Simpira v2: A Family of Efficient Permutations Using the AES Round Function

Shay Gueron[1,2] and Nicky Mouha[3,4,5]

[1] Department of Mathematics, University of Haifa, Israel.
[2] Intel Corporation, Israel Development Center, Haifa, Israel.
[3] Dept. Electrical Engineering-ESAT/COSIC, KU Leuven, Leuven and iMinds, Ghent, Belgium.
[4] Project-team SECRET, Inria, France.
[5] National Institute of Standards and Technology, Gaithersburg, MD, USA.
shay@math.haifa.ac.il, nicky@mouha.be

Abstract. This paper introduces Simpira, a family of cryptographic permutations that supports inputs of $128 \times b$ bits, where b is a positive integer. Its design goal is to achieve high throughput on virtually all modern 64-bit processors, that nowadays already have native instructions for AES. To achieve this goal, Simpira uses only one building block: the AES round function. For $b = 1$, Simpira corresponds to 12-round AES with fixed round keys, whereas for $b \geq 2$, Simpira is a Generalized Feistel Structure (GFS) with an F-function that consists of two rounds of AES. We claim that there are no structural distinguishers for Simpira with a complexity below 2^{128}, and analyze its security against a variety of attacks in this setting. The throughput of Simpira is close to the theoretical optimum, namely, the number of AES rounds in the construction. For example, on the Intel Skylake processor, Simpira has throughput below 1 cycle per byte for $b \leq 4$ and $b = 6$. For larger permutations, where moving data in memory has a more pronounced effect, Simpira with $b = 32$ (512 byte inputs) evaluates 732 AES rounds, and performs at 824 cycles (1.61 cycles per byte), which is less than 13 % off the theoretical optimum. If the data is stored in interleaved buffers, this overhead is reduced to less than 1 %. The Simpira family offers an efficient solution when processing wide blocks, larger than 128 bits, is desired.

Keywords. Cryptographic permutation, AES-NI, Generalized Feistel Structure (GFS), Beyond Birthday-Bound (BBB) security, hash function, Lamport signature, wide-block encryption, Even-Mansour.

1 Introduction

The introduction of AES instructions by Intel (subsequently by AMD, and recently ARM) has changed the playing field for symmetric-key cryptography on modern processors, which lead to a significant reduction of the encryption overheads. The performance of these instructions has been steadily improving in every new generation of processors. By now, on the latest Intel Architecture

Codename Skylake, the AESENC instruction that computes one round of AES has latency of 4 cycles and throughput of 1 cycle. The improved AES performance trend can be expected to continue, with the increasing demand for fast encryption of more and more data.

To understand the impact of the AES instructions in practice, consider for example the way that Google Chrome browser connects to `https://google.com`. In this situation, Google is in a privileged position, as it controls both the client and the server side. To speed up connections, Chrome (the client) is configured to identify the processor's capabilities. If AES-NI are available, it would offer (to the server) to use AES-128-GCM for performing authenticated encryption during the TLS handshake. The high-end server would accept the proposed cipher suite, due to the high performance of AES-GCM on its side. This would capture any recent 64-bit PC, tablet, desktop, or even smartphone. On older processors, or architectures without AES instructions, Chrome resorts to proposing the ChaCha20-Poly1305 algorithm during the secure handshake negotiation.

An advantage of AES-GCM is that the message blocks can be processed independently for encryption. This allows pipelining of the AES round instructions, so that the observed performance is dominated by their throughput, and not by their latency [43, 44]. We note that even if a browser negotiates to use an inherently sequential mode such as CBC encryption, the web server can process multiple independent data buffers in parallel to achieve high throughput (see [43, 44]), and this technique is already used in the recent OpenSSL version 1.0.2. This performance gain by collecting multiple independent encryption tasks and pipelining their execution, is important for the design rationale of Simpira.

Setting. This paper should be understood in the following setting. We focus only on processors with AES instructions. Assuming that several independent data sources are available, we explore several symmetric-key cryptographic constructions with the goal of achieving a high throughput. Our reported benchmarks are performed on the latest Intel processor, namely Architecture Codename Skylake, but we expect to achieve similar performance on any processor that has AES instructions with throughput 1.

In particular, we focus here on applications where the 128-bit block size of AES is not sufficient, and support for a wider range of block sizes is desired. This includes various use cases such as permutation-based hashing and wide-block encryption, or just to easily achieve security beyond 2^{64} input blocks without resorting to (often inefficient) modes of operation with "beyond birthday-bound" security. For several concrete suggestions of applications, we refer to Sect. 7.

Admittedly, our decision to focus on only throughput may result in unoptimized performance in certain scenarios where the latency is critical. However, we point out that this is not only a property of Simpira, but also of AES itself, when it is implemented on common architectures with AES instructions. To achieve optimal performance on such architectures, AES needs to be used in a parallelizable mode of operation, or in a protocol that supports processing

independent inputs. Similarly, this is the case for Simpira as well. In fact, for 128-bit inputs, Simpira is the same as 12-round AES with fixed round keys.

Origin of the name. Simpira is named after a mythical animal of the Peruvian Amazon. According to the legend, one of its front legs has the form of a spiral that can be extended to cover the entire surface of the earth [26]. In a similar spirit, the Simpira family of permutations extends itself to a very wide range of input sizes. Alternatively, Simpira can be seen as an acronym for "SIMple Permutations based on the Instruction for a Round of AES."

Update. This paper proposes Simpira v2. Compared to Simpira v1, the Type-1.x GFS by Yanagihara and Iwata was found to have a problem (see Sect. 8), and is replaced by a new construction that performs the same number of AESENCs. We also updated the round constants (see Sect. 4). Although no attack is currently known on Simpira v2 with the old rotation constants, the new constants seem to strengthen Simpira without affecting its performance in our benchmarks. Unless otherwise specified, Simpira in this document is assumed to refer to Simpira v2.

2   Related Work

Block ciphers that support wide input blocks have been around for a long time. Some of the earliest designs are Bear and Lion [2], and Beast [61]. They are higher-level constructions, in the sense that they use hash functions and stream ciphers as underlying components.

Perhaps the first wide-block block cipher that is not a higher-level construction is the Hasty Pudding Cipher [74], which supports block sizes of any positive number of bits. Another early design is the Mercy block cipher that operates on 4096-bit blocks [27]. More recently, low-level constructions that can be scaled up to large input sizes are the SPONGENT [17, 18] permutations and the LowMC [1] block ciphers.

Our decision to use only the AES round function as a building block for Simpira means that some alternative constructions are not considered in this paper. Of particular interest are the EGFNs [7] used in Lilliput [6], the AESQ permutation of PAEQ [13], and Haraka[6] [55]. The security claims and benchmark targets of these designs are very different from those of Simpira. We only claim security up to 2^{128} blocks of input. However unlike Haraka, we consider all distinguishing attacks up to this bound. Also, we focus only on throughput, and not on latency. An interesting topic for future work is to design variants of these constructions with similar security claims, and to compare their security and implementation properties with Simpira.

---

[6] The first version of Haraka was vulnerable to an attack by Jérémy Jean [51] due to a bad choice of round constants. We therefore refer to the second version of Haraka, which prevents the attack.

3 Design Rationale of Simpira

AES [31] is a block cipher that operates on 128-bit blocks. It iterates the AES round function 10, 12 or 14 times, using round keys that are derived from a key of 128, 192 or 256 bits, respectively. On Intel (and AMD) processors, the AES round function is implemented by the `AESENC` instruction. It takes a 128-bit state and a 128-bit round key as inputs, and returns a 128-bit output that is the result of applying the `SubBytes`, `ShiftRows`, `MixColumns` and `AddRoundKey` operations. An algorithmic description of `AESENC` is given in Alg. 1 of Sect. 4, where we give the full specification of Simpira.

A cryptographic permutation can be obtained by setting the AES round keys to fixed, publicly-known values. It is a bad idea to set all round keys to zero. Such a permutation can easily be distinguished from random: if all input bytes are equal to each other, the AES rounds preserve this property. Such problems are avoided when round constants are introduced: this breaks the symmetry inside every round, as well as the symmetry between rounds. Several ciphers are vulnerable to attacks resulting from this property, such as the CAESAR candidate PAES [52,53] and the first version of Haraka [51]. The aforementioned design criterion, already present in Simpira v1, excludes the round constants of these designs.

We decided to use two rounds of AES in Simpira as the basic building block. As the `AESENC` instruction includes an XOR with a round key, this can be used to introduce a round constant in one AES round, and to do a "free XOR" in the other AES round. An added advantage is that two rounds of AES achieve *full bit diffusion*: every output bit depends on every input bit, and every input bit depends on every output bit.

Another design choice that we made, is to use only AES round functions in our construction, and no other operations. Our hope is that this design would maximize the contribution of every instruction to the security of the cryptographic permutation. It also simplifies the analysis and the implementation. From the performance viewpoint, the theoretically optimal software implementation would be able to dispatch a new `AESENC` instruction in every CPU clock cycle. A straightforward way to realize this design strategy is to use a (Generalized) Feistel Structure (GFS) for $b \geq 2$ that operates on b input subblocks of 128 bits each, as shown in Fig. 1.

As with any design, our goal is to obtain a good trade-off between security and efficiency. In order to explore a large design space, we use simple metrics to quickly estimate whether a given design reaches a sufficient level of security, and to determine its efficiency. In subsequent sections, we will formally introduce the designs, and study them in detail to verify the accuracy of our estimates.

3.1 Design Criteria

Our design criteria are as follows. The significance of both criteria against cryptanalysis attacks will be explained in Sect. 6.

Fig. 1. Two common classes of Generalized Feistel Structures (GFSs) are the Type-1 GFS (left) and the Type-2 GFS (right). For each example, two rounds are shown of a GFS that operates on $b = 6$ subblocks. We will initially consider these GFSs in this paper, as well as other GFSs with a different number of F-functions per round, and other *subblock shuffles* at the end of every round. At a later stage, we will consider more advanced constructions as well.

– **Security**: We calculate the number of Feistel rounds to achieve either *full bit diffusion*, as well as the number of Feistel rounds to achieve at least 25 (linearly or differentially) active S-boxes. To ensure a sufficient security margin against known attacks, we require that the number of rounds is three times the largest of these two numbers.
– **Efficiency**: As explained in Sect. 1, we will only focus on throughput. Given that we use no other operations besides the AES round function, we will use the number of AES round functions as an estimate for the total number of cycles.

Suzaki and Minematsu [75] formally defined DRmax to calculate how many Feistel rounds are needed for an input subblock to affect all the output subblocks. We will say that *full subblock diffusion* is achieved after DRmax rounds of the permutation or its inverse, whichever is greater. To achieve the strictly stronger criterion of *full bit diffusion*, one or two additional Feistel rounds may be required.

To obtain a lower bound for the minimum number of active S-boxes, we use a simplified representation that assigns one bit to every pair of bytes, to indicate whether or not they contain a non-zero difference (or linear mask). This allows us to use the Mixed-Integer Linear Programming (MILP) technique introduced by Mouha et al. [69] to quickly find a lower bound for the minimum number of active S-boxes.

3.2 Design Space Exploration

For each input size of the permutation, we explore a range of designs, and choose the one that maximizes the design criteria. If the search returns several alternatives, it does not really matter which one we choose. In that case, we arbitrarily choose the "simplest" design. The resulting Simpira design is shown in Fig. 2–3.

We will restrict ourselves to "simple" designs, such as for example constructions with identical round functions, instead of exhaustively searching for the optimal design that satisfies the design criteria. This is meant to simplify the cryptanalysis, as well as the implementation. We revisit this assumption in App. C.

Case $b = 1$. Full bit diffusion is reached after two rounds of AES, and four rounds of AES ensures at least 25 active S-boxes [31]. Following the design criteria, we select a design with 12 AES rounds.

Case $b = 2$. This is a (standard) Feistel structure. Full subblock diffusion is achieved after two Feistel rounds, and three Feistel rounds are needed to reach full bit diffusion. We find that five rounds ensures that there are at least 25 active S-boxes (see Fig. 5). Consequently, we select a design with 15 Feistel rounds.

Case $b = 3$. There are several designs that are optimal according to our criteria. They have either one or two F-functions per Feistel round, and various possibilities exist to reorder the subblocks at the end of every Feistel round. We choose what is arguably the simplest design: a Type-1 GFS according to Zheng et al.'s classification [82]. Full subblock diffusion requires five Feistel rounds, and at least six Feistel rounds are needed to ensure that there are least 25 active S-boxes. As seven Feistel rounds are needed to achieve full bit diffusion, we select a design with 21 Feistel rounds.

Case $b \geq 4$. The Type-1 GFS does not scale well for larger b, as diffusion becomes the limiting factor. More formally, Yanagihara and Iwata [78, 79] proved that the number of rounds required to reach full subblock diffusion is (at best) quadratic in the number of subblocks, regardless of how the subblocks are reordered at the end of every Feistel round.

In Simpira v1, the Yanagihara and Iwata's Type-1.x (b,2) GFS [80] was used for $b \geq 4$, except for $b = 6$ and $b = 8$. This is a design with two F-functions per round, where the number of rounds for full subblock diffusion is linear in b. Unfortunately, as we will explain in Sect. 8, this GFS is problematic as the same input subblock can be processed by more than one F-function. This general observation enabled attacks on Simpira v1 by Dobraunig et al. [35] and by Rønjom [73].

The Simpira v2 in this paper addresses this problem by ensuring that every subblock will enter an F-function only once. We do this by means of a new GFS construction. It uses $4b - 6$ F-functions to reach full bit diffusion, and ensures that at least 30 S-boxes are active (see also App. A). This construction is iterated three times, resulting in a design with $12b - 18$ F-functions, which is the same number of F-function as in Simpira v1.

We could also have used this construction for $b = 4$. However, we instead chose to go for a Type-2 GFS with 15 rounds. This not only results in a simpler

6

Fig. 2. One round of the Simpira construction for $b \in \{1, 2, 3, 4, 6, 8\}$. The total number of rounds is 6 for $b = 1$, 15 for $b = 2$, $b = 4$ and $b = 6$, 21 for $b = 3$, and 18 for $b = 8$. F is shorthand for $F_{c,b}$, where c is a counter that is initialized by one, and incremented after every evaluation of $F_{c,b}$. Every $F_{c,b}$ consists of two AES round evaluations, where the round constants that are derived from (c, b). The last round is special: the `MixColumns` is omitted when $b = 1$, and the final subblocks may be output in a different order. See Sect. 4 for a full specification.

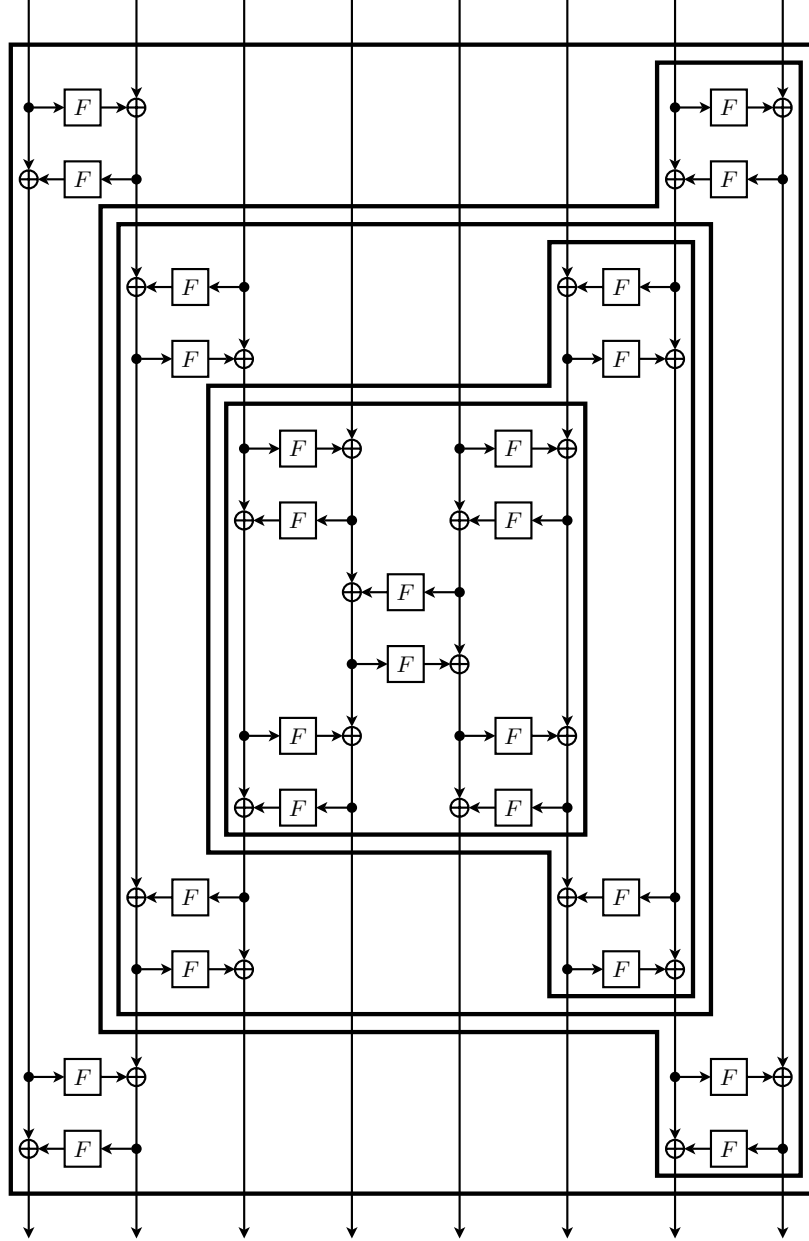**Fig. 3.** The Simpira construction for $b \notin \{1, 2, 3, 4, 6, 8\}$. F is shorthand for $F_{c,b}$, which consists of two rounds of AES as specified in Alg. 2. A generic construction is shown for all $b \geq 4$, however for $b \in \{4, 6, 8\}$ we will use the construction of Fig. 2. By convention, the leftmost F-function is from left to right; when this is not the case in the diagram, the direction of every F-function should be inverted. The full-round Simpira iterates the construction in this diagram three times. See Sect. 4 for a full specification.

construction, but also has the advantage ensuring at least 40 active S-boxes (instead of only 30) after five rounds.

But even if we had considered Yanagihara and Iwata's Type-1.x (b,2) GFS, we should also consider GFSs with more than two F-functions per Feistel round, which reach full subblock diffusion even quicker. However, this seems to come at the cost of using more F-functions in total. Looking only at the tabulated values of $\mathrm{DRmax}(\pi)$ and $\mathrm{DRmax}(\pi^{-1})$ in literature [75, 78–80], we can immediately rule out almost all alternative designs. Nevertheless, two improved Type-2 GFS designs by Suzaki and Minematsu [75] turned out to be superior. Instead of a cyclic left shift, they reorder the subblocks in a different way at the end of every Feistel round. We now explore these in detail.

Case $b = 6$. Let the *subblock shuffle* at the end of every Feistel round be presented by a list of indices that indicates which input subblock is mapped to which output subblock, e.g. $\{b - 1, 0, 1, 2, \ldots, b - 2\}$ denotes a cyclic left shift. Suzaki and Minematsu's improved Type-2 GFS with subblock shuffle $\{3, 0, 1, 4, 5, 2\}$ reaches full subblock diffusion and full bit diffusion after five Feistel rounds. At least 25 active S-boxes (in fact at least 30) are reached after four Feistel rounds. Following the design criteria, we end up with a design with 15 Feistel rounds. As this design has three F-functions in every Feistel round, it evaluates $3 \cdot 15 = 45$ F-functions. This is less than the general $b \geq 4$ case that requires $6b - 9$ Feistel rounds with 2 F-functions per round, which corresponds to $(6 \cdot 6 - 9) \cdot 2 = 54$ F-functions.

Case $b = 8$. Suzaki and Minematsu's improved Type-2 GFS with subblock shuffle $\{3, 0, 7, 4, 5, 6, 1, 2\}$ ensures both full subblock diffusion and full bit diffusion after six rounds. After four Feistel rounds, there are at least 25 active S-boxes (in fact at least 30). According to the design criteria, we end up with a design with 18 Feistel rounds, or $18 \cdot 4 = 72$ F-functions in total. The general $b \geq 4$ design would have required $(6b - 9) \cdot 2$ F-functions, which for $b = 8$ corresponds to $(6 \cdot 8 - 9) \cdot 2 = 78$ F-functions.

3.3 Design Alternatives

Until now, the only designs we discussed were GFS constructions where the F-function consists of two rounds of AES. We now take a step back, and briefly discuss alternative design choices.

As explained earlier, it is convenient to use two rounds of AES as a building block. It not only means that we reach full bit diffusion, but also that a "free XOR" is available to add a round constant on Intel and AMD architectures.

It is nevertheless possible to consider GFS designs with an F-function that consists of only one AES round. A consequence of this design choice is that extra XOR instructions will be needed to introduce round constants, which could increase the cycle count. But this design choice also complicates the analysis. For example when $b = 2$, we find that 25 Feistel rounds are then needed to ensure

at least 25 linearly active S-boxes. As shown in Fig. 4, this is because the tool can only ensure one active S-box for every Feistel round. Using two rounds of AES avoids this problem (see Fig. 5), and also significantly speeds up the tool: it makes bounding the minimum number of active S-boxes rather easy, instead of becoming increasingly complicated for a reasonably large value of b.


Fig. 4. A linear characteristic for an AES-based Feistel that uses only one round of AES inside its F-function. Crosshatches represent bytes with non-zero linear masks. The AES round consists of the `AddConstant` (AC), `SubBytes` (AC), `ShiftRows` (SR), and `MixColumns` (MC) operations. This round has only one active S-box. Therefore, 25 rounds are needed to ensure that there are least 25 linearly active S-boxes.

Likewise, we could also consider designs with more than two AES rounds per F-function. In our experiments, we have not found any cases where this results in a design where the total number of AES rounds is smaller. The intuition is as follows: the number of Feistel rounds to reach full subblock diffusion is independent of the F-function, therefore adding more AES rounds to every F function is not expected to result in a better trade-off.

If we take another step back, we might consider to use other instructions besides `AESENC`. Clearly, `AESDEC` can be used as an alternative, and the security properties and the benchmarks will remain the same. In fact, we use `AESDEC` when $b = 1$, to implement the inverse permutation. We do not use the `AESENCLAST` and `AESDECLAST` instructions, as they omit the `MixColumns` (resp. `InvMixColumns`) operation that is crucial to the *wide trail design strategy* [30] of AES. We do, however, use only one `AESENCLAST` for the very last round of the $b = 1$ permutation, as this makes an efficient implementation of the inverse permutation possible on Intel architectures. This is equivalent to applying a linear transformation to the output of the $b = 1$ permutation, therefore it does not reduce its cryptographic properties.

Fig. 5. A linear characteristic for one round of Simpira with $b = 2$ with 5 active S-boxes. Crosshatches represent bytes with non-zero linear masks. As Simpira uses two AES rounds per F-function, it can reach 25 active S-boxes in only 5 Feistel rounds, corresponding to 10 AES rounds in total.

Of course, it is possible to use non-AES instructions, possibly in combination with AES instructions. Actually, we do not need to be restricted to (generalized) Feistel designs for $b \geq 2$. However, such considerations are outside of the scope of this paper.

4   Specification of Simpira

An algorithmic specification of the Simpira design of Fig. 2–3 is given in Fig. 9–11. It uses one round of AES as a building block, which corresponds to the `AESENC` instruction on Intel processors (see Alg. 1). Its input is a 128-bit xmm register, which stores the AES 4×4 matrix of bytes as shown in Fig. 6. For additional details, we refer to [44].

s_0	s_4	s_8	s_{12}
s_1	s_5	s_9	s_{13}
s_2	s_6	s_{10}	s_{14}
s_3	s_7	s_{11}	s_{15}

Fig. 6. The internal state of AES can be represented by a 4×4 matrix of bytes, or as a 128-bit xmm register value $s = s_{15} \| \ldots \| s_0$, where s_0 is the least significant byte.

The F-function is specified in Alg. 2. It is parameterized by a counter c and by the number of subblocks b. Here, `SETR_EPI32` converts four 32-bit values into a 128-bit value, using the same byte ordering as the `_mm_setr_epi32()` compiler intrinsic. Fig. 7 shows how the constants can be expressed using the 4×4 byte matrix of AES.

$\texttt{0x00} \oplus c_0 \oplus b_0$	$\texttt{0x10} \oplus c_0 \oplus b_0$	$\texttt{0x20} \oplus c_0 \oplus b_0$	$\texttt{0x30} \oplus c_0 \oplus b_0$
$c_1 \oplus b_1$	$c_1 \oplus b_1$	$c_1 \oplus b_1$	$c_1 \oplus b_1$
$c_2 \oplus b_2$	$c_2 \oplus b_2$	$c_2 \oplus b_2$	$c_2 \oplus b_2$
$c_3 \oplus b_3$	$c_3 \oplus b_3$	$c_3 \oplus b_3$	$c_3 \oplus b_3$

Fig. 7. The constants used inside the $F_{c,b}$ function of Alg. 2, expressed as a 4×4 matrix of bytes. Here, $c = c_4 \| \ldots \| c_0$ and $b = b_4 \| \ldots \| b_0$ are 32-bit integers, where the least significant byte is c_0 and b_0 respectively.

Note that the constants have been updated in Simpira v2. The old constants of Simpira v1 are shown in Fig. 8. This update can be seen as "Grøstl strengthening," as it is inspired by the new round constants of the final-round Grøstl SHA-3 candidate [41]. No attack is currently known on Simpira v2 with the old rotation constants. Nevertheless, this change seems to strengthen Simpira without affecting its performance in our benchmarks.

c_0	b_0	0	0
c_1	b_1	0	0
c_2	b_2	0	0
c_3	b_3	0	0

Fig. 8. The old Simpira v1 constants used inside the $F_{c,b}$ function of Alg. 2, expressed as a 4×4 matrix of bytes. Again, $c = c_4 \| \ldots \| c_0$ and $b = b_4 \| \ldots \| b_0$ are 32-bit integers, where the least significant byte is x_0 and c_0 respectively.

Both the input and output of Simpira consist of b subblocks of 128 bits. The arrays use zero-based numbering, and array subscripts should be taken modulo the number of elements of the array. The subblock shuffle is done implicitly: we do not reorder the subblocks at the end of a Feistel round, but instead we apply the F-functions to other subblock inputs in the subsequent round. It is rather straightforward to implement the cyclic left shift in this way. For $b = 6$ and $b = 8$, the implementation of the subblock shuffle uses a decomposition into disjoint cycles.

As a result of this implementation choice, for $b \in \{2, 3, 4, 6, 8\}$, Simpira and its reduced-round variants are not always equivalent to a (generalized) Feistel

with identical rounds. For example, for $b = 2$ the F-function is alternatingly applied from left to right and from right to left. When the number of rounds is odd, this is not equivalent to a Feistel with identical rounds: the two output subblocks will be swapped.

When $b = 1$, an extra `InvMixColumns` operation is applied to the output. This is equivalent to omitting the `MixColumns` operation in the last round, and is required to efficiently implement the inverse Simpira permutation using Intel's AES instructions. For details on how to efficiently implement both Simpira and Simpira^{-1} when $b = 1$, see App. B.

The design strategy of Simpira is intended to be very conservative. Because we think that the security of Simpira with very large b may not yet be well understood, we recommend to use Simpira with $b \leq 65536$, corresponding to inputs of at most one megabyte. However, the external cryptanalysis of Simpira for any value of b is highly encouraged.

5 Benchmarks

We measured the performance of Simpira on the latest Intel processor, Architecture Codename Skylake. On this platform, the latency of `AESENC` is 4 cycles, and its throughput is 1 cycle. It follows that the software can be written in a way that fills the pipeline, by operating on four independent inputs. To obtain maximum throughput for all permutation sizes, we wrote functions that compute Simpira on four independent inputs. All Simpira permutations are benchmarked in the same setting, to make the results comparable.

Note that when $b = 4$, Simpira uses two independent F-functions, which means that maximum throughput could already be reached with only two independent inputs. For $b = 8$, where Simpira has four independent F-functions, even a single-stream Simpira implementation would fill the pipeline.

The measurements are performed as follows. We benchmark a function that evaluates Simpira for four independent inputs, and computed the number of cycles to carry out 256 calls to this function, as a "unit." This provides us with the throughput of Simpira. The results were obtained by using the `RDTSCP` instruction, 250 repetitions as a "warmup" phase, averaging the measurement on subsequent 1000 runs. Finally, this experiment was repeated 30 times, and the best result was selected. The platform was set up with Hyperthreading and Turbo Boost disabled.

The four data inputs can be stored sequentially at different pointers, or in an interleaved way (i.e. `A[0]B[0]C[0]D[0]A[1]B[1]C[1]D[1]...`). We benchmarked both settings. The results are shown in Table 1. We present only benchmarks for the forward Simpira permutation; the benchmarks for Simpira^{-1} turned out to be very similar.

We refer to App. D for a comparison with other constructions.

Algorithm 1 AESENC (see [44])

1: **procedure** AESENC(state, key)
2:     state ← SubBytes(state)
3:     state ← ShiftRows(state)
4:     state ← MixColumns(state)
5:     state ← state $\oplus$ key
6:     **return** state
7: **end procedure**

Algorithm 2 $F_{c,b}(x)$

1: **procedure** $F_{c,b}(x)$
2:     C ← SETR_EPI32(0x00 $\oplus$ c $\oplus$ b,
3:             0x10 $\oplus$ c $\oplus$ b,
4:             0x20 $\oplus$ c $\oplus$ b,
5:             0x30 $\oplus$ c $\oplus$ b)
6:     **return** AESENC(AESENC(x, C), 0)
7: **end procedure**

Algorithm 3 Simpira ($b = 1$)

1: **procedure** SIMPIRA(x_0)
2:     $R \leftarrow 6$
3:     **for** $c = 1, \ldots, R$ **do**
4:         $x_0 \leftarrow F_{c,b}(x_0)$
5:     **end for**
6:     InvMixColumns(x_0)
7:     **return** x_0
8: **end procedure**

Algorithm 4 Simpira^{-1} ($b = 1$)

1: **procedure** SIMPIRA(x_0)
2:     $R \leftarrow 6$
3:     MixColumns(x_0)
4:     **for** $c = R, \ldots, 1$ **do**
5:         $x_0 \leftarrow F_{c,b}^{-1}(x_0)$
6:     **end for**
7:     **return** x_0
8: **end procedure**

Algorithm 5 Simpira ($b \in \{2, 3, 4\}$)

1: **procedure** SIMPIRA($x_0, \ldots, x_{b-1}$)
2:     **if** $(b = 2) \vee (b = 3)$ **then**
3:         $R \leftarrow 6b + 3$
4:     **else**
5:         $R \leftarrow 15$
6:     **end if**
7:     $c \leftarrow 1$
8:
9:     **for** $r = 0, \ldots, R - 1$ **do**
10:         $x_{r+1} \leftarrow x_{r+1} \oplus F_{c,b}(x_r)$
11:         $c \leftarrow c + 1$
12:         **if** $b = 4$ **then**
13:             $x_{r+3} \leftarrow x_{r+3} \oplus F_{c,b}(x_{r+2})$
14:             $c \leftarrow c + 1$
15:         **end if**
16:     **end for**
17:     **return** $(x_0, x_1, \ldots, x_{b-1})$
18: **end procedure**

Algorithm 6 Simpira^{-1} ($b \in \{2, 3, 4\}$)

1: **procedure** SIMPIRA$^{-1}(x_0, \ldots, x_{b-1})$
2:     **if** $(b = 2) \vee (b = 3)$ **then**
3:         $R \leftarrow 6b + 3$
4:         $c \leftarrow R$
5:     **else**
6:         $R \leftarrow 15$
7:         $c \leftarrow 2R$
8:     **end if**
9:     **for** $r = R - 1, \ldots, 0$ **do**
10:         **if** $b = 4$ **then**
11:             $x_{r+3} \leftarrow x_{r+3} \oplus F_{c,b}(x_{r+2})$
12:             $c \leftarrow c - 1$
13:         **end if**
14:         $x_{r+1} \leftarrow x_{r+1} \oplus F_{c,b}(x_r)$
15:         $c \leftarrow c - 1$
16:     **end for**
17:     **return** $(x_0, x_1, \ldots, x_{b-1})$
18: **end procedure**

Fig. 9. Alg. 2 specifies $F_{c,b}$ using the AESENC operation that is defined in Alg. 1. Alg. 3–6 specify Simpira and its inverse for $b \leq 4$, where the input and output consist of b subblocks of 128 bits. Note that all arrays use zero-based numbering, and array subscripts should be taken modulo the number of elements of the array.

Algorithm 7 Simpira ($b = 6$)

1: **procedure** SIMPIRA($x_0, \ldots, x_5$)
2:      $R \leftarrow 15$
3:      $c \leftarrow 1$
4:      $s \leftarrow (0, 1, 2, 5, 4, 3)$
5:      **for** $r = 0, \ldots, R - 1$ **do**
6:         $x_{s_{r+1}} \leftarrow x_{s_{r+1}} \oplus F_{c,b}(x_{s_r})$
7:         $c \leftarrow c + 1$
8:         $x_{s_{r+5}} \leftarrow x_{s_{r+5}} \oplus F_{c,b}(x_{s_{r+2}})$
9:         $c \leftarrow c + 1$
10:        $x_{s_{r+3}} \leftarrow x_{s_{r+3}} \oplus F_{c,b}(x_{s_{r+4}})$
11:        $c \leftarrow c + 1$
12:      **end for**
13:      **return** $(x_0, x_1, \ldots, x_5)$
14: **end procedure**

Algorithm 8 Simpira^{-1} ($b = 6$)

1: **procedure** SIMPIRA^{-1}($x_0, \ldots, x_5$)
2:      $R \leftarrow 15$
3:      $c \leftarrow 45$
4:      $s \leftarrow (0, 1, 2, 5, 4, 3)$
5:      **for** $r = R - 1, \ldots, 0$ **do**
6:         $x_{s_{r+3}} \leftarrow x_{s_{r+3}} \oplus F_{c,b}(x_{s_{r+4}})$
7:         $c \leftarrow c - 1$
8:         $x_{s_{r+5}} \leftarrow x_{s_{r+5}} \oplus F_{c,b}(x_{s_{r+2}})$
9:         $c \leftarrow c - 1$
10:        $x_{s_{r+1}} \leftarrow x_{s_{r+1}} \oplus F_{c,b}(x_{s_r})$
11:        $c \leftarrow c - 1$
12:      **end for**
13:      **return** $(x_0, x_1, \ldots, x_5)$
14: **end procedure**

Algorithm 9 Simpira ($b = 8$)

1: **procedure** SIMPIRA($x_0, \ldots, x_7$)
2:      $R \leftarrow 18$
3:      $c \leftarrow 1$
4:      $s \leftarrow (0, 1, 6, 5, 4, 3)$
5:      $t \leftarrow (2, 7)$
6:      **for** $r = 0, \ldots, R - 1$ **do**
7:         $x_{s_{r+1}} \leftarrow x_{s_{r+1}} \oplus F_{c,b}(x_{s_r})$
8:         $c \leftarrow c + 1$
9:         $x_{s_{r+5}} \leftarrow x_{s_{r+5}} \oplus F_{c,b}(x_{t_r})$
10:        $c \leftarrow c + 1$
11:        $x_{s_{r+3}} \leftarrow x_{s_{r+3}} \oplus F_{c,b}(x_{s_{r+4}})$
12:        $c \leftarrow c + 1$
13:        $x_{t_{r+1}} \leftarrow x_{t_{r+1}} \oplus F_{c,b}(x_{s_{r+2}})$
14:        $c \leftarrow c + 1$
15:      **end for**
16:      **return** $(x_0, x_1, \ldots, x_7)$
17: **end procedure**

Algorithm 10 Simpira^{-1} ($b = 8$)

1: **procedure** SIMPIRA^{-1}($x_0, \ldots, x_7$)
2:      $R \leftarrow 18$
3:      $c \leftarrow 72$
4:      $s \leftarrow (0, 1, 6, 5, 4, 3)$
5:      $t \leftarrow (2, 7)$
6:      **for** $r = R - 1, \ldots, 0$ **do**
7:         $x_{t_{r+1}} \leftarrow x_{t_{r+1}} \oplus F_{c,b}(x_{s_{r+2}})$
8:         $c \leftarrow c - 1$
9:         $x_{s_{r+3}} \leftarrow x_{s_{r+3}} \oplus F_{c,b}(x_{s_{r+4}})$
10:        $c \leftarrow c - 1$
11:        $x_{s_{r+5}} \leftarrow x_{s_{r+5}} \oplus F_{c,b}(x_{t_r})$
12:        $c \leftarrow c - 1$
13:        $x_{s_{r+1}} \leftarrow x_{s_{r+1}} \oplus F_{c,b}(x_{s_r})$
14:        $c \leftarrow c - 1$
15:      **end for**
16:      **return** $(x_0, x_1, \ldots, x_7)$
17: **end procedure**

Fig. 10. Alg. 7–10 specify Simpira and its inverse for $b = 6$ and $b = 8$, using the $F_{c,b}$-function that is specified in Alg. 2. The input and the output consist of b subblocks of 128 bits. Note that all arrays use zero-based numbering, and array subscripts should be taken modulo the number of elements of the array.

Algorithm 11 Simpira
$(b \notin \{1, 2, 3, 4, 6, 8\})$

1: **procedure** SIMPIRA$(x_0, \ldots, x_{b-1})$
2:     $k \leftarrow 0$
3:     $d \leftarrow 2 \cdot \lfloor b/2 \rfloor$
4:     **for** $j = 1, \ldots, 3$ **do**
5:         **if** $d \neq b$ **then**
6:             TwoF$(b - 2, k)$
7:             $k \leftarrow k + 1$
8:         **end if**
9:         **for** $r = 0, \ldots, d - 2$ **do**
10:            TwoF(r, k)
11:            $k \leftarrow k + 1$
12:            **if** $r \neq d - r - 2$ **then**
13:                TwoF$(d - r - 2, k)$
14:                $k \leftarrow k + 1$
15:            **end if**
16:        **end for**
17:        **if** $d \neq b$ **then**
18:            TwoF$(b - 2, k)$
19:            $k \leftarrow k + 1$
20:        **end if**
21:    **end for**
22:    **return** $(x_0, x_1, \ldots, x_{b-1})$
23: **end procedure**

24: **procedure** TwoF(r, k)
25:    **if** $r \mod 2 = 0$ **then**
26:        $x_{r+1} \leftarrow x_{r+1} \oplus F_{2k+1,b}(x_r)$
27:        $x_r \leftarrow x_r \oplus F_{2k+2,b}(x_{r+1})$
28:    **else**
29:        $x_r \leftarrow x_r \oplus F_{2k+1,b}(x_{r+1})$
30:        $x_{r+1} \leftarrow x_{r+1} \oplus F_{2k+2,b}(x_r)$
31:    **end if**
32: **end procedure**

Algorithm 12 Simpira^{-1}
$(b \notin \{1, 2, 3, 4, 6, 8\})$

1: **procedure** SIMPIRA$^{-1}(x_0, \ldots, x_{b-1})$
2:     $k \leftarrow 6b - 10$
3:     $d \leftarrow 2 \cdot \lfloor b/2 \rfloor$
4:     **for** $j = 1, \ldots, 3$ **do**
5:         **if** $d \neq b$ **then**
6:             INVTwoF$(b - 2, k)$
7:             $k \leftarrow k - 1$
8:         **end if**
9:         **for** $r = d - 2, \ldots, 0$ **do**
10:            **if** $r \neq d - r - 2$ **then**
11:                INVTwoF$(d - r - 2, k)$
12:                $k \leftarrow k - 1$
13:            **end if**
14:            INVTwoF(r, k)
15:            $k \leftarrow k - 1$
16:        **end for**
17:        **if** $d \neq b$ **then**
18:            INVTwoF$(b - 2, k)$
19:            $k \leftarrow k - 1$
20:        **end if**
21:    **end for**
22:    **return** $(x_0, x_1, \ldots, x_{b-1})$
23: **end procedure**

24: **procedure** INVTwoF(r, k)
25:    **if** $r \mod 2 = 0$ **then**
26:        $x_r \leftarrow x_r \oplus F_{2k+2,b}(x_{r+1})$
27:        $x_{r+1} \leftarrow x_{r+1} \oplus F_{2k+1,b}(x_r)$
28:    **else**
29:        $x_{r+1} \leftarrow x_{r+1} \oplus F_{2k+2,b}(x_r)$
30:        $x_r \leftarrow x_r \oplus F_{2k+1,b}(x_{r+1})$
31:    **end if**
32: **end procedure**

Fig. 11. Alg. 11–12 specify Simpira and its inverse for $b \notin \{1, 2, 3, 4, 6, 8\}$, using the $F_{c,b}$-function that is specified in Alg. 2. Both the input and the output consist of b subblocks of 128 bits.

Table 1. Benchmarking results for the throughput of the Simpira permutations. For every b, we benchmark a function that applies the $128b$-bit permutation to four independent inputs. The data is either stored sequentially at different pointers, or in interleaved buffers. We give the number of cycles to process the four inputs, as well as the overhead compared the theoretical optimum of performing only `AESENC` instructions.

b	bits	# `AESENC`	non-interleaved		interleaved	
			cycles $(4\times)$	overhead	cycles $(4\times)$	overhead
1	128	12	50	$3\,\%$	50	$3\,\%$
2	256	30	122	$1\,\%$	122	$1\,\%$
3	384	42	171	$2\,\%$	171	$2\,\%$
4	512	60	241	$1\,\%$	241	$1\,\%$
6	768	90	362	$1\,\%$	362	$1\,\%$
8	1024	144	594	$3\,\%$	594	$3\,\%$
16	2048	348	1586	$14\,\%$	1400	$1\,\%$
32	4096	732	3295	$13\,\%$	2946	$1\,\%$
64	8192	1500	6791	$13\,\%$	6040	$1\,\%$
128	16384	3036	13942	$15\,\%$	12220	$1\,\%$
256	32768	6108	31444	$29\,\%$	24799	$2\,\%$

6 Cryptanalysis

The design criteria of Sect. 3 are not meant to be sufficient to guarantee security. In fact, it is not difficult to come up with trivially insecure constructions that satisfy (most of) the criteria. Rather, the design criteria are meant to assist us in identifying interesting constructions, which must then pass the scrutiny of cryptanalysis. Actually, during the design process of Simpira, we stumbled upon designs that were either insecure, or for which the security analysis was not so straightforward. When this happened, we adjusted the design criteria and repeated the search for constructions.

As such, we will not directly use the design criteria to argue the security of Simpira. Instead, we will use the fact that Simpira uses (generalized) Feistel structures and the AES round function, both of which have been extensively studied in literature. This allows us to focus our cryptanalysis efforts on the most promising attacks for this type of construction. We have tried to make this section easy to understand, which will hopefully convince the reader that Simpira should have a very comfortable security margin against all currently-known attacks.

Security claim. In what follows, we will only consider structural distinguishers [8] with a complexity up to 2^{128}. Simpira can be used in constructions that require a random permutation, however no statements can be made for adversaries that exceed 2^{128} queries. This type of security argument was first made by the SHA-3 [38] design team in response to high-complexity distinguishing

attacks on the underlying permutation [19–21], and has since been reused for other permutation-based designs.

Symmetry attacks. As explained in Sect. 3, the round constants are meant to avoid symmetry inside a Simpira round, as well as symmetry between rounds. The round constants also depend on b, which means that Simpira permutations of different widths should be indistinguishable from each other. The round constants are generated by a simple counter: this not only makes the design easy to understand and to implement, but also avoids any concerns that the constants may contain a backdoor. Every F-function has a different round constant: this does not seem to affect performance on recent Intel platforms, but greatly reduces the probability that a symmetry property can be maintained over several rounds.

Invariant subspace attacks. In its basic form, an invariant subspace attack [58] implies that there exists a coset of a vector space, so any number of iterations of the cryptographic round function maps to cosets of the same subspace. Rønjom [73] describes such an attack on Simpira v1 with $b = 4$, which is fixed in the current version. As explained in Sect. 9, no invariant subspace attacks were found for Simpira v2.

State collisions. For most block-cipher-based modes of operation, it is possible to define a "state," which is typically 128 bits long. This can be the chaining value for CBC mode, the counter for CTR mode, or the checksum in OCB. When a collision is found in this state, which is expected to happen around 2^{64} queries, the mode becomes insecure. For the Feistel-based Simpira ($b \geq 2$), there is no such concept of a "state." In fact: all subblocks receive roughly an equal amount of "processing." This allows Simpira to reach security beyond 2^{64} queries after a sufficient amount of Feistel rounds.

Linear and differential cryptanalysis. Simpira's security argument against linear [12] and differential [62] cryptanalysis (up to attacks with complexity 2^{128}) is the same as the argument for AES, which is based on counting the number of active S-boxes. As explained in [31], four rounds of AES have at least 25 (linearly or differentially) active S-boxes. Then any four-round differential characteristic holds with a probability less than $2^{-6 \cdot 25} = 2^{-150}$, and any four-round linear characteristic holds with a correlation less than $2^{-3 \cdot 25} = 2^{-75}$.

Here, 2^{-6} refers to the maximum difference propagation probability, and 2^{-3} is the maximum correlation amplitude of the S-box used in AES. The aforementioned reasoning makes the common assumptions that the probabilities of every round of a characteristic can be multiplied, and that this leads to a good estimate for the probability of the characteristic, and also of the corresponding differential.

18

The number of rounds typically needs to be slightly higher to account for partial key guesses (for keyed constructions), and to have a reasonable security margin. For any of the Simpira designs, we have at least three times the number of rounds required to reach 25 active S-boxes. This should give a sizable security margin against linear and differential cryptanalysis, and even against more advanced variants such as saturation and integral cryptanalysis [29]. In the case of integral cryptanalysis, of particular interest are the recently proposed integral distinguishers on Feistel and Generalized Feistel Networks by Todo [76] and by Zhang and Wenling [81].

Boomerang and differential-linear cryptanalysis. Instead of using one long characteristic, boomerang [77] and differential-linear [11, 57] cryptanalysis combine two shorter characteristics. But even combined with partial key guesses, the fact that Simpira has at least three times the number of rounds that result in 25 active S-boxes, should be more than sufficient to protect against this type of attacks.

Truncated and impossible differential cryptanalysis. When full bit diffusion is not reached, it is easy to construct a truncated differential [54] characteristic with probability one. A common way to construct an impossible differential [9,10] is the *miss in the middle* approach. It combines two probability-one truncated differentials, whose conditions cannot be met together.

However, every Simpira variant has at least three times the number of rounds to reach full bit diffusion. This should not only prevent truncated and impossible differential attacks, but result in a satisfactory security margin against such attacks.

Meet-in-the-middle and rebound attacks. Meet-in-the-middle-attacks [34] separate the equations that describe a symmetric-key primitive into two or three groups. This is done in such a way that some variables do not appear into at least one of these groups. A typical rebound attack [64] also splits a cipher into three parts: an inner part that is satisfied by meet-in-the-middle techniques (in the inbound phase), and two outer parts that are fulfilled in a probabilistic way (in the outbound phase).

With Simpira, splitting the construction in three parts will always result in one part that either has at least 25 active S-boxes, or that reaches full bit diffusion. This should not only prevent meet-in-the-middle and rebound attacks, but also provide a large security margin against these attacks.

On Simpira with $b = 1$ (corresponding to 12-round AES with fixed round keys), the best known distinguisher is a rebound attack by Gilbert and Peyrin [42] that attacks 8 rounds out of 12.

Generic attacks. A substantial amount of literature exists on generic attacks of Feistel structures. In particular, we are interested in attacks in Maurer et

19

al.'s indifferentiability setting [63], which is an extension of the indistinguishability notion for constructions that use publicly available oracles. In Simpira, the F-functions contain no secret key, and are therefore assumed to be publicly available.

Coron et al. [25] showed that five rounds of Feistel are not indifferentiable from a random permutation, and presented a indifferentiability proof for six rounds. Holenstein et al. [50] later showed that their proof is flawed, and provided a new indifferentiability proof for fourteen rounds. In very recent work, Dai and Steinberger [32] and independently Dachman-Soled et al. [28] announced an indifferentiability proof for the 10-round Feistel, which Dai and Steinberger subsequently improved to a proof for 8 rounds [33].

A problem with the aforementioned indifferentiability proofs is that they are rather weak: if the F-function is 128 bits wide, security is only proven up to about 2^{16} queries. The indistinguishability setting is better understood, where many proofs are available for not only Feistel, but also various generalized Feistel structures. But even in this setting, most proofs do not go beyond 2^{64} queries, and proving security with close to 2^{128} queries requires a very large number of rounds [49].

So although several of Simpira's Feistel-based permutations were proven to be indistinguishable from random permutations using [65,82], it is an open problem to prove stronger security bounds for Simpira and other generalized Feistel structures. Nevertheless, no generic attacks are known for Simpira, even when up to 2^{128} are made.

Note that strictly speaking, there is an exception to the previous sentence for Simpira with $b = 1$. It is guaranteed to be an even permutation [22, Thm. 4.8], and therefore $2^{128} - 1$ queries can distinguish it from a random permutation with advantage 0.5. We only mention this for completeness; actually all of Simpira's permutations can be shown to be even, but this is typically not considered to be more than just a mathematical curiosity.

Other attacks. We do not consider brute-force-like attacks [70], such as the biclique attacks on AES [16]: they perform exhaustive search on a smaller number of rounds, and therefore do not threaten the practical security of the cipher. However, it will be interesting to investigate such attacks in future work, as they give an indication of the security of the cipher in the absence of other attacks. We also do not look into algebraic attacks, as AES seems to very resistant against such attacks.

7   Applications

Simpira can be used in various scenarios where AES does not permit an efficient construction with security up to 2^{128} evaluations of the permutation. We present a brief overview possible applications.

A block cipher without round keys. The (single-key) Even-Mansour construction [37, 39, 40] uses a secret key K to turn a plaintext P into a ciphertext C as follows:

$$C = E_K(P) = \pi(P \oplus K) \oplus K \ , \tag{1}$$

where π is an n-bit permutation. As argued by Dunkelman et al. [37], the construction is minimal, in the sense that simplifying it, for example by removing one of its components, will render it completely insecure. Mouha and Luykx [67] showed that the Even-Mansour is in some sense optimal in the multi-key setting, where several keys are independently and uniformly drawn from the key space.

When D plaintext-ciphertexts are available, the secret key K of the Even-Mansour construction can be recovered in $2^n/D$ (off-line) evaluations of the permutation π [37]. This may be acceptable in lightweight authentication algorithms which rekey regularly, but may not be sufficient for encryption purposes [66, 67]. In order to achieve security up to about 2^{128} queries against all attacks in the multi-key setting, the Even-Mansour construction requires a permutation of at least 256 bits.

An important advantage of the Even-Mansour construction is that it avoids the need to precalculate round keys (and store them securely!) or to calculate them on the fly. Moreover, it also allows the easy construction of a tweakable block cipher. For a given tweak T, one can turn the Even-Mansour construction into a tweakable block cipher [59, 60]:

$$C = E_K(P) = \pi(P \oplus K \cdot T) \oplus K \cdot T \ , \tag{2}$$

that can be proven to be secure up to $2^{n/2}$ queries in the multi-key setting using the proof of [67, 68]. For concreteness, we use the multiplication $K \cdot T$ in $GF(2^n)$, which restricts the tweaks to $T \neq 0$. However, any ϵ-AXU hash function can be used instead of this multiplication [23].

If the cipher is computed in a parallelizable mode of operation, independent blocks can be pipelined, and the performance would be dominated by Simpira with the relevant value of b, plus the overhead of the key addition.

Permutation-based hashing. Achieving 128-bit collision resistance with a 128-bit permutation has been shown to be impossible [71]. Typically, a large permutation size is used to achieve a high throughput, for example 1600 bits in the sponge construction of SHA-3 [38]. The downside of using a large permutation is that performance is significantly reduced when many short messages need to be hashed, for example to compute a Lamport signature [56]. Simpira overcomes these problems by providing a family of efficient permutations with different input sizes.

In particular for hashing short messages, one may consider to use Simpira with a Davies-Meyer feed-forward: $\pi(x) \oplus x$. This construction has been shown to be optimally preimage and collision-resistant [14, 15], and even preimage aware [36], but not indifferentiable from a random oracle [24] as it is easy to find a fixed point: $\pi^{-1}(0)$. To match the intended application, padding of the input and/or truncation of the output of Simpira may be required.

Wide-block encryption and robust authenticated encryption. Wide-block encryption can be used to provide security against chosen ciphertext attacks when short (or even zero-length) authentication tags are used. In the context of full-disk encryption, there is usually no space to store an authentication tag. In an attempt to reduce the risk that ciphertext changes result in meaningful plaintext, a possibility is to use a wide block cipher to encrypt an entire disk sector, which typically has a size of 512 to 4096 bytes.

The same concern also exists when short authentication tags are used, and can be addressed by an encode-then-encipher approach [5]: add some bits of redundancy, and then encrypt with an arbitrary-input-length block cipher. Note that this technique achieves robust authenticated encryption [48].

Typical solutions for wide-block encryption such as the VIL [4], CMC [46] and EME [45, 47] modes of operation have the disadvantage that they are patented, and do not provide security beyond 2^{64} blocks of input. We are unaware of any patents related to Simpira.

When used in an Even-Mansour construction, Simpira with $b \geq 2$ can provide a wide block cipher that provides security up to 2^{128} blocks. When the block size exceeds the key size, the Even-Mansour construction can be generalized as follows:

$$C = E_K(P) = \pi(P \oplus (K \cdot T)\|0^*) \oplus (K \cdot T)\|0^*) \ , \tag{3}$$

where we set $T = 1$ if no tweak is provided. Note that this straightforward extension of the Even-Mansour construction appears in the proof for various sponge constructions. The first proof of security of this construction in the multikey setting was given by Andreeva et al. [3].

8 A Problem with Yanagihara and Iwata's GFS

For $b \geq 4$ (except $b = 6$ and $b = 8$), Simpira v1 used Yanagihara and Iwata's Type-1.x (b,2) GFS [80]. This is a GFS with two F-functions per round, shown in Fig. 12. Strictly speaking, we consider a variant of Yanagihara and Iwata's construction, that is identical up to a reordering of the input and output sub-blocks.

This construction has a problem. As can be seen from Fig. 12, the same value x_0 will eventually be processed by two F-functions. This clearly results in a redundant calculation, as the same F-function is evaluated twice on the same input.

In Simpira v1, the F-functions are not entirely identical due to the round constants. However, it can be seen that the problem in Yanagihara and Iwata's Type-1.x GFS also results in an attack on Simpira v1. In particular, the bounds on the number of active S-boxes were not correct, as the exact same S-box transitions were counted more than once. Dobraunig et al. [35] exploited the fact that the actual number of active S-boxes is much lower than expected, and constructed a series of attacks on the full 15-round Simpira v1 with $b = 4$, including a collision attack with complexity $2^{82.62}$ on Simpira when it is used in a truncated Davies-Meyer hash construction.

Fig. 12. Yanagihara and Iwata's Type-1.x (b,2) GFS [80], which was used in Simpira v1. Note that regardless of b, the same x_0 will eventually enter an F-function twice after a sufficient number of rounds.

The problem with Yanagihara and Iwata's construction was confirmed to us by its designers. It was pointed out to us that their Type-1.x GFS was implicitly assumed to use independent round keys, but that this assumption was unfortunately not mentioned in their paper [80].

When this assumption does not hold, the counts of active S-boxes can be incorrect. This occurs when various simple key schedules are used, such as for example the Even-Mansour construction. We avoid this problem in Simpira v2 by ensuring that the same input is never processed by more than one F-function. This can be seen to avoid attacks on GFS in block-cipher-based constructions, when used with a uniformly random key.

But Simpira is designed to be a family of cryptographic permutations, and should therefore also be secure in unkeyed settings. In the next section, we show how the unkeyed setting leads to invariant subspace attacks on Simpira v1 for $b = 4$.

9 Invariant Subspace Attacks

Leander et al. [58] introduced the term *invariant subspace attack*, which applies when there exists a (large) subspace, so that any coset of this subspace is mapped to itself when the round function is applied. We now explain such an attack applies to Yanagihara and Iwata Type-1.x (4,2) GFS. Again, strictly speaking Yanagihara and Iwata defined a variant of this construction, that is, however, identical up to a reordering of the input and output blocks. As illustrated in Fig. 13, we find that if the second and the last subblock of the input are identical, this property is preserved after any multiple of two rounds.

A similar observation also holds for Simpira v1 with $b = 4$, where only the round constants slightly destroy the symmetry property of the input. This is a consequence of the sparse round constants in Simpira v1, and the reuse of values into several F-functions, as explained in Sect. 8. In particular, for any even multiple of rounds up to 126, Simpira v1 round constants (see Fig. 8) only differ in the zeroth byte of the AES state. This means that if the second and the

Fig. 13. Yanagihara and Iwata Type-1.x (4,2) GFS [80], which was used in Simpira v1 for $b = 4$. We assume that all F-functions are identical. Here A, B and X can be any value. The leftmost input subblock enters the same F-function twice, and therefore guarantees that the output value Y will appear twice as well.

last subblock of the input are identical, this property will be preserved, except for the first column of corresponding AES states.

Rønjom [73] described an invariant subspace attack on Simpira v1 with $b = 4$. In particular, Rønjom identified a large subspace such that any coset of this space is invariant under two rounds of Simpira v2. This leads to a plaintext invariant over infinitely many even rounds. It can be seen as a generalization of the attack on Yanagihara-Iwata's Type-1.x GFS that is described in this section.

The property does not hold for an odd number of rounds, so it does not apply directly for Simpira v1 with $b = 4$, which consists of 15 Feistel rounds. For this reason, we did not detect any non-randomness in our test vectors, although it included the all-zero input that is an element of the coset of the invariant subspace. However, simply applying the permutation twice means that the total number of rounds is even, so that the distinguisher applies.

Do such invariant subspaces attacks also exist for Simpira v2? In an attempt to find such attacks, we first look for invariant subspaces when all F-functions are identical. This should give a good starting point to find invariant subspaces when the real (non-identical) F-functions of Simpira are used. More specifically, we select a random F-function, and consider four values for every input subblock: 0, $F(0)$, $F(F(0))$ and $F(F(0)) \oplus F(0)$. We then apply the Feistel round function several times, and use Gaussian elimination to check whether we stay within a particular linear subspace.

Using this technique, we found invariant subspaces for the GFS used in Simpira v2 when $b \in \{4, 6, 8\}$ (i.e. assuming identical F-functions), but not for other values of b. In fact, it can be seen that there is an invariant subspace for any

Type-2 GFS with an "even-odd shuffle [75]," that is, where even-numbered input subblocks are mapped to odd-numbered output subblocks and vice versa. For $b = 4$, such an invariant subspace is shown in Fig. 14. With the introduction of appropriate round constants, however, these invariant subspace attacks are avoided.


Fig. 14. The Type-2 GFS with $b = 4$, used in Simpira v2. We assume that all F-functions are identical. Here, A and B can be any value. If the odd-numbered input subblocks are equal, and the even-numbered input subblocks are equal, then this property is preserved for any number of rounds.

We chose to retain Type-2 GFS in Simpira v2 for $b \in \{4, 6, 8\}$, instead of replacing them by Generalized Feistel structures that "inherently" avoid invariant subspace attacks. This is because Type-2 GFS constructions are efficient and well-analyzed, and invariant subspaces can be avoided by using round constants.

We searched for invariant subspaces in all Simpira v2 variants, but were unable to find any. A similar search was also performed by Rønjom [72], who also could not identify invariant subspaces in the updated Simpira design. Unfortunately, currently no provable arguments against invariant subspace attacks are known. This is an interesting topic for future work.

10 Conclusion

We introduced Simpira, which is a family of cryptographic permutations that processes inputs of $128 \times b$ bits. It is intended to be a very conservative design that achieves high throughput on processors with AES instructions. We decided to use two rounds of AES as a building block, with the goal of simplifying

the design space exploration, and making the cryptanalysis and implementation straightforward.

With this building block, we explored a large number of generalized Feistel structures, and calculated how many rounds are required to reach either full bit diffusion, or 25 linearly or differentially active S-boxes, whichever is greater. To ensure a large security margin, we multiplied this number of rounds by three. Of all designs that we considered, we selected the ones with the lowest amount of F-functions in total.

Following these design criteria, Simpira resulted in seven different designs. For $b = 1$, we have AES with fixed round keys. Simpira uses a Feistel structure for $b = 2$, a Type-1 GFS for $b = 3$, and a Type-2 GFS for $b = 4$. The $b \geq 5$ design is a dedicated construction that we introduce in this paper. For $b = 6$ and $b = 8$, we use Suzaki and Minematsu's improved Type-2 GFS, as it has fewer F-functions than general construction for $b \geq 5$.

Our benchmarks on Intel Skylake showed that Simpira is close to the theoretical optimum of only executing `AESENC` instructions. For $b \leq 4$, Simpira is less than $3\,\%$ away from this optimum. For $b \leq 32$, corresponding to inputs of up to 512 bytes, Simpira is less than $13\,\%$ away from this optimum for a non-interleaved implementation, and less than $1\,\%$ away for an interleaved implementation.

It is unfortunate that many methods to encrypt wide input blocks, such as VIL, CMC, and EME, have not seen widespread adoption. The main obstacle appears to be that they are patented. We hope that Simpira can provide an interesting alternative: it is not only free from patent concerns, but offers security way beyond the 2^{64} limit for typical AES-based modes.

References

1. Albrecht, M.R., Rechberger, C., Schneider, T., Tiessen, T., Zohner, M.: Ciphers for MPC and FHE. In: EUROCRYPT 2015. LNCS, vol. 9056, pp. 430–454. Springer (2015)
2. Anderson, R.J., Biham, E.: Two Practical and Provably Secure Block Ciphers: BEARS and LION. In: FSE 1996. LNCS, vol. 1039, pp. 113–120. Springer (1996)
3. Andreeva, E., Daemen, J., Mennink, B., Assche, G.V.: Security of Keyed Sponge Constructions Using a Modular Proof Approach. In: FSE 2015. LNCS, vol. 9054, pp. 364–384. Springer (2015)
4. Bellare, M., Rogaway, P.: On the Construction of Variable-Input-Length Ciphers. In: FSE 1999. LNCS, vol. 1039, pp. 231–244. Springer (1999)
5. Bellare, M., Rogaway, P.: Encode-Then-Encipher Encryption: How to Exploit Nonces or Redundancy in Plaintexts for Efficient Cryptography. In: ASIACRYPT 2000. LNCS, vol. 1976, pp. 317–330. Springer (2000)
6. Berger, T.P., Francq, J., Minier, M., Thomas, G.: Extended Generalized Feistel Networks Using Matrix Representation to Propose a New Lightweight Block Cipher: Lilliput. IEEE Trans. Computers 65(7), 2074–2089 (2016)
7. Berger, T.P., Minier, M., Thomas, G.: Extended Generalized Feistel Networks Using Matrix Representation. In: SAC 2013. LNCS, vol. 8282, pp. 289–305. Springer (2013)
8. Bertoni, G., Daemen, J., Peeters, M., Van Assche, G.: Cryptographic Sponge Functions, available at `http://sponge.noekeon.org/CSF-0.1.pdf`
9. Biham, E., Biryukov, A., Shamir, A.: Cryptanalysis of Skipjack Reduced to 31 Rounds Using Impossible Differentials. In: EUROCRYPT 1999. LNCS, vol. 1592, pp. 12–23. Springer (1999)
10. Biham, E., Biryukov, A., Shamir, A.: Cryptanalysis of Skipjack Reduced to 31 Rounds Using Impossible Differentials. J. Cryptology 18(4), 291–311 (2005)
11. Biham, E., Dunkelman, O., Keller, N.: Enhancing Differential-Linear Cryptanalysis. In: ASIACRYPT 2002. LNCS, vol. 2501, pp. 254–266. Springer (2002)
12. Biham, E., Shamir, A.: Differential Cryptanalysis of DES-like Cryptosystems. J. Cryptology 4(1), 3–72 (1991)
13. Biryukov, A., Khovratovich, D.: PAEQ: Parallelizable Permutation-Based Authenticated Encryption. In: ISC 2014. LNCS, vol. 8783, pp. 72–89. Springer (2014)
14. Black, J., Rogaway, P., Shrimpton, T.: Black-Box Analysis of the Block-Cipher-Based Hash-Function Constructions from PGV. In: CRYPTO 2002. LNCS, vol. 2442, pp. 320–335. Springer (2002)
15. Black, J., Rogaway, P., Shrimpton, T., Stam, M.: An Analysis of the Blockcipher-Based Hash Functions from PGV. J. Cryptology 23(4), 519–545 (2010)
16. Bogdanov, A., Khovratovich, D., Rechberger, C.: Biclique Cryptanalysis of the Full AES. In: ASIACRYPT 2011. LNCS, vol. 7073. Springer (2011)
17. Bogdanov, A., Knežević, M., Leander, G., Toz, D., Varıcı, K., Verbauwhede, I.: SPONGENT: A Lightweight Hash Function. In: CHES 2011. LNCS, vol. 6917, pp. 312–325. Springer (2011)
18. Bogdanov, A., Knežević, M., Leander, G., Toz, D., Varıcı, K., Verbauwhede, I.: SPONGENT: The Design Space of Lightweight Cryptographic Hashing. IEEE Trans. Computers 62(10), 2041–2053 (2013)
19. Boura, C., Canteaut, A.: A zero-sum property for the KECCAK-f permutation with 18 rounds. In: ISIT 2010. pp. 2488–2492. IEEE (2010)

20. Boura, C., Canteaut, A.: Zero-Sum Distinguishers for Iterated Permutations and Application to KECCAK-f and Hamsi-256. In: SAC 2010. LNCS, vol. 6544, pp. 1–17. Springer (2011)
21. Boura, C., Canteaut, A., De Cannière, C.: Higher-order differential properties of Keccak and Luffa. Cryptology ePrint Archive, Report 2010/589 (2010)
22. Cid, C., Murphy, S., Robshaw, M.J.B.: Algebraic Aspects of the Advanced Encryption Standard. Springer (2006)
23. Cogliati, B., Lampe, R., Seurin, Y.: Tweaking Even-Mansour Ciphers. In: CRYPTO 2015. LNCS, vol. 9215, pp. 189–208. Springer (2015)
24. Coron, J., Dodis, Y., Malinaud, C., Puniya, P.: Merkle-Damgård Revisited: How to Construct a Hash Function. In: CRYPTO 2005. LNCS, vol. 3621, pp. 430–448. Springer (2005)
25. Coron, J., Patarin, J., Seurin, Y.: The Random Oracle Model and the Ideal Cipher Model Are Equivalent. In: CRYPTO 2008. LNCS, vol. 5157, pp. 1–20. Springer (2008)
26. Cossíos, D.: Breve Bestiario Peruano. Editorial Casatomada, second edn. (2008)
27. Crowley, P.: Mercy: A Fast Large Block Cipher for Disk Sector Encryption. In: FSE 2000. LNCS, vol. 1978, pp. 49–63. Springer (2000)
28. Dachman-Soled, D., Katz, J., Thiruvengadam, A.: 10-Round Feistel is Indifferentiable from an Ideal Cipher. In: EUROCRYPT 2016. LNCS, vol. 9666, pp. 649–678. Springer (2016)
29. Daemen, J., Knudsen, L.R., Rijmen, V.: The Block Cipher Square. In: FSE 1997. LNCS, vol. 1267, pp. 149–165. Springer (1997)
30. Daemen, J., Rijmen, V.: The Wide Trail Design Strategy. In: IMA 2001. LNCS, vol. 2260, pp. 222–238. Springer (2001)
31. Daemen, J., Rijmen, V.: The Design of Rijndael: AES - The Advanced Encryption Standard. Springer (2002)
32. Dai, Y., Steinberger, J.: Indifferentiability of 10-Round Feistel Networks. Cryptology ePrint Archive, Report 2015/874 (2015)
33. Dai, Y., Steinberger, J.P.: Indifferentiability of 8-Round Feistel Networks. In: CRYPTO 2016. LNCS, vol. 9814, pp. 95–120. Springer (2016)
34. Diffie, W., Hellman, M.E.: Exhaustive Cryptanalysis of the NBS Data Encryption Standard. Computer 10(6), 74–84 (1977)
35. Dobraunig, C., Eichlseder, M., Mendel, F.: Cryptanalysis of Simpira. Cryptology ePrint Archive, Report 2016/244 (2016)
36. Dodis, Y., Ristenpart, T., Shrimpton, T.: Salvaging Merkle-Damgård for Practical Applications. In: EUROCRYPT 2009. LNCS, vol. 5479, pp. 371–388. Springer (2009)
37. Dunkelman, O., Keller, N., Shamir, A.: Minimalism in Cryptography: The Even-Mansour Scheme Revisited. In: EUROCRYPT 2012. LNCS, vol. 7237, pp. 336–354. Springer (2012)
38. Dworkin, M.J.: SHA-3 Standard: Permutation-Based Hash and Extendable-Output Functions. Federal Inf. Process. Stds. (NIST FIPS) - 202 (August 2015)
39. Even, S., Mansour, Y.: A Construction of a Cipher From a Single Pseudorandom Permutation. In: ASIACRYPT 1991. LNCS, vol. 739. Springer (1993)
40. Even, S., Mansour, Y.: A Construction of a Cipher from a Single Pseudorandom Permutation. J. Cryptology 10(3), 151–162 (1997)
41. Gauravaram, P., Knudsen, L.R., Matusiewicz, K., Mendel, F., Rechberger, C., Schläffer, M., Thomsen, S.S.: Grøstl – a SHA-3 candidate. Submission to the NIST SHA-3 Competition (Round 3) (2011), `http://www.groestl.info/Groestl.pdf`

42. Gilbert, H., Peyrin, T.: Super-Sbox Cryptanalysis: Improved Attacks for AES-Like Permutations. In: FSE 2010. LNCS, vol. 6147, pp. 365–383. Springer (2010)

43. Gueron, S.: Intel's New AES Instructions for Enhanced Performance and Security. In: FSE 2009. LNCS, vol. 5665, pp. 51–66. Springer (2009)

44. Gueron, S.: Intel® Advanced Encryption Standard (AES) New Instructions Set. Available at: `https://software.intel.com/en-us/articles/intel-advanced-encryption-standard-aes-instructions-set` (September 2012), Revision 3.01

45. Halevi, S.: EME*: Extending EME to Handle Arbitrary-Length Messages with Associated Data. In: INDOCRYPT 2004. LNCS, vol. 3348, pp. 315–327. Springer (2004)

46. Halevi, S., Rogaway, P.: A Tweakable Enciphering Mode. In: CRYPTO 2003. LNCS, vol. 2729, pp. 482–499. Springer (2003)

47. Halevi, S., Rogaway, P.: A Parallelizable Enciphering Mode. In: CT-RSA 2004. LNCS, vol. 2964, pp. 292–304. Springer (2004)

48. Hoang, V.T., Krovetz, T., Rogaway, P.: Robust Authenticated-Encryption AEZ and the Problem That It Solves. In: EUROCRYPT 2015. LNCS, vol. 9056, pp. 15–44. Springer (2015)

49. Hoang, V.T., Rogaway, P.: On Generalized Feistel Networks. In: CRYPTO 2010. LNCS, vol. 6223, pp. 613–630. Springer (2010)

50. Holenstein, T., Künzler, R., Tessaro, S.: The equivalence of the random oracle model and the ideal cipher model, revisited. In: STOC 2011. pp. 89–98. ACM (2011)

51. Jean, J.: Cryptanalysis of Haraka. Cryptology ePrint Archive, Report 2016/396 (2016)

52. Jean, J., Nikolić, I., Sasaki, Y., Wang, L.: Practical Cryptanalysis of PAES. In: SAC 2014. LNCS, vol. 8781, pp. 228–242. Springer (2014)

53. Jean, J., Nikolić, I., Sasaki, Y., Wang, L.: Practical Forgeries and Distinguishers against PAES. IEICE Transactions 99-A(1), 39–48 (2016)

54. Knudsen, L.R.: Truncated and Higher Order Differentials. In: FSE 1994. LNCS, vol. 1008, pp. 196–211. Springer (1994)

55. Kölbl, S., Lauridsen, M.M., Mendel, F., Rechberger, C.: Haraka - Efficient Short-Input Hashing for Post-Quantum Applications. Cryptology ePrint Archive, Report 2016/098 (2016)

56. Lamport, L.: Constructing Digital Signatures from a One Way Function. Tech. Rep. SRI-CSL-98, SRI International Computer Science Laboratory (October 1979)

57. Langford, S.K., Hellman, M.E.: Differential-Linear Cryptanalysis. In: CRYPTO 1994. LNCS, vol. 839, pp. 17–25. Springer (1994)

58. Leander, G., Abdelraheem, M.A., AlKhzaimi, H., Zenner, E.: A Cryptanalysis of PRINTcipher: The Invariant Subspace Attack. In: CRYPTO 2011. LNCS, vol. 6841, pp. 206–221. Springer (2011)

59. Liskov, M., Rivest, R.L., Wagner, D.: Tweakable Block Ciphers. In: CRYPTO 2002. LNCS, vol. 2442, pp. 31–46. Springer (2002)

60. Liskov, M., Rivest, R.L., Wagner, D.: Tweakable Block Ciphers. J. Cryptology 24(3), 588–613 (2011)

61. Lucks, S.: BEAST: A Fast Block Cipher for Arbitrary Blocksizes. In: CMS 1996. IFIP Conference Proceedings, vol. 70, pp. 144–153. Chapman & Hall (1996)

62. Matsui, M.: Linear Cryptanalysis Method for DES Cipher. In: EUROCRYPT 1993. LNCS, vol. 765, pp. 386–397. Springer (1994)

63. Maurer, U.M., Renner, R., Holenstein, C.: Indifferentiability, Impossibility Results on Reductions, and Applications to the Random Oracle Methodology. In: TCC 2004. LNCS, vol. 2951, pp. 21–39. Springer (2004)

64. Mendel, F., Rechberger, C., Schläffer, M., Thomsen, S.S.: The Rebound Attack: Cryptanalysis of Reduced Whirlpool and Grøstl. In: FSE 2009. LNCS, vol. 5665, pp. 260–276. Springer (2009)

65. Moriai, S., Vaudenay, S.: On the Pseudorandomness of Top-Level Schemes of Block Ciphers. In: ASIACRYPT 2000. LNCS, vol. 1976, pp. 289–302. Springer (2000)

66. Mouha, N.: The Design Space of Lightweight Cryptography. Cryptology ePrint Archive, Report 2015/303 (2015)

67. Mouha, N., Luykx, A.: Multi-key Security: The Even-Mansour Construction Revisited. In: CRYPTO 2015. LNCS, vol. 9215, pp. 209–223. Springer (2015)

68. Mouha, N., Mennink, B., Herrewege, A.V., Watanabe, D., Preneel, B., Verbauwhede, I.: Chaskey: An Efficient MAC Algorithm for 32-bit Microcontrollers. In: SAC 2014. LNCS, vol. 8781, pp. 306–323. Springer (2014)

69. Mouha, N., Wang, Q., Gu, D., Preneel, B.: Differential and Linear Cryptanalysis Using Mixed-Integer Linear Programming. In: Inscrypt 2011. LNCS, vol. 7537, pp. 57–76. Springer (2011)

70. Rechberger, C.: On Bruteforce-Like Cryptanalysis: New Meet-in-the-Middle Attacks in Symmetric Cryptanalysis. In: ICISC 2012. LNCS, vol. 7839, pp. 33–36. Springer (2013)

71. Rogaway, P., Steinberger, J.P.: Security/Efficiency Tradeoffs for Permutation-Based Hashing. In: EUROCRYPT 2008. LNCS, vol. 4965, pp. 220–236. Springer (2008)

72. Rønjom, S.: Personal Communication (March 2016)

73. Rønjom, S.: Invariant subspaces in Simpira. Cryptology ePrint Archive, Report 2016/248 (2016)

74. Schroeppel, R.: The Hasty Pudding Cipher – A Tasty Morsel (1998), submission to the NIST AES competition

75. Suzaki, T., Minematsu, K.: Improving the Generalized Feistel. In: FSE 2010. LNCS, vol. 6147, pp. 19–39. Springer (2010)

76. Todo, Y.: Structural Evaluation by Generalized Integral Property. In: EUROCRYPT 2015. LNCS, vol. 9056. Springer (2015)

77. Wagner, D.: The Boomerang Attack. In: FSE 1999. LNCS, vol. 1636, pp. 156–170. Springer (1999)

78. Yanagihara, S., Iwata, T.: On Permutation Layer of Type 1, Source-Heavy, and Target-Heavy Generalized Feistel Structures. In: CANS 2011. LNCS, vol. 7092, pp. 98–117. Springer (2011)

79. Yanagihara, S., Iwata, T.: Improving the Permutation Layer of Type 1, Type 3, Source-Heavy, and Target-Heavy Generalized Feistel Structures. IEICE Transactions 96-A(1), 2–14 (2013)

80. Yanagihara, S., Iwata, T.: Type 1.x Generalized Feistel Structures. IEICE Transactions 97-A(4), 952–963 (2014)

81. Zhang, H., Wu, W.: Structural Evaluation for Generalized Feistel Structures and Applications to LBlock and TWINE. In: INDOCRYPT 2015. LNCS, vol. 9462, pp. 218–237. Springer (2015)

82. Zheng, Y., Matsumoto, T., Imai, H.: On the Construction of Block Ciphers Provably Secure and Not Relying on Any Unproved Hypotheses. In: CRYPTO 1989. LNCS, vol. 435, pp. 461–480. Springer (1990)

A Simpira for $b \notin \{1, 2, 3, 4, 6, 8\}$: Full Bit Diffusion and Active S-Boxes

For the Simpira construction with $b \notin \{1, 2, 3, 4, 6, 8\}$, it is very straightforward to show that $4b - 6$ F-functions are sufficient to reach full bit diffusion, as well as to ensure that at least 30 S-boxes are active. This can be proven by induction.

Full Bit Diffusion. Recall that full bit diffusion means that every output bit depends on every input bit, and every input bit depends on every output bit. Observe that the construction of Fig. 3 reaches full bit diffusion for $b = 4$. From the same figure, it can also easily be seen that the four additional F-functions ensure that full bit diffusion is reached for $b + 1$, provided that full bit diffusion was already reached for b.

Active S-boxes. The MILP tool shows us that the construction of Fig. 3 reaches at least 30 (linearly or differentially) active S-boxes for $b = 4$. We now prove that if the construction for b has at least 30 active S-boxes, then it has at least 30 active S-boxes for $b + 1$. When a non-zero difference or a non-zero linear mask enters into the construction for b, the result follows directly. If this is not the case, this imposes a restriction on inputs and outputs for the construction of b: the difference (or linear mask) of every input and output subblock must be zero. In that case, the MILP tool proved that the four new F-functions that were added when going from b to $b + 1$ ensure that there will be at least 30 active S-boxes. One such case is illustrated in Fig. 15.

B Efficient Implementation For $b = 1$

We recall the four Intel instructions to implement one round of AES: `AESENC` (Alg. 13), `AESENCLAST` (Alg. 14), `AESDEC` (Alg. 15), and `AESENCLAST` (Alg. 16). The `AESIMC` instruction corresponds to the `InvMixColumns` operation. Then for $b = 1$, Simpira can be implemented as in Alg. 17, and Simpira^{-1} as in Alg. 18.

C Optimizing the Number of F-functions

In Sect. 3, we based ourselves mainly on designs that were published in literature, instead of exhaustively searching for the optimal design that satisfies the design criteria. Here, we revisit this assumption. In particular, we will consider the case where $b = 4$.

An exhaustive search for all GFS with $b = 4$ shows that full bit diffusion requires at least 8 F-functions. One such construction is shown in Fig. 17. However, we found that all designs with 8 F-functions have a lower bound of at most 10 for the number of active S-boxes, and therefore do not satisfy our design criteria.

Fig. 15. A differential characteristic for (reduced-round) Simpira with $b = 5$ that has 30 active S-boxes. A thick full line indicates a difference in every byte, a thick dotted line refers to a difference in only one byte – it does not matter which one. A normal line indicates that no difference is present. When non-zero, the number of active S-boxes is shown above every F-function.

Algorithm 13 AESENC (= Alg. 1)

1: **procedure** AESENC(state, key)
2:     state ← SubBytes(state)
3:     state ← ShiftRows(state)
4:     state ← MixColumns(state)
5:     state ← state ⊕ key
6:     **return** state
7: **end procedure**

Algorithm 14 AESENCLAST

1: **procedure** AESENCLAST(state, key)
2:     state ← SubBytes(state)
3:     state ← ShiftRows(state)
4:
5:     state ← state ⊕ key
6:     **return** state
7: **end procedure**

Algorithm 15 AESDEC

1: **procedure** AESDEC(state, key)
2:     state ← InvSubBytes(state)
3:     state ← InvShiftRows(state)
4:     state ← InvMixColumns(state)
5:     state ← state ⊕ key
6:     **return** state
7: **end procedure**

Algorithm 16 AESDECLAST

1: **procedure** AESDECLAST(state, key)
2:     state ← InvSubBytes(state)
3:     state ← InvShiftRows(state)
4:
5:     state ← state ⊕ key
6:     **return** state
7: **end procedure**

Algorithm 17 Simpira ($b = 1$) (= Alg. 3)

1: **procedure** SIMPIRA(x_0)
2:     $R \leftarrow 6$
3:     **for** $r = 1, \ldots, R - 1$ **do**
4:         $C \leftarrow$ SETR_EPI32(0x00 $\oplus r \oplus R$,
5:                 0x10 $\oplus r \oplus R$,
6:                 0x20 $\oplus r \oplus R$,
7:                 0x30 $\oplus r \oplus R$)
8:
9:         $x_0 \leftarrow$ AESENC(x, C)
10:        $x_0 \leftarrow$ AESENC($x, 0$)
11:        $c \leftarrow c + 1$
12:     **end for**
13:     $C \leftarrow$ SETR_EPI32(0x00 $\oplus R \oplus R$,
14:             0x10 $\oplus R \oplus R$,
15:             0x20 $\oplus R \oplus R$,
16:             0x30 $\oplus R \oplus R$)
17:
18:     $x_0 \leftarrow$ AESENC(x, C)
19:     $x_0 \leftarrow$ AESENCLAST($x, 0$)
20:     **return** x_0
21: **end procedure**

Algorithm 18 Simpira^{-1} ($b = 1$) (= Alg. 4)

1: **procedure** SIMPIRA(x_0)
2:     $R \leftarrow 6$
3:     **for** $r = R, \ldots, 1$ **do**
4:         $C \leftarrow$ SETR_EPI32(0x00 $\oplus r \oplus R$,
5:                 0x10 $\oplus r \oplus R$,
6:                 0x20 $\oplus r \oplus R$,
7:                 0x30 $\oplus r \oplus R$)
8:         $C \leftarrow$ AESIMC(C)
9:         $x_0 \leftarrow$ AESDEC(x, C)
10:        $x_0 \leftarrow$ AESDEC($x, 0$)
11:        $c \leftarrow c + 1$
12:     **end for**
13:     $C \leftarrow$ SETR_EPI32(0x00 $\oplus 1 \oplus R$,
14:             0x10 $\oplus 1 \oplus R$,
15:             0x20 $\oplus 1 \oplus R$,
16:             0x30 $\oplus 1 \oplus R$)
17:     $C \leftarrow$ AESIMC(C)
18:     $x_0 \leftarrow$ AESDEC(x, C)
19:     $x_0 \leftarrow$ AESDECLAST($x, 0$)
20:     **return** x_0
21: **end procedure**

Fig. 16. In Alg. 13–18, we recall the AES-NI instructions of [44] and show how they can be used to efficiently implement Simpira and its inverse for $b = 1$.

To reach both full bit diffusion and at least 25 active S-boxes, at least 9 F-functions are required. An example of such a construction is shown in Fig. 18. This construction has 35 active S-boxes, and would therefore lead to a slightly better construction that could replace the current choice for Simpira with $b = 4$.

However, the problem with this approach is all of these Feistels have a very random-looking structure. When there is no simple structure, the design becomes more difficult to cryptanalyze, and possibly also more difficult implement.

It therefore seems better to consider only GFS with either identical round functions ($b \in \{2, 3, 4, 6, 8\}$ of Simpira), or with the TwoF function that is used in Simpira for large b. With this restriction, the Simpira design for $b = 4$ with 10 F-functions is an optimal according to the design criteria.


Fig. 17. A GFS with $b = 4$ that uses 8 F-functions to reach full bit diffusion. This design has at least 10 active S-boxes.

Fig. 18. A GFS with $b = 4$ that uses 9 F-functions to reach full bit diffusion. This design has at least 35 active S-boxes.

D   Comparison with Other Constructions

For comparison, we now provide the throughput of SHA-256, SHA-512, and Rijndael256 (with a 256-bit block size), measured on the same platform, and using the same methodology. In the case of SHA-256 and SHA-512, we wrote

an optimized throughput-oriented implementation that uses the AVX2 architecture, available on the discussed platform. For SHA-256 and SHA-512, this implementation processes 4 and 8 independent (long) buffers respectively. For Rijndael256, we prepared optimized code that uses AES-NI (see details in [44]). We measured it in ECB mode, operating on 8 blocks in parallel, to get the highest throughput possible on this platform.

Under this setup, the throughput of SHA-256, SHA-512, and Rijndael256 is 2.35, 3.13, and 1.54 cycles per byte, respectively. Therefore, for $b = 2$, it is clearly much faster to use the Simpira permutation, which requires only 0.94 cycles per byte. This permutation is to be used inside an Even-Mansour construction (for encryption), or with a Davies-Meyer feedforward (for hashing); but these operations not change the throughput in a noticeable way.

For larger b, it is interesting to compare Simpira with two-pass constructions, for example for encryption. We cannot use AES as a building block in a typical two-pass construction, as it would be insecure beyond about 2^{64} input blocks, and we aim for security up to 2^{128} blocks. We may choose Simpira with $b = 2$ as a building block, as typical alternatives such as Rijndael256 are slower on our target platform.

In a double-pass mode of operation, Simpira with $b = 2$ requires at least $2 \cdot 30 \cdot b$ AESENC operations per $32b$ bytes, which is $30b$ AESENC operations per $16b$ bytes. When b is large, Simpira requires $24b - 36$ AESENC operations per $16b$ bytes, which is less than the previously mentioned double-pass mode of operation, even for very large b.

On a sidenote: Simpira with large b is faster than the two-pass construction, but cannot process the input blocks in parallel: to fill the pipeline, a sufficient number of independent messages is required. Therefore, which of these two Simpira-based constructions is better, depends on the application.