \mathcal{S}}(\mathcal{D}) = \left| \mathbb{P}(\mathcal{D}^{\mathcal{C}^{\mathcal{P}}, \mathcal{P}} = 1) - \mathbb{P}(\mathcal{D}^{\mathcal{R}, \mathcal{S}^{\mathcal{R}}} = 1) \right|.$$

Distinguisher \mathcal{D} can query both its *left oracle* λ (either \mathcal{C} or \mathcal{R}) and its *right oracle* ρ (either \mathcal{P} or \mathcal{S}). We refer to $\mathcal{C}^{\mathcal{P}}, \mathcal{P}$ as the *real world*, and to $\mathcal{R}, \mathcal{S}^{\mathcal{R}}$ as the *simulated world*; the distinguisher \mathcal{D} converses with either of these worlds and its goal is to tell both worlds apart.

3.3 PRF-Security

For keyed hash functions, the indistinguishability framework suffices, and we use it to express the PRF-security of a keyed hash function. In this work, we will consider security in the multi-key setting. We adopt the model of Mouha and Luykx [ML15] to PRF-security. We refer to Bellare et al. [BBT16] for a more general discussion on multi-key security of PRFs. In below definition, μ denotes the number of instantiations with which the adversary interacts, and \mathcal{K} the key space.

Definition 2. Let $\mu \geq 1$, and let $\mathbf{k} \stackrel{\$}{\leftarrow} (\mathcal{K})^{\mu}$. Let \mathcal{C} be a keyed construction with key space \mathcal{K} and with oracle access to an ideal primitive \mathcal{P} . Let $\mathcal{R}_1, \dots, \mathcal{R}_{\mu}$ be random functions

with the same domains and ranges as $\mathcal{C}_{k_1}, \dots, \mathcal{C}_{k_\mu}$. Let \mathcal{D} be a distinguisher. The PRF distinguishing advantage of \mathcal{D} is defined as

$$\text{Prf}_{\mathcal{C}^{\mathcal{P}}}(\mathcal{D}) = \left| \mathbb{P} \left(\mathcal{D}^{\mathcal{C}_{k_1}^{\mathcal{P}}, \dots, \mathcal{C}_{k_\mu}^{\mathcal{P}}, \mathcal{P}} = 1 \right) - \mathbb{P} \left(\mathcal{D}^{\mathcal{R}_1, \dots, \mathcal{R}_\mu, \mathcal{P}} = 1 \right) \right|.$$

4 Indifferentiability of BLAKE2 Compression Function

We will prove that the BLAKE2 compression function is indifferentiable from a random compression function up to about $2^{n/2}$ queries, under the assumption that the underlying block cipher is randomly drawn from $\text{Block}^*(2n)$. The bound is in fact tight: an adaptation of the differentiability attack of [ALM12] from BLAKE to BLAKE2 does the job. We have included the attack in Appendix A for completeness.

Theorem 1 (Indifferentiability of BLAKE2 Compression Function). *Let $E \xleftarrow{\$} \text{Block}^*(2n)$ be a weakly ideal cipher, and consider the BLAKE2 compression function F^E of (2) that internally uses E . There exists a simulator \mathcal{S} such that for any distinguisher \mathcal{D} with total complexity q ,*

$$\text{Indiff}_{F^E, \mathcal{S}}(\mathcal{D}) \leq \frac{\binom{q}{2}}{2^{2n}} + \frac{\binom{q}{2}}{2^n} + \frac{q}{2^{n/2}},$$

where \mathcal{S} makes at most q queries to \mathcal{R} .

Note that one compression function evaluation corresponds to one block cipher evaluation, and vice versa, hence there is no need to separate the distinguisher's complexity into construction and primitive queries. The bound shows similarities with the analysis of the MD6 compression function [DRRS09, Theorem 1], but multiple differences appear at a technical level: most importantly, our analysis is in the weakly ideal cipher model.

The proof of Theorem 1 consists of two steps: in Section 4.1 we design the simulator used in the proof, and the derivation of the bound is given in Section 4.2.

4.1 Simulator

The simulator will simulate the interface of a block cipher $E \in \text{Block}^*(2n)$, but our simulator will generate most of its responses as if it were a random function: while this gives a small degradation in the security bound via the appearance of collisions, this significantly simplifies the description of the simulator and of the proof. Likewise, the simulator will not obey the \mathcal{S} -subspace invariance: the probability that a random value hits \mathcal{W} is $2^{4w}/2^{2n} = 1/2^{3n/2}$.

In more detail, our simulator will always generate uniformly random responses from $\{0, 1\}^{2n}$, with two exceptions:

- (i) The bijective \mathcal{W} -subspace invariance property for evaluations of the form (5) is retained. In other words, in a forward query $(m, x) \in \text{WK} \times \mathcal{W}$, the response y is randomly and bijectively drawn from \mathcal{W} , and similar for inverse queries;
- (ii) A forward query (m, x) , where x can be parsed into

$$h \parallel 0^{n/2} \parallel t \parallel f \leftarrow \text{parse}_{IV}(x),$$

is responded with a randomly generated y that satisfies $h' = \text{L}(y) \oplus \text{R}(y) \oplus h$, where $h' = \mathcal{R}(h, m, t, f)$.

Note that for exception (i), bijectivity on \mathcal{W} for keys from WK is strictly necessary due to the small size of \mathcal{W} ; otherwise, a distinguisher can find collision for the simulator in q queries with probability $\binom{q}{2}/|\mathcal{W}|$. The following brief lemma shows that exceptions (i) and (ii) cannot apply to the same query simultaneously.

Lemma 1. For any $x \in W$ of (3), $\text{parse}_{IV,z}(x) \neq 0^{n/2}$.

Proof. Note that, by definition of $\text{parse}_{IV,z}$, we have

$$\text{parse}_{IV,z}(x) = L(R(\text{parse}_{IV}(x))) = L(R(x)) \oplus L(IV).$$

As $x \in W$, $L(R(x)) = cccc$ for some $c \in \{0, 1\}^w$. On the other hand, the IV of BLAKE2 satisfies that $L(IV) \neq aaaa$ for any word $a \in \{0, 1\}^w$ (see the beginning of Section 2). Therefore, $\text{parse}_{IV,z}(x) \neq 0^{n/2}$. \square

The formal simulator is given in Figure 2. It maintains a table \mathcal{L} in which all query-response tuples (m, x, y) are stored. For convenience, we write $\mathcal{L}_m^+(x) = y$ and $\mathcal{L}_m^-(y) = x$. Furthermore, write $\text{dom}(\mathcal{L}_m) = \{x \mid (m, x, \cdot) \in \mathcal{L}\}$ and $\text{rng}(\mathcal{L}_m) = \{y \mid (m, \cdot, y) \in \mathcal{L}\}$.

Simulator Forward \mathcal{S}

Input: $(m, x) \in \{0, 1\}^{2n} \times \{0, 1\}^{2n}$
Output: $y \in \{0, 1\}^{2n}$

- 1: **if** $\mathcal{L}_m^+(x) = \perp$ **then**
- 2: $h \| z \| t \| f \leftarrow \text{parse}_{IV}(x)$
- 3: **if** $z = 0^{n/2}$ **then**
- 4: $L(y) \stackrel{\$}{\leftarrow} \{0, 1\}^n$
- 5: $h' \leftarrow \mathcal{R}(h, m, t, f)$
- 6: $\mathcal{L}_m^+(x) \leftarrow L(y) \| (L(y) \oplus h \oplus h')$
- 7: **else if** $(m, x) \in \text{WK} \times W$ **then**
- 8: $\mathcal{L}_m^+(x) \stackrel{\$}{\leftarrow} W \setminus \text{rng}(\mathcal{L}_m)$
- 9: **else**
- 10: $\mathcal{L}_m^+(x) \stackrel{\$}{\leftarrow} \{0, 1\}^{2n}$
- 11: **end if**
- 12: **end if**
- 13: **return** $\mathcal{L}_m^+(x)$

Simulator Inverse \mathcal{S}^{-1}

Input: $(m, y) \in \{0, 1\}^{2n} \times \{0, 1\}^{2n}$
Output: $x \in \{0, 1\}^{2n}$

- 1: **if** $\mathcal{L}_m^-(y) = \perp$ **then**
- 2: **if** $(m, y) \in \text{WK} \times W$ **then**
- 3: $\mathcal{L}_m^-(y) \stackrel{\$}{\leftarrow} W \setminus \text{dom}(\mathcal{L}_m)$
- 4: **else**
- 5: $\mathcal{L}_m^-(y) \stackrel{\$}{\leftarrow} \{0, 1\}^{2n}$
- 6: **end if**
- 7: **end if**
- 8: **return** $\mathcal{L}_m^-(y)$

Figure 2: Simulator \mathcal{S} for the proof of Theorem 1

4.2 Proof

Let $E \stackrel{\$}{\leftarrow} \text{Block}^*(2n)$ and F be the BLAKE2 compression function of Section 2.2. Let \mathcal{S} be the simulator of Figure 2, and let \mathcal{D} be any distinguisher that makes at most q oracle queries. Recall from Definition 1 that the distinguisher has access to either (F, E) or $(\mathcal{R}, \mathcal{S})$:

$$\text{Indiff}_{FE, \mathcal{S}}(\mathcal{D}) = \left| \mathbb{P}(\mathcal{D}^{F^E, E} = 1) - \mathbb{P}(\mathcal{D}^{\mathcal{R}, \mathcal{S}^{\mathcal{R}}} = 1) \right|. \quad (7)$$

As a first step, we apply a URP-URF switch to the real world: we replace E by a functionality \bar{E} that always generates its responses from $\{0, 1\}^{2n}$, *except for inputs from* $\text{WK} \times W$. By the triangle inequality, we find for (7):

$$\text{Indiff}_{FE, \mathcal{S}}(\mathcal{D}) \leq \left| \mathbb{P}(\mathcal{D}^{F^{\bar{E}}, \bar{E}} = 1) - \mathbb{P}(\mathcal{D}^{\mathcal{R}, \mathcal{S}^{\mathcal{R}}} = 1) \right| + \frac{\binom{q}{2}}{2^{2n}}, \quad (8)$$

and we focus on the success of \mathcal{D} in distinguishing $(F^{\bar{E}}, \bar{E})$ from $(\mathcal{R}, \mathcal{S}^{\mathcal{R}})$. We assume without loss of generality that the distinguisher never makes trivial queries, i.e., repeating a query to any of the oracles, querying ρ^{-1} on input of the response from ρ , or vice

versa, where $\rho \in \{\bar{E}, \mathcal{S}\}$. The oracles are written out in detail in Figure 3 for convenience. In the description, the functionality \bar{E} maintains an initially empty list \mathcal{L} as before. \mathcal{R} maintains an initially empty list \mathcal{M} that stores all query-response tuples (h, m, t, f, h') , and we write $\mathcal{M}^+(h, m, t, f) = h'$. For the sake of the proof, the function F also maintains an initially empty list Y_F of all responses given so far. The description of the simulator

<hr/> <p style="text-align: center;">Compression Function F</p> <hr/> <p>Input: $(h, m, t, f) \in \{0, 1\}^{n+2n+n/4+n/4}$ Output: $h' \in \{0, 1\}^n$</p> <ol style="list-style-type: none"> 1: $x \leftarrow h \ 0^{n/2} \ t \ f \oplus 0^n \ IV$ 2: $y \leftarrow \bar{E}(m, x)$ 3: $Y_F \leftarrow Y_F \cup \{y\}$ ▷ administrative 4: $h' \leftarrow L(y) \oplus R(y) \oplus h$ 5: return h' <hr/> <p style="text-align: center;">Ideal Cipher \bar{E}</p> <hr/> <p>Input: $(m, x) \in \{0, 1\}^{2n} \times \{0, 1\}^{2n}$ Output: $y \in \{0, 1\}^{2n}$</p> <ol style="list-style-type: none"> 1: if $\mathcal{L}_m^+(x) = \perp$ then 2: if $(m, x) \in \text{WK} \times \text{W}$ then 3: $\mathcal{L}_m^+(x) \stackrel{\\$}{\leftarrow} \text{W} \setminus \text{rng}(\mathcal{L}_m)$ 4: else 5: $\mathcal{L}_m^+(x) \stackrel{\\$}{\leftarrow} \{0, 1\}^{2n}$ 6: end if 7: end if 8: return $\mathcal{L}_m^+(x)$ <hr/> <p style="text-align: center;">Ideal Cipher Inverse \bar{E}^{-1}</p> <hr/> <p>Input: $(m, y) \in \{0, 1\}^{2n} \times \{0, 1\}^{2n}$ Output: $x \in \{0, 1\}^{2n}$</p> <ol style="list-style-type: none"> 1: if $y \in Y_F$ then 2: bad1 3: end if ▷ administrative 4: if $\mathcal{L}_m^-(y) = \perp$ then 5: if $(m, y) \in \text{WK} \times \text{W}$ then 6: $\mathcal{L}_m^-(y) \stackrel{\\$}{\leftarrow} \text{W} \setminus \text{dom}(\mathcal{L}_m)$ 7: else 8: $\mathcal{L}_m^-(y) \stackrel{\\$}{\leftarrow} \{0, 1\}^{2n}$ 9: end if 10: end if 11: return $\mathcal{L}_m^-(y)$ <hr/>	<hr/> <p style="text-align: center;">Random Function \mathcal{R}</p> <hr/> <p>Input: $(h, m, t, f) \in \{0, 1\}^{n+2n+n/4+n/4}$ Output: $h' \in \{0, 1\}^n$</p> <ol style="list-style-type: none"> 1: if $\mathcal{M}^+(h, m, t, f) = \perp$ then 2: $\mathcal{M}^+(h, m, t, f) \stackrel{\\$}{\leftarrow} \{0, 1\}^n$ 3: end if 4: return $\mathcal{M}^+(h, m, t, f)$ <hr/> <p style="text-align: center;">Simulator Forward \mathcal{S}</p> <hr/> <p>Input: $(m, x) \in \{0, 1\}^{2n} \times \{0, 1\}^{2n}$ Output: $y \in \{0, 1\}^{2n}$</p> <ol style="list-style-type: none"> 1: if $\mathcal{L}_m^+(x) = \perp$ then 2: $h \ z \ t \ f \leftarrow \text{parse}_{IV}(x)$ 3: if $z = 0^{n/2}$ then 4: $L(y) \stackrel{\\$}{\leftarrow} \{0, 1\}^n$ 5: $h' \leftarrow \mathcal{R}(h, m, t, f)$ 6: $\mathcal{L}_m^+(x) \leftarrow L(y) \ (L(y) \oplus h \oplus h')$ 7: else if $(m, x) \in \text{WK} \times \text{W}$ then 8: $\mathcal{L}_m^+(x) \stackrel{\\$}{\leftarrow} \text{W} \setminus \text{rng}(\mathcal{L}_m)$ 9: else 10: $\mathcal{L}_m^+(x) \stackrel{\\$}{\leftarrow} \{0, 1\}^{2n}$ 11: end if 12: end if 13: return $\mathcal{L}_m^+(x)$ <hr/> <p style="text-align: center;">Simulator Inverse \mathcal{S}^{-1}</p> <hr/> <p>Input: $(m, y) \in \{0, 1\}^{2n} \times \{0, 1\}^{2n}$ Output: $x \in \{0, 1\}^{2n}$</p> <ol style="list-style-type: none"> 1: if $\mathcal{L}_m^-(y) = \perp$ then 2: if $(m, y) \in \text{WK} \times \text{W}$ then 3: $\mathcal{L}_m^-(y) \stackrel{\\$}{\leftarrow} \text{W} \setminus \text{dom}(\mathcal{L}_m)$ 4: else 5: $\mathcal{L}_m^-(y) \stackrel{\\$}{\leftarrow} \{0, 1\}^{2n}$ 6: end if 7: if $\text{parse}_{IV,z}(\mathcal{L}_m^-(y)) = 0^{n/2}$ then 8: bad2 9: end if ▷ administrative 10: end if 11: return $\mathcal{L}_m^-(y)$ <hr/>
--	--

Figure 3: Real world (F, \bar{E}) (left) and simulated world $(\mathcal{R}, \mathcal{S})$ (right). The statements “administrative” do not influence the operation of the oracles and are purely included for administrative reasons. The description of \mathcal{S} is identical to that of Figure 2, \mathcal{S}^{-1} differs only in the addition of lines 7-9

inverse in Figure 3 is now equipped with two **bad**-events. Note that this adjustment is purely administrative and does not influence the procedures of \bar{E}^{-1} and \mathcal{S}^{-1} . We will prove that, as long as \bar{E}^{-1} does not set **bad1** and \mathcal{S}^{-1} does not set **bad2**, both oracles are indistinguishable. More formally, by the principles of game playing, we obtain for (8):

$$\begin{aligned} & \left| \mathbb{P} \left(\mathcal{D}^{F^{\bar{E}}, \bar{E}} = 1 \right) - \mathbb{P} \left(\mathcal{D}^{\mathcal{R}, \mathcal{S}^{\mathcal{R}}} = 1 \right) \right| \\ & \leq \left| \mathbb{P} \left(\mathcal{D}^{F^{\bar{E}}, \bar{E}} = 1 \mid \neg \mathbf{bad1} \right) - \mathbb{P} \left(\mathcal{D}^{\mathcal{R}, \mathcal{S}^{\mathcal{R}}} = 1 \mid \neg \mathbf{bad2} \right) \right| + \mathbb{P}(\mathbf{bad1}) + \mathbb{P}(\mathbf{bad2}) . \end{aligned} \quad (9)$$

bad1 is set if the distinguisher makes an inverse query $\bar{E}^{-1}(m, y)$ that has already been defined in an earlier compression function call. Event **bad2** captures the case where $\mathcal{S}^{-1}(m, y)$ satisfies $z = 0^{n/2}$ by accident, where $h \| z \| t \| f \leftarrow \text{parse}_{IV}(x)$. It is straightforward to see that $\mathbb{P}(\mathbf{bad1}) \leq \binom{q}{2}/2^n$ and that $\mathbb{P}(\mathbf{bad2}) \leq q/2^{n/2}$. In the remainder of the proof, we will show that

$$\left| \mathbb{P} \left(\mathcal{D}^{F^{\bar{E}}, \bar{E}} = 1 \mid \neg \mathbf{bad1} \right) - \mathbb{P} \left(\mathcal{D}^{\mathcal{R}, \mathcal{S}^{\mathcal{R}}} = 1 \mid \neg \mathbf{bad2} \right) \right| = 0, \quad (10)$$

and the proof follows from (7)-(10). To prove (10), we will consider any query made by the distinguisher, either to $\lambda \in \{F, \mathcal{R}\}$, $\rho \in \{\bar{E}, \mathcal{S}\}$, or $\rho^{-1} \in \{\bar{E}^{-1}, \mathcal{S}^{-1}\}$, and show that for every query the responses from the real or ideal world are identical. Without loss of generality, we assume that the distinguisher never makes a repeat query, i.e., to which it knows the answer in advance.

- **Query $h' \leftarrow \lambda(h, m, t, f)$.** Write $x = h \| 0^{n/2} \| t \| f \oplus 0^n \| IV$. We make the following case distinction:
 - $\mathcal{L}_m^+(x) = \perp$. In the real world, this means that $\bar{E}(m, x)$ has never been queried. We will thus have $y \stackrel{\$}{\leftarrow} \{0, 1\}^{2n}$, and hence, $h' = L(y) \oplus R(y) \oplus h \stackrel{\$}{\leftarrow} \{0, 1\}^n$. In the simulated world, the condition implies that \mathcal{S} has never been evaluated on (m, x) , and hence, it never queried \mathcal{R} on input of (h, m, t, f) . Consequently, the call to \mathcal{R} is new, and responded with $h' \stackrel{\$}{\leftarrow} \{0, 1\}^n$;
 - $\mathcal{L}_m^+(x) \neq \perp$. Denote $y = \mathcal{L}_m^+(x)$. In the real world, we necessarily have $h' = L(y) \oplus R(y) \oplus h$. In the simulated world, the tuple (m, x, y) must have been added to \mathcal{L} in a *forward* query to \mathcal{S} : if it were an inverse query, it would have set **bad2**. Thus, following the algorithm of \mathcal{S} , we have $L(y) \oplus R(y) = h \oplus h'$. Thus, the responses are identically distributed;
- **Query $y \leftarrow \rho(m, x)$.** Parse $h \| z \| t \| f \leftarrow \text{parse}_{IV}(x)$. We make the following case distinction:
 - $z \neq 0^{n/2}$. The response is distributed identically in both worlds: if $(m, x) \in \text{WK} \times \text{W}$ the response y is generated uniformly at random from W without replacement; otherwise it simply satisfies $y \stackrel{\$}{\leftarrow} \{0, 1\}^{2n}$;
 - $z = 0^{n/2}$ and $\mathcal{M}^+(h, m, t, f) = \perp$. The condition implies that \bar{E} has never been evaluated on (m, x) . Additionally, by Lemma 1, we necessarily have $x \notin \text{W}$. In the real world, the response thus satisfies $y \stackrel{\$}{\leftarrow} \{0, 1\}^{2n}$. In the simulated world, \mathcal{S} will generate $L(y) \stackrel{\$}{\leftarrow} \{0, 1\}^n$ and query $h' \leftarrow \mathcal{R}(h, m, t, f)$, which will also be uniformly randomly drawn. Concluding, the entire output $y \stackrel{\$}{\leftarrow} \{0, 1\}^{2n}$;
 - $z = 0^{n/2}$ and $\mathcal{M}^+(h, m, t, f) \neq \perp$. Denote $h' = \mathcal{M}^+(h, m, t, f)$. In the real world, an earlier construction query has already specified the call, and we necessarily have $L(y) \oplus R(y) \oplus h = h'$. In the simulated world, \mathcal{S} will define y as $L(y) \| (L(y) \oplus h \oplus h')$ for $L(y) \stackrel{\$}{\leftarrow} \{0, 1\}^n$. Thus, in both cases, $y \stackrel{\$}{\leftarrow} \{0, 1\}^{2n} \setminus \{\bar{y} \in \{0, 1\}^{2n} \mid L(\bar{y}) \oplus R(\bar{y}) \neq h \oplus h'\}$;

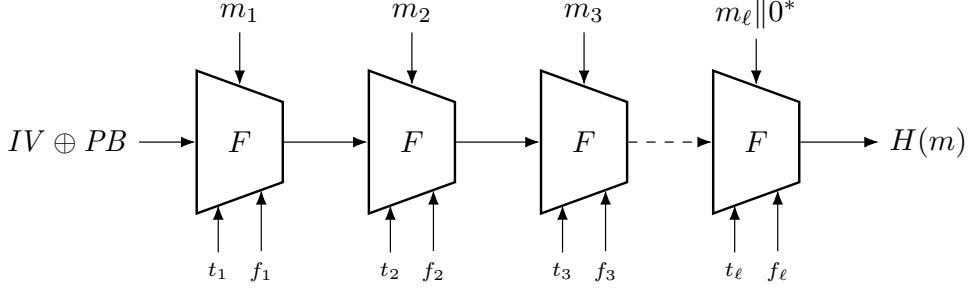


Figure 4: BLAKE2 hash function

- **Query** $x \leftarrow \rho^{-1}(m, y)$. If $\mathcal{L}_m^-(y) = \perp$, both oracles behave identically. Note that in the simulated world, this condition always holds (by our assumption that the distinguisher never makes trivial queries). For the real world, $\mathcal{L}_m^-(y) \neq \perp$ if and only if $y \in Y_F$, in which case the query would trigger **bad1**. \square

Remark 1. In the simulator of Figure 3, **bad2** is set if the response from $\mathcal{S}^{-1}(m, y)$ satisfies $z = 0^{n/2}$, where $h \| z \| t \| f \leftarrow \text{parse}_{IV}(x)$. However, BLAKE2 evaluates its compression function for at most 4 distinct flags f , rather than $2^{n/8}$. This means that it suffices to set **bad2** if $z = 0^{n/2}$ AND f is a valid flag. This happens with probability at most $q/2^{n/2+n/8-2}$. For the sake of generality, we have opted not to include this optimization in the proof.

5 BLAKE2 Hashing Modes

The BLAKE2 hashing mode differs from the one of BLAKE mostly in the use of a *parameter block* $PB \in \{0, 1\}^n$. Half of the parameter block, $n/2$ bits, consists of a salt and personalization data, both of which can be freely chosen by the user. The remaining half consists of mode-specific parameters (such as digest size, key size, tree parameters, etc.) and are merely determined by the mode.

In more detail, the BLAKE2 mode H gets as input a parameter block PB and a message m of size at most $2^{n/4}$ bytes. The message is padded into $m_1 \| \dots \| m_\ell \leftarrow m \| 0^*$ in such a way that $m_i \in \{0, 1\}^{2n}$ for $i = 1, \dots, \ell$. The HAIFA counter t_1, \dots, t_ℓ and the flags f_1, \dots, f_ℓ are set in such a way that $m \mapsto (m_1 \| t_1 \| f_1, \dots, m_\ell \| t_\ell \| f_\ell)$ is injective, suffix-free, and prefix-free. We refer to [ANWW13, SA15] for details regarding the flags and to [BD07] for details regarding the counter. The data is then hashed as follows:

$$h_0 \leftarrow IV \oplus PB \quad (11a)$$

$$\mathbf{for} \ i = 1, \dots, \ell \quad (11b)$$

$$h_i \leftarrow F(h_{i-1}, m_i, t_i, f_i) \quad (11c)$$

$$\mathbf{end\ for} \quad (11d)$$

$$\mathbf{return} \ h_\ell. \quad (11e)$$

The mode is depicted in Figure 4.

5.1 Security Analysis

Coron et al. [CDMP05] gave an indistinguishability analysis of prefix-free hash functions, based on the randomness of the underlying compression function. Improved bounds were obtained by Chang et al. [CLNY06] and Bhattacharyya et al. [BMN09]. These analyses

apply to the HAIFA structure, however, they assume a fixed IV of the state. For the BLAKE2 hashing mode, the initial state value is $IV \oplus PB$, where PB is a parameter block, which for a large part consists of data freely choosable by the user.

To simplify our analysis, we will henceforth simply consider $IV \oplus PB =: m_0$, with input to the hash function being (m_0, m) where m is as above. We henceforth relax our security games and simply consider a user that can freely choose (m_0, m) for every query. Bagheri et al. [BGKZ12] presented an indistinguishability analysis of sequential hashing with free IV (or, in our terminology, “free m_0 ”) if the first and last compression function are different from the remaining ones. We will use the result on sufficient conditions for tree hashing by Bertoni et al. [BDPV14]. Although the result focuses on trees, it is directly applicable to the sequential mode of BLAKE2 (as a sequential mode is a specific type of tree). We will state the result from [BDPV14] in generality.

Lemma 2 (Bertoni et al. [BDPV14]). *Let H be an ideal hash function, and consider a tree mode T^H that internally uses H . Assume that the tree mode is tree-decodable, message-complete, and final-node separable. There exists a simulator \mathcal{S} such that for any distinguisher \mathcal{D} with total complexity q ,*

$$\text{Indiff}_{T^H, \mathcal{S}}(\mathcal{D}) \leq \frac{\binom{q}{2}}{2^n},$$

where \mathcal{S} makes at most $O(q^3)$ queries to \mathcal{R} .

Here, the total complexity counts the number of evaluations of H induced by all queries by \mathcal{D} . The three conditions on the tree informally imply that every tree is uniquely parseable, that every message bit is compressed, and that the final call to H is domain-separated from the other calls to H . We refer to [BDPV14] for a formal discussion of these conditions.

Note that Lemma 2 can particularly be used in case H is a fixed-input-length compression function. As such, the BLAKE2 hash function (11) defines a tree that is tree-decodable, message-complete, and final-node separable. The final condition is particularly covered as the flag f_ℓ is distinct from $f_1, \dots, f_{\ell-1}$. From Theorem 1 and Lemma 2 we henceforth obtain the following result.

Corollary 1 (Indistinguishability of BLAKE2 Hashing Mode). *Let $E \stackrel{\$}{\leftarrow} \text{Block}^*(2n)$ be a weakly ideal cipher, and consider the BLAKE2 hash function H^E of (11) that internally uses E . There exists a simulator \mathcal{S} such that for any distinguisher \mathcal{D} with total complexity q ,*

$$\text{Indiff}_{H^E, \mathcal{S}}(\mathcal{D}) \leq \frac{\binom{q}{2}}{2^{2n}} + \frac{2\binom{q}{2}}{2^n} + \frac{q}{2^{n/2}},$$

where \mathcal{S} makes at most $O(q^3)$ queries to \mathcal{R} .

Here, the total complexity counts the number of evaluations of E induced by all queries by \mathcal{D} . Note that the combined simulator corresponds to the simulator for the compression function (Theorem 1) which interacts with that of the mode (Lemma 2) which queries \mathcal{R} , and hence it has complexity $O(q^3)$. This will be inherited in applications of BLAKE2 due to the composition result of Maurer et al. [MRH04].

5.2 Tree/Parallel Hashing Mode

The BLAKE2 specification [ANWW13] also supports tree hashing or parallel hashing, performed on top of the BLAKE2 hash function (11). These modes are designed along the methodology by Bertoni et al. [BDPV14]. They satisfy tree-decodability, message-completeness, and final-node separability by design,² and Lemma 2 directly applies. In

²Particularly, final-node separability is achieved as the finalization flags are defined in such a way that the flag to the final evaluation of H is distinct from the other flags.

more detail, if T^E denotes either the tree or parallel mode based on a weakly ideal cipher E , from Lemma 2 and Corollary 1, we can obtain the following result. We remark that if tree/parallel hashing is done in such a way that only one compression function call per node is performed, a direct application of Lemma 2 to the compression function result from Theorem 1 applies, and one can obtain a slightly better bound.

Corollary 2 (Indifferentiability of BLAKE2 Tree/Parallel Mode). *Let $E \stackrel{\$}{\leftarrow} \text{Block}^*(2n)$ be a weakly ideal cipher, and consider the BLAKE2 tree/parallel hash function T^E that internally uses E . There exists a simulator \mathcal{S} such that for any distinguisher \mathcal{D} with total complexity q ,*

$$\text{Indiff}_{T^E, \mathcal{S}}(\mathcal{D}) \leq \frac{\binom{q}{2}}{2^{2n}} + \frac{3\binom{q}{2}}{2^n} + \frac{q}{2^{n/2}},$$

where \mathcal{S} makes at most $O(q^3)$ queries to \mathcal{R} .

6 BLAKE2 Keyed Hashing Mode

BLAKE2 supports keyed hashing by simply prepending the key to the message:

$$KH_k(PB, m) = H(PB, k \| 0^{2n-\kappa} \| m), \quad (12)$$

where $\kappa \leq 2n$ denotes the key size. In other words, the key gets processed as other data, and the HAIFA counter and flags are designated to the key in a similar fashion as if they were for normal data blocks. As claimed by the designers [ANWW13], the usage of the HAIFA counter makes the need of a HMAC-like mode unnecessary.

6.1 Security Analysis

We first derive a generic PRF-security result for KH provided that H is a random oracle. Recall from Definition 2 that we consider multi-key security, where the distinguisher gets access to $\mu \geq 1$ independent instances.

Lemma 3. *Let $\mu \geq 1$, and let $\mathbf{k} \stackrel{\$}{\leftarrow} (\{0, 1\}^\kappa)^\mu$. Let H be an ideal hash function, and consider the keyed hashing mode KH^H of (12) that internally uses H . For any distinguisher \mathcal{D} with total complexity q ,*

$$\text{Prf}_{KH^H}(\mathcal{D}) \leq \frac{\mu q}{2^\kappa} + \frac{\binom{\mu}{2}}{2^\kappa}.$$

Proof. Let \mathcal{D} be any distinguisher that makes at most q oracle queries. Let G be an ideal hash function with the same domain and range as H . Starting from Definition 2:

$$\begin{aligned} \text{Prf}_{KH^H}(\mathcal{D}) &= \left| \mathbb{P}\left(\mathcal{D}^{KH_{k_1}^H, \dots, KH_{k_\mu}^H, H} = 1\right) - \mathbb{P}\left(\mathcal{D}^{\mathcal{R}_1, \dots, \mathcal{R}_\mu, H} = 1\right) \right| \\ &\leq \left| \mathbb{P}\left(\mathcal{D}^{KH_{k_1}^H, \dots, KH_{k_\mu}^H, H} = 1\right) - \mathbb{P}\left(\mathcal{D}^{KH_{k_1}^G, \dots, KH_{k_\mu}^G, H} = 1\right) \right| \quad (13) \\ &\quad + \left| \mathbb{P}\left(\mathcal{D}^{KH_{k_1}^G, \dots, KH_{k_\mu}^G, H} = 1\right) - \mathbb{P}\left(\mathcal{D}^{\mathcal{R}_1, \dots, \mathcal{R}_\mu, H} = 1\right) \right|. \quad (14) \end{aligned}$$

Distance (13) is bound by the event that the distinguisher queries H directly on one of the μ keys, which happens with probability at most $\mu q / 2^\kappa$. Distance (14) is bound by the event that two distinct keys k_i and k_j collide, which happens with probability at most $\binom{\mu}{2} / 2^\kappa$. \square

We immediately obtain the following from the hash function indifferentiability of Corollary 1 and from Lemma 3.

Corollary 3 (PRF-Security of BLAKE2 Keyed Hashing Mode). *Let $\mu \geq 1$, and let $\mathbf{k} \xleftarrow{\$} (\{0, 1\}^\kappa)^\mu$. Let $E \xleftarrow{\$} \text{Block}^*(2n)$ be a weakly ideal cipher, and consider the keyed hashing mode KH^E of (12) that internally uses H of (11) that internally uses E . For any distinguisher \mathcal{D} with total complexity q ,*

$$\text{Prf}_{KH^E}(\mathcal{D}) \leq \frac{\binom{q}{2}}{2^{2n}} + \frac{2\binom{q}{2}}{2^n} + \frac{q}{2^{n/2}} + \frac{\mu q}{2^\kappa} + \frac{\binom{\mu}{2}}{2^\kappa}.$$

We remark that Dinur and Leurent [DL14] presented a state recovery attack on a HAIFA MAC function, with complexity $2^{4n/5}$. In our model, however, we consider indistinguishability, a much weaker attack.

ACKNOWLEDGMENTS. This work was supported in part by the Research Council KU Leuven: GOA TENSE (GOA/11/007). Atul Luykx is supported by a Ph.D. Fellowship from the Institute for the Promotion of Innovation through Science and Technology in Flanders (IWT-Vlaanderen). Bart Mennink is a Postdoctoral Fellow of the Research Foundation – Flanders (FWO). The authors would like to thank the anonymous reviewers of FSE for their comments and suggestions.

References

- [AABS14] Leonardo C. Almeida, Ewerton R. Andrade, Paulo S. L. M. Barreto, and Marcos A. Simplicio Jr. Lyra: password-based key derivation with tunable memory and processing costs. *J. Cryptographic Engineering*, 4(2):75–89, 2014.
- [AGK⁺10a] Jean-Philippe Aumasson, Jian Guo, Simon Knellwolf, Krystian Matusiewicz, and Willi Meier. Differential and invertibility properties of BLAKE. In Hong and Iwata [HI10], pages 318–332.
- [AGK⁺10b] Jean-Philippe Aumasson, Jian Guo, Simon Knellwolf, Krystian Matusiewicz, and Willi Meier. Differential and invertibility properties of BLAKE (full version). Cryptology ePrint Archive, Report 2010/043, 2010.
- [ALM12] Elena Andreeva, Atul Luykx, and Bart Mennink. Provable security of BLAKE with non-ideal compression function. In Lars R. Knudsen and Huapeng Wu, editors, *Selected Areas in Cryptography, 19th International Conference, SAC 2012, Windsor, ON, Canada, August 15-16, 2012, Revised Selected Papers*, volume 7707 of *Lecture Notes in Computer Science*, pages 321–338. Springer, 2012.
- [AMP10] Elena Andreeva, Bart Mennink, and Bart Preneel. On the indifferentiability of the Grøstl hash function. In Juan A. Garay and Roberto De Prisco, editors, *Security and Cryptography for Networks, 7th International Conference, SCN 2010, Amalfi, Italy, September 13-15, 2010. Proceedings*, volume 6280 of *Lecture Notes in Computer Science*, pages 88–105. Springer, 2010.
- [ANWW13] Jean-Philippe Aumasson, Samuel Neves, Zooko Wilcox-O’Hearn, and Christian Winnerlein. BLAKE2: simpler, smaller, fast as MD5. In Michael J. Jacobson Jr., Michael E. Locasto, Payman Mohassel, and Reihaneh Safavi-Naini, editors, *Applied Cryptography and Network Security - 11th International Conference, ACNS 2013, Banff, AB, Canada, June 25-28, 2013. Proceedings*, volume 7954 of *Lecture Notes in Computer Science*, pages 119–135. Springer, 2013.

- [BBT16] Mihir Bellare, Daniel J. Bernstein, and Stefano Tessaro. Hash-function based PRFs: AMAC and its multi-user security. In Marc Fischlin and Jean-Sébastien Coron, editors, *Advances in Cryptology - EUROCRYPT 2016 - 35th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Vienna, Austria, May 8-12, 2016, Proceedings, Part I*, volume 9665 of *Lecture Notes in Computer Science*, pages 566–595. Springer, 2016.
- [BCC⁺09] Emmanuel Bresson, Anne Canteaut, Benoît Chevallier-Mames, Christophe Clavier, Thomas Fuhr, Aline Gouget, Thomas Icart, Jean-François Misarsky, Maria Naya-Plasencia, Pascal Paillier, Thomas Pornin, Jean-René Reinhard, Céline Thuillet, and Marion Videau. Indifferentiability with distinguishers: Why Shabal does not require ideal ciphers. *Cryptology ePrint Archive*, Report 2009/199, 2009.
- [BD07] Eli Biham and Orr Dunkelman. A framework for iterative hash functions - HAIFA. *Cryptology ePrint Archive*, Report 2007/278, 2007.
- [BDK16] Alex Biryukov, Daniel Dinu, and Dmitry Khovratovich. Argon2: New generation of memory-hard functions for password hashing and other applications. In *IEEE European Symposium on Security and Privacy, EuroS&P 2016, Saarbrücken, Germany, March 21-24, 2016*, pages 292–302. IEEE, 2016.
- [BDPV08] Guido Bertoni, Joan Daemen, Michaël Peeters, and Gilles Van Assche. On the indifferentiability of the sponge construction. In Nigel P. Smart, editor, *Advances in Cryptology - EUROCRYPT 2008, 27th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Istanbul, Turkey, April 13-17, 2008. Proceedings*, volume 4965 of *Lecture Notes in Computer Science*, pages 181–197. Springer, 2008.
- [BDPV14] Guido Bertoni, Joan Daemen, Michaël Peeters, and Gilles Van Assche. Sufficient conditions for sound tree and sequential hashing modes. *Int. J. Inf. Sec.*, 13(4):335–353, 2014.
- [BFL10] Charles Bouillaguet, Pierre-Alain Fouque, and Gaëtan Leurent. Security analysis of SIMD. In Alex Biryukov, Guang Gong, and Douglas R. Stinson, editors, *Selected Areas in Cryptography - 17th International Workshop, SAC 2010, Waterloo, Ontario, Canada, August 12-13, 2010, Revised Selected Papers*, volume 6544 of *Lecture Notes in Computer Science*, pages 351–368. Springer, 2010.
- [BGKZ12] Nasour Bagheri, Praveen Gauravaram, Lars R. Knudsen, and Erik Zenner. The suffix-free-prefix-free hash function construction and its indifferentiability security analysis. *Int. J. Inf. Sec.*, 11(6):419–434, 2012.
- [BHH⁺15] Daniel J. Bernstein, Daira Hopwood, Andreas Hülsing, Tanja Lange, Ruben Niederhagen, Louiza Papachristodoulou, Michael Schneider, Peter Schwabe, and Zooko Wilcox-O’Hearn. SPHINCS: practical stateless hash-based signatures. In Oswald and Fischlin [OF15], pages 368–397.
- [BKL⁺09] Mihir Bellare, Tadayoshi Kohno, Stefan Lucks, Niels Ferguson, Bruce Schneier, Doug Whiting, Jon Callas, and Jesse Walker. Provable security support for the Skein hash family. 2009.
- [BMN09] Rishiraj Bhattacharyya, Avradip Mandal, and Mridul Nandi. Indifferentiability characterization of hash functions and optimal bounds of popular domain extensions. In Bimal K. Roy and Nicolas Sendrier, editors, *Progress in*

- Cryptology - INDOCRYPT 2009, 10th International Conference on Cryptology in India, New Delhi, India, December 13-16, 2009. Proceedings*, volume 5922 of *Lecture Notes in Computer Science*, pages 199–218. Springer, 2009.
- [BMN10] Rishiraj Bhattacharyya, Avradip Mandal, and Mridul Nandi. Security analysis of the mode of JH hash function. In Hong and Iwata [HI10], pages 168–191.
- [BNR11] Alex Biryukov, Ivica Nikolić, and Arnab Roy. Boomerang attacks on BLAKE-32. In Antoine Joux, editor, *Fast Software Encryption - 18th International Workshop, FSE 2011, Lyngby, Denmark, February 13-16, 2011, Revised Selected Papers*, volume 6733 of *Lecture Notes in Computer Science*, pages 218–237. Springer, 2011.
- [BS01] Mike Boyle and Chris Salter. Dual counter mode, July 2001. <https://cryptome.org/dctr-spec.pdf>.
- [BS16] Raphael Bost and Olivier Sanders. Trick or tweak: On the (in)security of OTR’s tweaks. Cryptology ePrint Archive, Report 2016/234, 2016.
- [CDMP05] Jean-Sébastien Coron, Yevgeniy Dodis, Cécile Malinaud, and Prashant Puniya. Merkle-Damgård revisited: How to construct a hash function. In Victor Shoup, editor, *Advances in Cryptology - CRYPTO 2005: 25th Annual International Cryptology Conference, Santa Barbara, California, USA, August 14-18, 2005, Proceedings*, volume 3621 of *Lecture Notes in Computer Science*, pages 430–448. Springer, 2005.
- [CJMS14] Donghoon Chang, Arpan Jati, Sweta Mishra, and Somitra Kumar Sanadhya. Rig: A simple, secure and flexible design for password hashing. In Lin et al. [LYZ15], pages 361–381.
- [CLNY06] Donghoon Chang, Sangjin Lee, Mridul Nandi, and Moti Yung. Indifferentiable security analysis of popular hash functions with prefix-free padding. In Xuejia Lai and Kefei Chen, editors, *Advances in Cryptology - ASIACRYPT 2006, 12th International Conference on the Theory and Application of Cryptology and Information Security, Shanghai, China, December 3-7, 2006, Proceedings*, volume 4284 of *Lecture Notes in Computer Science*, pages 283–298. Springer, 2006.
- [CNY11] Donghoon Chang, Mridul Nandi, and Moti Yung. Indifferentiability of the hash algorithm BLAKE. Cryptology ePrint Archive, Report 2011/623, 2011.
- [CPB⁺12] Shu-jen Chang, Ray Perlner, William E. Burr, Meltem Sönmez Turan, John M. Kelsey, Souradyuti Paul, and Lawrence E. Bassham. Third-Round Report of the SHA-3 Cryptographic Hash Algorithm Competition. NISTIR 7896, National Institute for Standards and Technology, November 2012.
- [DFF⁺14] Jean Paul Degabriele, Victoria Fehr, Marc Fischlin, Tommaso Gagliardoni, Felix Günther, Giorgia Azzurra Marson, Arno Mittelbach, and Kenneth G. Paterson. Unpicking PLAID - A cryptographic analysis of an ISO-standards-track authentication protocol. In Liqun Chen and Chris J. Mitchell, editors, *Security Standardisation Research - First International Conference, SSR 2014, London, UK, December 16-17, 2014. Proceedings*, volume 8893 of *Lecture Notes in Computer Science*, pages 1–25. Springer, 2014.
- [DGW01] Pompiliu Donescu, Virgil D. Gligor, and David Wagner. A note on NSA’s dual counter mode of encryption, September 2001. <http://www.cs.berkeley.edu/~daw/papers/dcm-prelim.pdf>.

- [DK11] Orr Dunkelman and Dmitry Khovratovich. Iterative differentials, symmetries, and message modification in BLAKE-256. In *ECRYPT2 Hash Workshop*, 2011.
- [DL14] Itai Dinur and Gaëtan Leurent. Improved generic attacks against hash-based MACs and HAIFA. In Juan A. Garay and Rosario Gennaro, editors, *Advances in Cryptology - CRYPTO 2014 - 34th Annual Cryptology Conference, Santa Barbara, CA, USA, August 17-21, 2014, Proceedings, Part I*, volume 8616 of *Lecture Notes in Computer Science*, pages 149–168. Springer, 2014.
- [DRRS09] Yevgeniy Dodis, Leonid Reyzin, Ronald L. Rivest, and Emily Shen. Indifferentiability of permutation-based compression functions and tree-based modes of operation, with applications to MD6. In Orr Dunkelman, editor, *Fast Software Encryption, 16th International Workshop, FSE 2009, Leuven, Belgium, February 22-25, 2009, Revised Selected Papers*, volume 5665 of *Lecture Notes in Computer Science*, pages 104–121. Springer, 2009.
- [EFK15] Thomas Espitau, Pierre-Alain Fouque, and Pierre Karpman. Higher-order differential meet-in-the-middle preimage attacks on SHA-1 and BLAKE. In Gennaro and Robshaw [GR15], pages 683–701.
- [FLW14] Christian Forler, Stefan Lucks, and Jakob Wenzel. Memory-demanding password scrambling. In Palash Sarkar and Tetsu Iwata, editors, *Advances in Cryptology - ASIACRYPT 2014 - 20th International Conference on the Theory and Application of Cryptology and Information Security, Kaohsiung, Taiwan, R.O.C., December 7-11, 2014, Proceedings, Part II*, volume 8874 of *Lecture Notes in Computer Science*, pages 289–305. Springer, 2014.
- [GKN⁺14] Jian Guo, Pierre Karpman, Ivica Nikolić, Lei Wang, and Shuang Wu. Analysis of BLAKE2. In Josh Benaloh, editor, *Topics in Cryptology - CT-RSA 2014 - The Cryptographer's Track at the RSA Conference 2014, San Francisco, CA, USA, February 25-28, 2014. Proceedings*, volume 8366 of *Lecture Notes in Computer Science*, pages 402–423. Springer, 2014.
- [GR15] Rosario Gennaro and Matthew Robshaw, editors. *Advances in Cryptology - CRYPTO 2015 - 35th Annual Cryptology Conference, Santa Barbara, CA, USA, August 16-20, 2015, Proceedings, Part I*, volume 9215 of *Lecture Notes in Computer Science*. Springer, 2015.
- [Hao14] Yonglin Hao. The boomerang attacks on BLAKE and BLAKE2. In Lin et al. [LYZ15], pages 286–310.
- [HBHW16] Daira Hopwood, Sean Bowe, Taylor Hornby, and Nathan Wilcox. Zcash protocol specification, 2016. <https://github.com/zcash/zips/blob/master/protocol/protocol.pdf>.
- [HI10] Seokhie Hong and Tetsu Iwata, editors. *Fast Software Encryption, 17th International Workshop, FSE 2010, Seoul, Korea, February 7-10, 2010, Revised Selected Papers*, volume 6147 of *Lecture Notes in Computer Science*. Springer, 2010.
- [HKR15] Viet Tung Hoang, Ted Krovetz, and Phillip Rogaway. Robust authenticated-encryption AEZ and the problem that it solves. In Oswald and Fischlin [OF15], pages 15–44.

- [JAA⁺15] Marcos A. Simplicio Jr., Leonardo C. Almeida, Ewerton R. Andrade, Paulo C. F. dos Santos, and Paulo S. L. M. Barreto. Lyra2: Password hashing scheme with improved security against time-memory trade-offs. *Cryptology ePrint Archive*, Report 2015/136, 2015.
- [Jut01] Charanjit S. Jutla. Encryption modes with almost free message integrity. In Birgit Pfitzmann, editor, *Advances in Cryptology - EUROCRYPT 2001, International Conference on the Theory and Application of Cryptographic Techniques, Innsbruck, Austria, May 6-10, 2001, Proceeding*, volume 2045 of *Lecture Notes in Computer Science*, pages 529–544. Springer, 2001.
- [KLT15] Jonathan Katz, Stefan Lucks, and Aishwarya Thiruvengadam. Hash functions from defective ideal ciphers. In Kaisa Nyberg, editor, *Topics in Cryptology - CT-RSA 2015, The Cryptographer’s Track at the RSA Conference 2015, San Francisco, CA, USA, April 20-24, 2015. Proceedings*, volume 9048 of *Lecture Notes in Computer Science*, pages 273–290. Springer, 2015.
- [KNP⁺15] Dmitry Khovratovich, Ivica Nikolić, Josef Pieprzyk, Przemyslaw Sokolowski, and Ron Steinfeld. Rotational cryptanalysis of ARX revisited. In Gregor Leander, editor, *Fast Software Encryption - 22nd International Workshop, FSE 2015, Istanbul, Turkey, March 8-11, 2015, Revised Selected Papers*, volume 9054 of *Lecture Notes in Computer Science*, pages 519–536. Springer, 2015.
- [Lis06] Moses Liskov. Constructing an ideal hash function from weak ideal compression functions. In Eli Biham and Amr M. Youssef, editors, *Selected Areas in Cryptography, 13th International Workshop, SAC 2006, Montreal, Canada, August 17-18, 2006 Revised Selected Papers*, volume 4356 of *Lecture Notes in Computer Science*, pages 358–375. Springer, 2006.
- [LYZ15] Dongdai Lin, Moti Yung, and Jianying Zhou, editors. *Information Security and Cryptology - 10th International Conference, Inscrypt 2014, Beijing, China, December 13-15, 2014, Revised Selected Papers*, volume 8957 of *Lecture Notes in Computer Science*. Springer, 2015.
- [Men13] Bart Mennink. Indifferentiability of double length compression functions. In Martijn Stam, editor, *Cryptography and Coding - 14th IMA International Conference, IMACC 2013, Oxford, UK, December 17-19, 2013. Proceedings*, volume 8308 of *Lecture Notes in Computer Science*, pages 232–251. Springer, 2013.
- [Min14] Kazuhiko Minematsu. Parallelizable rate-1 authenticated encryption from pseudorandom functions. In Phong Q. Nguyen and Elisabeth Oswald, editors, *Advances in Cryptology - EUROCRYPT 2014 - 33rd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Copenhagen, Denmark, May 11-15, 2014. Proceedings*, volume 8441 of *Lecture Notes in Computer Science*, pages 275–292. Springer, 2014.
- [ML15] Nicky Mouha and Atul Luykx. Multi-key security: The Even-Mansour construction revisited. In Gennaro and Robshaw [GR15], pages 209–223.
- [MLMI13] Kazuhiko Minematsu, Stefan Lucks, Hiraku Morita, and Tetsu Iwata. Attacks and security proofs of EAX-Prime. In Shiho Moriai, editor, *Fast Software Encryption - 20th International Workshop, FSE 2013, Singapore, March 11-13, 2013. Revised Selected Papers*, volume 8424 of *Lecture Notes in Computer Science*, pages 327–347. Springer, 2013.

- [MP15] Bart Mennink and Bart Preneel. On the impact of known-key attacks on hash functions. In Tetsu Iwata and Jung Hee Cheon, editors, *Advances in Cryptology - ASIACRYPT 2015 - 21st International Conference on the Theory and Application of Cryptology and Information Security, Auckland, New Zealand, November 29 - December 3, 2015, Proceedings, Part II*, volume 9453 of *Lecture Notes in Computer Science*, pages 59–84. Springer, 2015.
- [MPS16] Dustin Moody, Souradyuti Paul, and Daniel Smith-Tone. Improved indistinguishability security bound for the JH mode. *Des. Codes Cryptography*, 79(2):237–259, 2016.
- [MRH04] Ueli M. Maurer, Renato Renner, and Clemens Holenstein. Indistinguishability, impossibility results on reductions, and applications to the random oracle methodology. In Moni Naor, editor, *Theory of Cryptography, First Theory of Cryptography Conference, TCC 2004, Cambridge, MA, USA, February 19-21, 2004, Proceedings*, volume 2951 of *Lecture Notes in Computer Science*, pages 21–39. Springer, 2004.
- [OF15] Elisabeth Oswald and Marc Fischlin, editors. *Advances in Cryptology - EUROCRYPT 2015 - 34th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Sofia, Bulgaria, April 26-30, 2015, Proceedings, Part I*, volume 9056 of *Lecture Notes in Computer Science*. Springer, 2015.
- [Per16] Trevor Perrin. The Noise protocol framework, 2016. <https://noiseprotocol.org/noise.html>.
- [Rog04] Phillip Rogaway. Efficient instantiations of tweakable blockciphers and refinements to modes OCB and PMAC. In Pil Joong Lee, editor, *Advances in Cryptology - ASIACRYPT 2004, 10th International Conference on the Theory and Application of Cryptology and Information Security, Jeju Island, Korea, December 5-9, 2004, Proceedings*, volume 3329 of *Lecture Notes in Computer Science*, pages 16–31. Springer, 2004.
- [Ros13] Alexander Roshal. RAR 5.0 archive format, 2013. <http://www.rarlab.com/technote.htm>.
- [RSS11] Thomas Ristenpart, Hovav Shacham, and Thomas Shrimpton. Careful with composition: Limitations of the indistinguishability framework. In Kenneth G. Paterson, editor, *Advances in Cryptology - EUROCRYPT 2011 - 30th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Tallinn, Estonia, May 15-19, 2011. Proceedings*, volume 6632 of *Lecture Notes in Computer Science*, pages 487–506. Springer, 2011.
- [SA15] Markku-Juhani Saarinen and Jean-Philippe Aumasson. The BLAKE2 cryptographic hash and message authentication code (MAC). Request for Comments (RFC) 7693, November 2015. <https://tools.ietf.org/html/rfc7693>.

A Differentiability Attack on the BLAKE2 Compression Function

We will derive a differentiability attack on the BLAKE2 compression function up to approximately $2^{n/2}$ queries. The proof is a fair translation of the differentiability attack on BLAKE by Andreeva et al. [ALM12] to BLAKE2. Note that we consider the compression

function based on an ideal cipher $E \stackrel{\$}{\leftarrow} \text{Block}(2n)$, rather than $E \stackrel{\$}{\leftarrow} \text{Block}^*(2n)$. This is without loss of generality. Note that, in below proof, the distinguisher may indeed opt to take the messages so that $m_j \notin \text{WK}$.

Theorem 2 (Differentiability of BLAKE2 Compression Function). *Let $E \stackrel{\$}{\leftarrow} \text{Block}(2n)$ be an ideal cipher, and consider the BLAKE2 compression function F^E of (2) that internally uses E . For any simulator \mathcal{S} that makes at most $q_{\mathcal{S}} \leq 2^{n-3}$ queries to \mathcal{R} , there exists a distinguisher \mathcal{D} that makes at most $2^{n/2} + 1$ queries to its oracles, such that*

$$\text{Indiff}_{F^E, \mathcal{S}}(\mathcal{D}) \geq 1 - e^{-1} - \frac{q_{\mathcal{S}}}{2^n} \geq 0.5.$$

Proof. Consider any simulator \mathcal{S} making at most $q_{\mathcal{S}}$ queries to the random oracle \mathcal{R} . We will construct a distinguisher \mathcal{D} that has access to either (F, E) and $(\mathcal{R}, \mathcal{S})$, and can distinguish those with significant probability. Denote its oracles by $(\lambda, \rho, \rho^{-1})$ (either (F, E, E^{-1}) or $(\mathcal{R}, \mathcal{S}, \mathcal{S}^{-1})$). \mathcal{D} operates as follows, where a return of 0 corresponds to guessing that it is talking to the real world (F, E) and a return 1 that it is talking to the simulated world $(\mathcal{R}, \mathcal{S})$:

1. \mathcal{D} selects $2^{n/2}$ distinct messages m_j , queries $x_j \leftarrow \rho^{-1}(m_j, 0)$, and parses

$$h_j \| z_j \| t_j \| f_j \leftarrow \text{parse}_{IV}(x_j);$$

2. If $z_j \neq 0^{n/2}$ for all $j \in \{1, \dots, 2^{n/2}\}$, then \mathcal{D} returns 1;
3. Let $j \in \{1, \dots, 2^{n/2}\}$ be such that $z_j = 0^{n/2}$. \mathcal{D} queries

$$h \leftarrow \lambda(h_j, m_j, t_j, f_j).$$

If $h = h_j$, \mathcal{D} returns 0, otherwise it returns 1.

The distinguisher guesses its oracles correctly *except* if one of the following events occur:

$$\begin{aligned} \mathbf{E}_1 : \quad & \forall j \in \{1, \dots, 2^{n/2}\} : z_j \neq 0^{n/2} & | \quad (\lambda, \rho) = (F, E); \\ \mathbf{E}_2 : \quad & \exists j \in \{1, \dots, 2^{n/2}\} : z_j = 0^{n/2} \text{ and } h = h_j & | \quad (\lambda, \rho) = (\mathcal{R}, \mathcal{S}). \end{aligned}$$

In particular, $\text{Indiff}_{F^E, \mathcal{S}}(\mathcal{D}) \geq 1 - \mathbb{P}(\mathbf{E}_1) - \mathbb{P}(\mathbf{E}_2)$.

Consider $\mathbb{P}(\mathbf{E}_1)$, and suppose that $(\lambda, \rho) = (F, E)$. Because E is an ideal cipher and the message blocks m_j are all pairwise distinct, we have

$$\begin{aligned} \mathbb{P}\left(\forall j \in \{1, \dots, 2^{n/2}\} : z_j \neq 0^{n/2}\right) &= \prod_{j=1}^{2^{n/2}} \mathbb{P}\left(z_j \neq 0^{n/2}\right) = \prod_{j=1}^{2^{n/2}} 1 - \mathbb{P}\left(z_j = 0^{n/2}\right) \\ &= \prod_{j=1}^{2^{n/2}} 1 - \frac{1}{2^{n/2}} = \left(1 - \frac{1}{2^{n/2}}\right)^{2^{n/2}} \leq e^{-1}. \end{aligned}$$

Next, consider $\mathbb{P}(\mathbf{E}_2)$, and suppose that $(\lambda, \rho) = (\mathcal{R}, \mathcal{S})$. The event implies that \mathcal{S} has generated an evaluation $\mathcal{R}(h_j, m_j, t_j, f_j) = h_j$, i.e., a fixed-point in the first n bits of the input to \mathcal{R} . As \mathcal{S} makes at most $q_{\mathcal{S}}$ queries, it can find such a fixed-point with probability at most $q_{\mathcal{S}}/2^n$.

We have obtained that $\text{Indiff}_{F^E, \mathcal{S}}(\mathcal{D}) \geq 1 - e^{-1} - q_{\mathcal{S}}/2^n \geq 0.5$ for $q_{\mathcal{S}} \leq 2^{n-3}$. \square