

The PHOTON Family of Lightweight Hash Functions

Jian Guo¹, Thomas Peyrin^{*2}, and Axel Poschmann^{*2}

¹ Institute for Infocomm Research, Singapore

² Nanyang Technological University, Singapore
{ntu.guo,thomas.peyrin}@gmail.com, aposchmann@ntu.edu.sg

Abstract. RFID security is currently one of the major challenges cryptography has to face, often solved by protocols assuming that an on-tag hash function is available. In this article we present the PHOTON lightweight hash-function family, available in many different flavors and suitable for extremely constrained devices such as passive RFID tags. Our proposal uses a sponge-like construction as domain extension algorithm and an AES-like primitive as internal unkeyed permutation. This allows us to obtain the most compact hash function known so far (about 1120 GE for 64-bit collision resistance security), reaching areas very close to the theoretical optimum (derived from the minimal internal state memory size). Moreover, the speed achieved by PHOTON also compares quite favorably to its competitors. This is mostly due to the fact that unlike for previously proposed schemes, our proposal is very simple to analyze and one can derive tight AES-like bounds on the number of active Sboxes. This kind of AES-like primitive is usually not well suited for ultra constrained environments, but we describe in this paper a new method for generating the column mixing layer in a serial way, lowering drastically the area required. Finally, we slightly extend the sponge framework in order to offer interesting trade-offs between speed and preimage security for small messages, the classical use-case in hardware.

Key words: lightweight, hash function, sponge function, AES.

1 Introduction

RFID tags are likely to be deployed widely in many different situations of everyday life and they represent a great business opportunity for various markets. However, this rising technology also provides new security challenges that the cryptography community has to handle. RFID tags can be used to fight product counterfeiting by authenticating them and on the other hand, we would also like to guarantee the privacy of the users.

These two security aspects have already been studied considerably and, interestingly, in most of the privacy-preserving RFID protocols proposed [5, 36, 46] a hash function is required. Informally, such a primitive is a function that takes an arbitrary length input and outputs a fixed-size value. While no secret is involved in the computation, one would like that finding collisions (two distinct messages hashing to the same value) or (second)-preimages (a message input that hashes to a given challenge output value) is computationally intractable for an attacker. More precisely, for an n -bit ideal hash function we expect to perform $2^{n/2}$ and 2^n computations in order to find a collision and a (second)-preimage respectively. While not as mature as block-ciphers, the research on hash functions saw a rapid development lately, mainly due to the groundbreaking attacks on standardized primitives [71, 69, 70]. At the present time, most of the attention of the symmetric key cryptography academic community is focused on the SHA-3 competition organized by NIST [55], which should provide a potential replacement of the MD-SHA family.

In parallel, nice advances have also been made in the domain of lightweight symmetric key primitives in the last years. Protocol designers now have at disposal PRESENT [16], a 64-bit block-cipher with 80-bit key whose security has already been analyzed intensively and that can be as compact as 1075 GE [64]. Stream-ciphers are not outcast with implementations [33] with 80-bit security requiring about 1300 GE and 2600 GE reported for GRAIN [35] and TRIVIUM [23] respectively, two candidates selected in the final eSTREAM hardware portfolio. However, the situation is not as bright in the case of hash functions.

* The authors were supported in part by the Singapore National Research Foundation under Research Grant NRF-CRP2-2007-03.

As already pointed out in [30] and echoed in [17], the community lacks very compact hash functions. Standardized primitives such as SHA-1 [53] or SHA-2 [54] are much too large to fit in very constrained hardware (5527 GE [56] and 10868 GE [30] for 80 and 128-bit aimed security respectively) and even compact-oriented proposals such as MAME [72] require 8100 GE for 128-bit security. While hardware is not an important criteria in the selection process, one can not expect the SHA-3 finalists to be much more compact. At the present time, all SHA-3 finalists require more than 12000 GE for 128-bit security (smaller versions of KECCAK that have not been submitted to the competition provide for example 64-bit security with 5090 GE [40]). Note that a basic RFID tag may have a total gate count of anywhere from 1000-10000 gates, with only 200-2000 gates budgeted for security [39].

This compactness problem in hash algorithms is partly due to the fact that it widely depends on the memory registers required for the computation. Most hash functions proposed so far are software-oriented and output at least 256 bits in order to be out of reach of any generic collision search in practice. While such an output size makes sense where high level and long-term security are needed, RFID use-cases could bear much smaller security parameters. This is for example the path taken in [17], where the authors instantiate lightweight hash functions using literature-based constructions [37, 62] with the compact block-cipher PRESENT [16]. With SQUASH [66], Shamir proposed a compact keyed hash function inspired by the Rabin encryption scheme that processes short messages (at most 64-bit inputs) and that provides 64 bits of preimage security, without being collision resistant. At CHES 2010, the lightweight hash-function family ARMADILLO [6] was proposed, but has recently been shown to present serious security weaknesses [15]. At the same conference, Aumasson *et al.* published the hash function QUARK [4], using sponge functions [7] as domain extension algorithm, and an internal permutation inspired from the stream-cipher GRAIN [35] and the block-cipher KATAN [22]. Using sponge functions as operating mode is another step towards compactness. Indeed, classical n -bit hash function constructions like the MD-SHA family utilize a Merkle-Damgård [51, 28] domain extension algorithm with a compression function h built upon an n -bit block-cipher E in Davies-Meyer mode ($h(CV, M) = E_M(CV) \oplus CV$), where CV stands for the chaining variable and M for the current message block. Avoiding any feed-forward like for sponge constructions saves a lot of memory registers at the cost of an invertible iterative process which induces a lower (second)-preimage security for the same internal state size. All in all, designers have to deal with a trade-off between security and memory requirements.

In this article, we describe a new hardware-oriented hash-function family: PHOTON. We chose to use the sponge functions framework in order to keep the internal memory size as low as possible. However, we extend this framework so as to provide very interesting trade-offs in hardware between preimage security and small messages hashing speed (small message scenario is a classical use-case and can be problematic for sponge functions because of their squeezing process that can be very slow in practice). The internal permutations of PHOTON can be seen as AES-like primitives especially derived for hardware: our columns mixing layer can be computed in a serial way while maintaining optimal diffusion properties. Overall, as shown in Table 5 in Section 5.3, not only PHOTON is easily the smallest hash function known so far, but it also achieves excellent area/throughput trade-offs.

In terms of security, it is particularly interesting to use AES-like permutations as we can fully leverage all the previous cryptanalysis performed on AES and on AES-based hash functions. Moreover, we can directly derive very simple bounds on the number of active Sboxes for 4 rounds of the permutation. These bounds being tight, we can confidently set an appropriate number of rounds that ensures a comfortable security margin.

2 Design Choices

In tag-based applications, one typically does not require high security primitives, such as a 512-bit output hash function. In contrary, 64 or 80-bit security is often appropriate considering the value of objects an RFID tag is protecting and the use cases. Moreover, a designer should use exactly the level that he expects from his primitive, so as to avoid any waste of area or computing power. This is the reason why we chose to precisely instantiate several security levels for PHOTON, ranging from 64-bit preimage resistance security to 128-bit collision resistance security.

2.1 Extended Sponge functions

Sponge functions have been introduced by Bertoni *et al.* [7] as a new way of building hash functions from a fixed permutation (later more applications were proposed [10]). The internal state S of t bits, composed

of the c -bit capacity and the r -bit bitrate ($t = c + r$), is first initialized with some fixed value. Then, after having appropriately padded and split the message into r -bit chunks, one simply and iteratively processes all r -bit message chunks by xoring them to the bitrate part of the internal state and then applying the t -bit permutation P . Once all message chunks have been handled by this absorbing phase, one successively outputs r bits of the final hash value by extracting r bits from the bitrate part of the internal state and then applying the permutation P on it (squeezing process).

When the internal permutation P is modeled as a randomly chosen permutation, a sponge function has been proven to be indifferentiable from a random oracle [8] up to $2^{c/2}$ calls to P . More precisely, for an n -bit sponge hash function with capacity c and bitrate r , when the internal primitive is modeled as a random permutation, one obtains $\min\{2^{n/2}, 2^{c/2}\}$ as collision resistance bound and $\min\{2^n, 2^{c/2}\}$ as (second)-preimage bound. However, in the case of preimage, there exists a gap between this bound and the best known generic attack³. Therefore, we expect the following complexities in the generic case:

- **Collision:** $\min\{2^{n/2}, 2^{c/2}\}$
- **Second-preimage:** $\min\{2^n, 2^{c/2}\}$
- **Preimage:** $\min\{2^{\min\{n,t\}}, \max\{2^{\min\{n,t\}-r'}, 2^{c/2}\}\}$

Moreover, sponge functions can be used as a Message Authentication Code with $MAC_K(M) = H(K||M)$, where $K \in \{0, 1\}^k$ stands for the key and M for the message. It has been shown [11] that as long as the amount of message queries is limited to 2^a with $a \ll c/2$, then no attack better than exhaustive key search exists if $c \geq k + a + 1$.

Sponge functions seem a natural choice in order to minimize the amount of memory registers in hardware since they can offer speed/area/security trade-offs. Indeed, the only memory required for the internal state is $t = c + r$ bits, while for a classical Davies-Meyer construction using an m -bit block cipher with a k -bit key input one needs to store $2m + k$ bits, out of which m bits are required for the feed-forward. For an equivalent ideal collision security level (thus setting $m = c = n$) and by minimizing the area (r and k are very small), the sponge function requires only about half of the memory. Note that if one looks for a perfectly (second)-preimage resistant hash function (up to the 2^n ideal bound), then it is required that $c \geq 2 \cdot n$ (which implies that the n -bit hash function built is indifferentiable from an n -bit random oracle anyway). In that particular case the sponge functions are not better than the Davies-Meyer construction in terms of area requirements and therefore in this work we will not focus on this scenario. Instead, we will build hash functions that may have ideal resistance to collision, but not for (second)-preimage. The typical shape will be a capacity c equal to the hash output n and a very small bitrate r . This security/area trade-off, already utilized by the QUARK designers, will allow us to aim at extremely low area requirements, while maintaining security expectations very close to ideal.

In [17], the authors identify that in most RFID applications the user will not hash a large amount of data, *i.e.* in general less than 256 bits. Consider for example the *electronic product code (EPC)* number, which is a 96-bit string that is meant to identify globally any tag/product. In this particular case of small messages, sponge functions with a small bitrate r seem to be slow since one needs to call $(\lceil n/r \rceil - 1)$ times the internal permutation to complete the final squeezing process. This is for example the case with U-QUARK, that has a throughput of 1.47 kbps for very long messages which drops to 0.63 kbps for 96-bit inputs. On the other side, this “small messages” effect is reduced by the fact that having a small bitrate will reduce the amount of padding actually hashed (the padding simply consists in adding a “1” and as many “0” required to fill the last message block). Note that lightweight proposals based on classical Davies-Meyer construction that include the message length as suffix padding are also slow for small messages: DM-PRESENT-80 has a throughput of 14.63 kbps for very long messages which drops to 5.85 kbps for 96-bit inputs, because in the latter case many of the compression function calls are spent in order to handle padding blocks.

In order to allow more flexibility about this issue, we propose to slightly extend the sponge framework by allowing the number r' of bits extracted during each iteration of the squeezing process to be different from the bitrate r ⁴ (see Figure 1). Increasing r' will directly reduce the time spent in the squeezing process, but

³ The $2^{\min\{n,t\}-r}$ term for preimage comes from the fact that in order to invert the hash function the attacker will have to invert the squeezing process and the best known generic attack to solve this “multiblock constrained-input constrained-output problem” [9] requires 2^{n-r} computations when $t \geq n$, and 2^{t-r} otherwise. When the internal state just before the squeezing process has been recovered, the attacker can run a meet-in-the-middle attack with $2^{c/2}$ computations. If the attacker does not invert the squeezing process, then he will have to pay the generic preimage cost 2^n when $t \geq n$, and 2^t otherwise.

⁴ A recent work from Andreeva *et al.* [2] also independently proposed such an extension of the sponge model.




Fig. 1. The extended sponge framework, the domain extension algorithm used by the PHOTON hash-function family.

might also reduce the preimage security. On the contrary, decreasing r' might improve the preimage bound at the cost of a speed drop for small messages. As long as the preimage security remains in an acceptable bound, this configuration can be interesting in many scenarios where only tiny inputs are to be hashed. More precisely, in this new model, the best known generic attacks require the following amount of computations:

- **Collision:** $\min\{2^{n/2}, 2^{c/2}\}$
- **Second-preimage:** $\min\{2^n, 2^{c/2}\}$
- **Preimage:** $\min\{2^{\min\{n,t\}}, \max\{2^{\min\{n,t\}-r'}, 2^{c/2}\}\}$

Finally, in most tag-based applications the collision resistance is not a requirement, while only the one-wayness of the function must be ensured. However, as we previously explained, for lightweight scenarios the sponge construction does not maintain the (second)-preimage security at the full level of its capacity c . This is due to the output process of the sponge operating mode. Of course, performing a Davies-Meyer like feed-forward just after the final truncation would do the job, but that would also double the memory area required (which is precisely what we are trying to avoid). The nice trick of squeezing in the sponge functions framework permits to avoid any feed-forward while somehow rendering the process non-invertible, up to some extend (see multiblock constrained-input constrained-output problem in [9]). One solution to reach the full capacity preimage security would be to add one more squeezing iteration, thus increasing the output size of the hash by r' bits.⁵ Then, the best known generic preimage attack for this $(n + r')$ -bit hash function will run in

$$\min\{2^{\min\{n+r',t\}}, \max\{2^{\min\{n,t-r'\}}, 2^{c/2}\}\} \geq 2^n \text{ when } c + r - r' \geq n$$

and one has to note that this hash output extension has no influence on the second-preimage resistance.

In this article, we will provide five sizes of internal permutations and one PHOTON flavor for each of them. The four biggest versions fit the classical sponge model and will ensure $2^{n/2}$ collision and second preimage resistance and 2^{n-r} concerning preimage. However, in order to illustrate the powerful trade-offs allowed by our extended model, the smaller PHOTON variant will have different input/output bitrates and an extended hash size. Using the five permutations defined in the next Section, one can derive its own PHOTON flavor depending on the collision / (second)-preimage / MAC security required, the maximal area and the maximal hash output size allowed. The process to obtain the optimal parameters is given in Appendix A. Note that the area required will only depend on the internal permutation chosen.

2.2 An AES-like internal permutation

We define an AES-like function to be a fixed key permutation P applied on an internal state of d^2 elements of s bits each, which can be represented as a $(d \times d)$ matrix. P is composed of N_r rounds, each containing four layers as depicted in Figure 3: AddConstants (AC), SubCells (SC), ShiftRows (ShR), and MixColumnsSerial

⁵ This generalization has been independently utilized by the QUARK designers in a revised version of their original article.

(MCS). Informally, `AddConstants` simply consists in adding fixed values to the cells of the internal state, while `SubCells` applies an s -bit Sbox to each of them. `ShiftRows` rotates the position of the cells in each of the rows and `MixColumnsSerial` linearly mixes all the columns independently.

We chose to use AES-like permutations because they offer much confidence in the design strategy as one can leverage previous cryptanalysis works done on AES and on AES-like hash functions. Moreover, AES-like permutations allow to derive very simple proofs on the number of active Sboxes over four rounds of the primitive. More precisely, if the matrix underlying the `MixColumnsSerial` layer is Maximum Distance Separable (MDS), then one can immediately show that at least $(d + 1)^2$ Sboxes will be active for any 4-round differential path [27]. This bound is tight, and we already know differential paths with only $(d + 1)^2$ active Sboxes for four rounds (we will use them later for security analysis purposes). Moreover, note that the permutations we will design are fixed-key, so we naturally get rid of related-key attacks or any issue that might arise from the construction of a key-schedule [12, 13].




Fig. 2. One round of a PHOTON permutation.

AddConstants. The constants have been chosen such that each of the N_r round computations are different, and such that the classical symmetry between columns in AES-like designs are destroyed (without the `AddConstants` layer, an input with all columns equal would maintain this property through any number of rounds). Also, the round constants can be generated by a combination of very compact Linear Feedback Shift Registers. For performance reasons, only the first column of the internal state is involved.

SubCells. Our choice of the Sboxes was mostly motivated by their hardware quality. 4-bit Sboxes can be very compact in hardware while the acceptable upper limit on the cell size is $s = 8$. We avoided to use an Sbox size s which is odd, because this leads to odd message block size or capacity when d is also odd. This leaves us with $s = 4, 6, 8$, but we also believe that reusing some already trusted and well analyzed components increases the confidence in the security of the scheme and saves a lot of time for cryptanalysts. Finally, we will use two types of Sboxes: the 4-bit PRESENT Sbox $SBOX_{PRE}$ and the 8-bit AES Sbox $SBOX_{AES}$ the latter being only utilized for high security levels (at least 128 bits of collision resistance). Note also that $s = 4, 8$ allows simpler and faster software implementations.

ShiftRows. The choice of the `ShiftRows` constants is very simple for PHOTON since our internal state is always a square of cells. Therefore, row i will classically be rotated by i positions to the left, i counts from 0.

MixColumnsSerial. The matrix underlying the AES `MixColumns` function is a circulant matrix with low hamming weight coefficients. Even if those coefficients and the irreducible polynomial used to create the Galois field for the AES `MixColumns` function have been chosen so as to improve the hardware footprint of the cipher, it can not be implemented in an extremely compact way. One of the main reason is that the byte-serial implementation of this function is not compact. Said in other words, if we write the AES `MixColumns` matrix as the composition of d operations each updating a single byte at a time in a serial way, then the coefficients of these d matrices will be very bad for small area implementations.

In order to solve this issue, we took the problem the other way round. Let A be the matrix that updates the last cell of the column vector with a linear combination of all of the vector cells and then rotates the vector by one position towards the top. Our new MixColumnsSerial layer will be composed of d applications of this matrix to the input column vector. More formally, let $X = (x_0, \dots, x_{d-1})^T$ be an input column vector of MixColumnsSerial and $Y = (y_0, \dots, y_{d-1})^T$ be the corresponding output. Then, we have $Y = A^d \times X$, where A is a $(d \times d)$ matrix of the form:

$$A = \begin{pmatrix} 0 & 1 & 0 & 0 & \cdots & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & \cdots & 0 & 0 & 0 & 0 \\ & \vdots & & & & & & \vdots & \\ 0 & 0 & 0 & 0 & \cdots & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & \cdots & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & \cdots & 0 & 0 & 0 & 1 \\ Z_0 & Z_1 & Z_2 & Z_3 & \cdots & Z_{d-4} & Z_{d-3} & Z_{d-2} & Z_{d-1} \end{pmatrix}$$

where coefficients (Z_0, \dots, Z_{d-1}) can be chosen freely. We denote by $Serial(Z_0, \dots, Z_{d-1})$ such a matrix. Of course, we would like the final matrix A^d to be MDS, so as to maintain as much diffusion as for the AES initial design strategy. For each square size d we picked during the design of PHOTON, we used MAGMA [18] to test all the possible values of Z_0, \dots, Z_{d-1} and picked the most compact candidate making A^d an MDS matrix. We also chose the irreducible polynomial with compactness as main criterion.

For design strategy comparison purposes, we can take as an example the AES case. By using our new mixing layer design method, we were able to find the matrix $A = Serial(1, 2, 1, 4)$ which gives the following MDS final matrix:

$$(A)^4 = \begin{pmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 1 & 2 & 1 & 4 \end{pmatrix}^4 = \begin{pmatrix} 1 & 2 & 1 & 4 \\ 4 & 9 & 6 & 17 \\ 17 & 38 & 24 & 66 \\ 66 & 149 & 100 & 11 \end{pmatrix}$$

The smallest AES hardware implementation requires 2400 GE [52], for which 263 GE are dedicated to MixColumns. It is possible to implement MixColumns of AES in a byte-by-byte fashion, which requires only 81 GE to calculate one byte of the output column. However, since AES uses a circulant matrix, at least three additional 8-bit registers (144 GE), are required to hold the output, plus additional control logic, which increases the area requirements significantly. That is why [52] does not use a serial MixColumns, but rather processes one column at a time.

Please note that in general the choice of non-zero constants for any $d \times d$ MDS matrix on s -bit cells has only a minor impact of the area consumption, since a multiplication by x consists of w XOR gates, where w denotes the Hamming weight of the irreducible polynomial used. At the same time, $(d-1) \cdot s$ XOR gates are required to sum up the d individual terms of s bits each. It is no surprise, that multiplying with the constants above accounts for only 21.3 GE out of the 74 GE required. In fact, the efficiency of our approach lies in the shifting property of A , since this allows to re-use the existing memory with neither temporary storage nor additional control logic required.

All in all, using our approach would provide a tweaked AES cipher with the very same diffusion properties as the original one (the matrix being MDS), but that can fit in only 2210 GE, a total saving of around 8%. Moreover, for the deciphering process, a slightly modified hardware can be used in order to unroll the MixColumnsSerial, further reducing the area footprint of such a PHOTON-based cipher. One might think that the software implementations will suffer from this new layer. While our goal is to make a hardware-oriented primitive, we would like to remark that most AES software implementations are precomputed tables-based (applying both the Sbox and the MixColumns coefficients at the same time) and the very same method can be applied to PHOTON. This is confirmed by our first software implementations, whose benchmarks are given in Section 5.4.

3 The PHOTON Hash-Function Family

We describe in this section the PHOTON family of hash functions. Each variant will be fully defined by its hash output size $64 \leq n \leq 256$, its input and its output bitrate r and r' respectively. Therefore we denote each function PHOTON- $n/r/r'$. The internal state size $t = (c + r)$ depends on the hash output size and can take only 5 distinct values: 100, 144, 196, 256 and 288 bits. As a consequence, we only have to define 5 internal permutations P_t , one for each internal state size.

In order to cover a wide spectrum of applications, we propose five different flavors of PHOTON, one for each internal state size: PHOTON-80/20/16, PHOTON-128/16/16, PHOTON-160/36/36, PHOTON-224/32/32 and PHOTON-256/32/32 will use internal permutations P_{100} , P_{144} , P_{196} , P_{256} and P_{288} respectively. Note that the first proposal is special in the sense that it is designed for the specific cases where 64-bit preimage security and 64-bit key MAC are considered to be sufficient.⁶ In contrary, the last proposal provides a high security level of 128-bit collision resistance, thus making it suitable for generic applications.

3.1 The domain extension algorithm

The message M to hash is first padded by appending a “1” bit and as many zeros (possibly none) such that the total length is a multiple of the bitrate r and we can finally obtain l message blocks m_0, \dots, m_{l-1} of r bits each. The t -bit internal state S is initialized by setting it to the value $S_0 = IV = \{0\}^{t-24} || n/4 || r || r'$, where $||$ denotes the concatenation and each value is coded on 8 bits. For implementation purposes, note that each byte is interpreted in big-endian form.

Then, as for the classical sponge strategy, at iteration i we absorb the message block m_i on leftmost part of the internal state S_i ⁷ and then apply the permutation P_t , *i.e.* $S_{i+1} = P_t(S_i \oplus (m_i || \{0\}^c))$. Once all l message blocks have been absorbed, we build the hash value by concatenating the successive r' -bit output blocks z_i until we reach the appropriate output size n :

$$hash = z_0 || \dots || z_{l'-1}$$

where l' denotes the number of squeezing iterations, that is $l' = \lceil n/r' \rceil - 1$. More precisely, z_i is the r' leftmost bits of the internal state $S_{l'+i}$ and we have $S_{l'+i+1} = P_t(S_{l'+i})$ for $0 \leq i < l'$. If the hash output size is not a multiple of r' , one just truncates $z_{l'-1}$ to $n \bmod r'$ bits.

3.2 The internal permutations

We define here the internal permutations P_t , where $t \in \{100, 144, 196, 256, 288\}$. The internal state of the N_r -round permutation is viewed as a $(d \times d)$ matrix of s -bit cells and the corresponding values depending of t are given in Table 1. Note that we will always use a cell size of 4 bits, except for the largest version for which we use 8-bit cells, and that the number of rounds is always $N_r = 12$, whatever the value of t is. The internal state cell located at row i and column j is denoted $S[i, j]$ with $0 \leq i, j < d$.

Table 1. The parameters of the internal permutations P_t , together with the internal constants IC_d , the irreducible polynomials and the Z_i coefficients for the MixColumnsSerial computation.

	t	d	s	N_r	$IC_d(\cdot)$	irr. polynomial	Z_i coefficients
P_{100}	100	5	4	12	[0, 1, 3, 6, 4]	$x^4 + x + 1$	(1, 2, 9, 9, 2)
P_{144}	144	6	4	12	[0, 1, 3, 7, 6, 4]	$x^4 + x + 1$	(1, 2, 8, 5, 8, 2)
P_{196}	196	7	4	12	[0, 1, 2, 5, 3, 6, 4]	$x^4 + x + 1$	(1, 4, 6, 1, 1, 6, 4)
P_{256}	256	8	4	12	[0, 1, 3, 7, 15, 14, 12, 8]	$x^4 + x + 1$	(2, 4, 2, 11, 2, 8, 5, 6)
P_{288}	288	6	8	12	[0, 1, 3, 7, 6, 4]	$x^8 + x^4 + x^3 + x + 1$	(2, 3, 1, 2, 1, 4)

⁶ By sponge keying and using the security bound from [11], PHOTON-80/20/16 provides a secure 64-bit key MAC as long as the number of messages to be computed is lower than 2^{15} . For a secure 64-bit key MAC handling more messages (up to 2^{27}), one can for example go for a very similar PHOTON-80/8/8 version that also uses P_{100} . This version with capacity $c = 92$ would require the same area as PHOTON-80/20/16 but would be slower.

⁷ Please refer to Appendix E for a visualization.

One round is composed of four layers (see Figure 2): AddConstant (AC), SubCell (SC), ShiftRows (ShR) and MixColumnsSerial (MCS).

AddConstant. At round number v (starting the counting from 1), we first XOR a round constant $RC(v)$ to each cell $S[i, 0]$ of the first column of the internal state. Then, we XOR distinct internal constants $IC_d(i)$ to each cell $S[i, 0]$ of the same first column. Overall, for round v we have $S'[i, 0] = S[i, 0] \oplus RC(v) \oplus IC_d(i)$ for all $0 \leq i < d$. The round constants are $RC(v) = [1, 3, 7, 14, 13, 11, 6, 12, 9, 2, 5, 10]$. The internal constants depend on the square size d and on the row position i . They are given in Table 1. We give in the Appendix D all the constants for all square sizes, round numbers, row positions and how they have been generated.

SubCells. This layer simply applies an s -bit Sbox to each of the cells of the internal state, *i.e.* $S'[i, j] = \text{SBOX}(S[i, j])$ for all $0 \leq i, j < d$. In the case of 4-bit cells, we use the PRESENT Sbox SBOX_{PRE} [16] while for the 8-bit cells case we use the AES Sbox SBOX_{AES} [27]. Both are given in the Appendix B.

ShiftRows. As for the AES, for each row i this layer rotates all cells to the left by i column positions. Namely, $S'[i, j] = S[i, (j + i) \bmod d]$ for all $0 \leq i, j < d$.

MixColumnsSerial. The final mixing layer is applied to each of the columns of the internal state independently. For each column j input vector $(S[0, j], \dots, S[d - 1, j])^T$, we apply d times the matrix $A_t = \text{Serial}(Z_0, \dots, Z_{d-1})$. That is, for all $0 \leq j < d$:

$$(S'[0, j], \dots, S'[d - 1, j])^T = A_t^d \times (S[0, j], \dots, S[d - 1, j])^T$$

where the coefficients Z_0, \dots, Z_{d-1} are given in Table 1. In the case of 4-bit cells, the irreducible polynomial we chose is $x^4 + x + 1$, while for the 8-bit case we chose the AES one, *i.e.* $x^8 + x^4 + x^3 + x + 1$. The matrices A_t together with the overall MixColumnsSerial matrices A_t^d for each internal state size t are given in the Appendix C. Note that all A_t^d matrices are Maximum Distance Separable.⁸

4 Security Analysis

The sponge-like domain extension algorithm allows us to fully trust the security of the PHOTON hash functions as long as the internal permutation P_t does not present any structural flaw whatsoever (so-called “hermetic sponge strategy”). In other words, the cryptanalysis work is made easy since in order to cover any instantiation of PHOTON, one just has to study the security of the five internal permutations P_{100} , P_{144} , P_{196} , P_{256} and P_{288} . However, we do not have to consider an adversary bounded to 2^t computations, since for the flat sponge claim⁹ no resistance is claimed for attacks requiring a workload of more than $2^{c/2}$ operations. Thus, in case of PHOTON we only require the internal permutation to be indistinguishable from a randomly chosen permutation on the same domain $\{0, 1\}^t$ up to $2^{c/2}$ operations.

The PHOTON hash functions security is extremely conservative: we chose a number of rounds with a comfortable security margin such that this assumption on the internal permutations is fulfilled. However, even if an attacker could find a structural flaw on P_t , this is not likely to become an issue for the whole hash function. This argument is particularly true for our “small- r ” sponge-like shape. Indeed, the amount of freedom degrees available at the input of each internal permutation call during the absorbing phase is extremely small. Thus, even if a flaw is found for the internal permutation, the amount of freedom degrees is so thin that utilizing this flaw will very likely turn out to be intractable. The utilization of freedom degrees has always been one of the most powerful cryptanalyst tool (for MD-SHA family of hash functions [71, 70, 69] or even for sponge-like hash functions [60, 31]), thus reducing this ability as much as possible greatly increases the confidence in PHOTON’s security.

The PHOTON hash functions are very simple to analyze so as to make the cryptanalyst work as easy as possible. In the following subsections, we will study different attack scenarios for the internal permutations P_t , in which we consider the attacker has all freedom degrees possible, that is all t bits of internal state.

⁸ One could wonder why we did not propose a version with $d = 9$ and $s = 4$. The reason is that there is no matrix fulfilling the desired “serial MDS” properties for those parameters, whatever the irreducible polynomial chosen.

⁹ Flat sponge claim [8] with capacity c : the success probability of any attack should be smaller than or equal to the maximum of that for a random oracle and of $1 - e^{-N^2 \cdot 2^{-c+1}}$, with N the number of calls to the underlying function (or its inverse).

4.1 Differential/Linear cryptanalysis

PHOTON’s internal primitives are AES-like permutations. Therefore, by reusing the extensive work done in the past years on AES, it is very easy to compute a bound on the best differential path probability (where all differences on the input and output of all rounds are specified) or even the best differential probability (where only the input and output differences are specified).

We can obtain a bound on the number of active Sboxes (*i.e.* Sboxes with non-zero difference) for four rounds of the PHOTON internal permutation by simply adapting the wide-trail strategy [27] to our parameters: since our matrices underlying the MixColumnsSerial layer are MDS, at least $(d + 1)^2$ Sboxes will be active for any non-null differential path. Note that this bound is tight since we know four-round differential paths with $(d + 1)^2$ active Sboxes, as we will see in the next Section. As the best differential probability of the PRESENT Sbox is 2^{-2} , and for the AES Sbox it is 2^{-6} , we can directly deduce that the best differential path probability on four rounds of an internal permutation of PHOTON is upper bounded by $2^{-2 \cdot (d+1)^2}$ when $s = 4$ and by $2^{-6 \cdot (d+1)^2}$ when $s = 8$.

Since the diffusion in PHOTON is achieved with 2 rounds, a simple freedom degrees utilization (like in [25] for SHA-0) is likely to allow the control of two rounds. However, as shown in the next sections, more involved methods were recently introduced that allow to control up to three rounds. Even if we consider those methods can someday be pushed up to four rounds, this does not seem to endanger P_t since eight rounds provide at least $(d + 1) \cdot (d + 3)$ active Sboxes (eight consecutive rounds provide $2 \cdot (d + 1)^2$ active Sboxes, but when divided into two separate sub-paths they provide at least $(d + 1) \cdot (d + 3)$).

By adapting the work from [58], we can also show that the maximum differential probability for 4 rounds of a PHOTON internal permutation is upper bounded by

$$\max \left\{ \max_{1 \leq u \leq 2^s - 1} \sum_{j=1}^{2^s - 1} \{DP^S(u, j)\}^{d+1}, \max_{1 \leq u \leq 2^s - 1} \sum_{j=1}^{2^s - 1} \{DP^S(j, u)\}^{d+1} \right\}^d$$

where $DP^S(i, j)$ stands for the differential probability of the Sbox to map the difference i to j .

The duality between linear and differential attacks allows us to apply the same approaches to compute a bound on the best linear approximation or even the best linear hulls. We summarize in Table 2 the upper bounds on the best differential path probability, the best differential probability, the best linear approximation probability and the best linear hull probability for four rounds of the five PHOTON permutations. Note that such a reasoning assumes that random subkeys are added each round in order to make the Sboxes inputs independant. Yet, in the case of PHOTON the subkeys are simulated by the round constants addition and thus these bounds give a very good indication of the quality of the PHOTON internal permutations with regard to linear and differential cryptanalysis.

Table 2. Upper bounds on the best differential path probability, best differential probability, best linear approximation probability and best linear hull probability for 4 rounds and for the full version of the five PHOTON internal permutations.

	P_{100}		P_{144}		P_{196}		P_{256}		P_{288}	
	4 rds	full	4 rds	full	4 rds	full	4 rds	full	4 rds	full
differential path probability	2^{-72}	2^{-216}	2^{-98}	2^{-294}	2^{-128}	2^{-384}	2^{-162}	2^{-486}	2^{-294}	2^{-882}
differential probability	2^{-50}		2^{-72}		2^{-98}		2^{-128}		2^{-246}	
linear approx. probability	2^{-72}	2^{-216}	2^{-98}	2^{-294}	2^{-128}	2^{-384}	2^{-162}	2^{-486}	2^{-294}	2^{-882}
linear hull probability	2^{-50}		2^{-72}		2^{-98}		2^{-128}		2^{-246}	

Another very important security argument is also applicable to PHOTON: the internal permutations are fixed and allow no key input. Of course, this comes at a cost of lowered efficiency, but on the other hand it avoids any attack leveraging a weakness in the key schedule. This is particularly important as it was shown that such a control can lead for example to theoretical attacks against AES-192 and AES-256 in the related-key model [13, 12], or distinguishers for the full WHIRLPOOL compression function [45].

4.2 Rebound and Super-Sbox attacks

The original rebound attack [50] and its improved variants (start-from-the-middle attack [49] and Super-Sbox cryptanalysis [32, 45]) have been introduced very recently and provide the currently best known methods to analyze AES-like permutations in a hash functions setting, where no secret is involved. It is therefore very important to test what is the resistance of the PHOTON hash functions regarding those new tools. We provide in this section only an overview of the application of the rebound and Super-Sbox attacks to the PHOTON internal permutations and we invite the reader to look for the original articles for more details.

At the present time, in order to distinguish a fixed-key AES-like permutation from an ideal one, the best results are obtained by using a non-full active differential path [65], while the freedom degrees will be utilized in a Super-Sbox fashion [32, 45]. This allows to reach 8 rounds and we give in Figure 3 the differential path considered. Note that this trail is actually meeting the bound on the minimal number of active Sboxes for four rounds (round 3 to round 6), thus confirming its quality. Moreover, this trail is perfectly fit for a rebound-kind of attack, since most of its complexity is located at the same place, right in the middle of the path, which allows the attacker to concentrate all available freedom degrees on this precise part.




Fig. 3. 8-round differential path for PHOTON internal permutations, where the black cells represent active cells (*i.e.* containing a difference) while the white ones stand for inactive cells. We depict here the case $d = 8$, $n_L^{in} = n_L^{out} = n_R^{out} = 5$ and $n_R^{in} = 4$.

Here we slightly extend the non-full linear framework from [65]: we call n_L^{in} the number of active Sboxes after application of the first round and n_L^{out} after the second in the path from Figure 3. Identically, we denote n_R^{in} and n_R^{out} the number of active Sboxes after application of the fifth and sixth rounds respectively. Because of the MDS property of the mixing layer, we have the constraints that $n_L^{out} + n_R^{in} \geq d + 1$, $n_L^{in} + n_L^{out} \geq d + 1$ and $n_R^{in} + n_R^{out} \geq d + 1$. By using the freedom degrees counting method from [32], we can check that $2^{s(n_L^{in} + n_L^{out} + n_R^{in} + n_R^{out} - 2d)}$ solution pairs can be found for this differential path. This is always greater than one if we fulfill the previous constraints. Moreover, with the tweaked Super-Sbox method [65], one can find a solution with one operation on average for the 3 controlled round in the middle (depicted with dashed lines) and with a very limited minimal cost of 2^s computations and memory. Overall, obtaining a solution for the complete path will cost the attacker $2^{s(2d - n_L^{in} - n_R^{out})}$ operations because of the uncontrolled differential transitions in round 2 and 6. Considering the limited-birthday paradox problem [32], with a generic algorithm (modeling the permutation as a black box) one can find a solution pair with the same input/output properties with $\max\{2^{sd(d - n_L^{in})/2}, 2^{sd(d - n_R^{out})/2}, 2^{sd(d - n_L^{in} - n_R^{out})}\}$ computations.

We choose the best n_L^{in} , n_L^{out} , n_R^{in} , n_R^{out} values in order to minimize the distinguishing attack cost and the results are given in Table 3. As long as they satisfy the constraints, n_L^{out} and n_R^{in} have no impact on the complexity except for the controlled part (which we always ensure has a minimal cost of 2^s). Moreover, we observed that the best results were obtained when $n_L^{in} = n_R^{out}$. Overall, all versions of PHOTON provide a very comfortable security margin against state-of-the-art rebound-style attacks since only 8 rounds can be reached. Though the attacks complexities are low for 8 rounds, they can not be extended as is to more rounds, even if we allow more computation power to the attacker. Note that there is no key input for the PHOTON internal permutation and this avoids any potential improvement based on a weak key schedule or by the increased amount of freedom degrees provided by the key. Moreover, the recently introduced internal differential attack [61] that was applied on Grøst1 can not work here since only one single internal permutation is used (whereas the internal differential attack looks for properties between two very related permutations).

Table 3. Computation and memory complexities of the application of rebound-like attacks on the internal permutations of PHOTON. The attacks allow to distinguish 8 rounds of the permutation from an ideal permutation of the same size and can not be extended as is to more rounds even with higher complexity.

	P_{100}	P_{144}	P_{196}	P_{256}	P_{288}
n_L^{in}, n_R^{out}	4	5	6	7	5
computations	2^8	2^8	2^8	2^8	2^{16}
memory	2^4	2^4	2^4	2^4	2^8
generic	2^{10}	2^{12}	2^{14}	2^{16}	2^{24}

4.3 Cube testers and algebraic attacks

We applied the most recent developed cube testers [3] and its zero-sum distinguishers to the PHOTON permutations, the best we could find within practical time complexity is at most 3 rounds for all PHOTON variants. Note, in case of AES, “zero-sum” property is also referred as “balanced”, found by the AES designers [26], in which 3-round balanced property is shown. To the best of our knowledge, there is no balanced property found for more than 3 AES rounds.

There are two types of Sboxes used in PHOTON, the PRESENT Sbox and the AES Sbox, having an algebraic degree of $a = 3$ and $a = 7$ respectively. We showed in the previous section that the number of active Sbox is at least $(d + 1)^2$ for any consecutive 4-round trail of PHOTON. One can easily check that $a \cdot (d + 1)^2 \cdot 3 \gg t$ for all PHOTON variants. However, better bounds on the algebraic degree were recently published [20] and are summarized in Table 4 for the five PHOTON internal permutations (the P_{288} bounds are refined using a Super-Sbox view of the permutation [19]). Overall, one can construct zero-sum partitions of size 2^{98} for full P_{100} , 2^{138} for full P_{144} , 2^{184} for full P_{196} , 2^{237} for full P_{256} and 2^{283} for 8-round reduced P_{288} . However, all these zero-sum partition sizes are much above the claimed attack complexities for the PHOTON variants. For a size of the zero-sums partitions comparable to half the internal state size (which is the maximal possible flat sponge claim), not more than 8 rounds can be reached.

Table 4. Bounds on the algebraic degree from [20] applied to the PHOTON internal permutations, depending on the number of rounds considered.

number of rounds	1	2	3	4	5	6	7	8	9
P_{100}	3	9	27	75	91	97	99	99	99
P_{144}	3	9	27	81	123	137	141	143	143
P_{196}	3	9	27	81	157	183	191	194	195
P_{256}	3	9	27	81	197	236	249	253	255
P_{288}	7	42	252	282	287	287	287	287	287

Concerning algebraic attacks, the PRESENT Sbox is described by $e = 21$ quadratic equations in the $v = 8$ input/output-bit variables over $GF(2)$, while the AES Sbox is described by $e = 40$ quadratic equations in the $v = 16$ input/output-bit variables over $GF(2)$. The entire system for the internal permutations of PHOTON therefore consists of $(d^2 \cdot N_r \cdot e)$ quadratic equations in $(d^2 \cdot N_r \cdot v)$ variables. For example, in the case of P_{144} used for PHOTON-128/16/16, we end up with 9072 equations in 3456 variables. In comparison, the entire system for a fixed-key AES permutation consists of 6400 equations in 2560 variables. While the applicability of algebraic attacks on AES remains unclear, those numbers tend to indicate that PHOTON offers a comparable level of protection.

4.4 Other cryptanalysis

The slide attack is originally a block cipher cryptanalysis technique [14], but was recently applied to sponge-like hash functions [34]. The idea is to exploit the degree of self-similarity of a permutation. In the case of PHOTON, all rounds of the internal permutation are made different thanks to the round-dependent constants addition. Thus the slide attack is impossible to perform at the permutation level. Moreover, the slide attack

at the operating mode level from [34] is impossible to apply here since the padding rule from PHOTON forces the last message block to be different from zero (which prevent any sliding event).

Rotational cryptanalysis [41] was proven to be quite successful against Addition-Rotation-XOR (ARX) primitives and nice advances were made on some SHA-3 candidates [42]. The idea is to study the evolution of a rotated variant of some input words through the round process. However, PHOTON is an Sbox-oriented hash function and any rotation property in a cell will be directly removed by the application of the Sbox layer. One could look for rotation of cell positions in the internal state, but this is unlikely to lead to an attack since the constants used in a PHOTON round are all distinct and any position rotation property between columns or lines is removed after the application of two rounds.

Integral attacks are quite efficient against AES-based permutations and one can directly adapt the known-key variant from [43] to the PHOTON internal permutation cases. However, those attacks can only reach seven rounds with complexity $2^{s(2d-1)}$, which is worse than what can be obtained with rebound-style attacks from Section 4.2.

5 Performances and Comparison

Before we detail the hardware architectures and the optimizations done, we first describe the tools used. Finally we compare our results to previous work.

5.1 Design flow

We used *Mentor Graphics ModelSimXE 6.4b* and *Synopsys DesignCompiler A-2007.12-SP1* for functional simulation and synthesis of the designs to the *Virtual Silicon (VST)* standard cell library *UMCL18G212T3* [68], which is based on the *UMC L180 0.18 μ m 1P6M* logic process with a typical voltage of 1.8 V. We used *Synopsys Power Compiler* version *A-2007.12-SP1* to estimate the power consumption of our ASIC implementations. For synthesis and for power estimation we advised the compiler to keep the hierarchy and use a clock frequency of 100 KHz. Note that the wire-load model used, though it is the smallest available for this library, still simulates the typical wire-load of a circuit with a size of around 10 000 GE.

5.2 Architectures

To substantiate our claims on the hardware efficiency of our PHOTON family, we have implemented the flavors specified in Section 3 in VHDL and simulated their post-synthesis performance. We designed two architectures: one is fully serialized, *i.e.* performing operations on one cell per clock cycle, and aims for the smallest area possible; the second one is a d times parallelization of the first architecture, thus performing operations on one row in one clock cycle, resulting in a significant speed-up. As can be seen in Figure 4, our serialized design consists of six modules: **MCS**, **State**, **IO**, **AC**, **SC**, and **Controller**.

IO allows to 1) initialize our implementation with an all ‘0’ vector, 2) input the IV, 3) absorb message chunks, and 4) forward the output of the **State** module to the **AC** module without further modification. Instead of using two Multiplexer and an XOR gate, we used two NAND and one XOR gate thereby reducing the gate count required from $s \cdot 7.33$ to $s \cdot 4.67$ GE.

State comprises a $d \cdot d$ array of flip-flop cells storing s bits each. Every row constitutes a shift-register using the output of the last stage, *i.e.* column 0, as the input to the first stage (column $d - 1$) of the same row and the next row. Using this feedback functionality ShiftRows can be performed in $d - 1$ clock cycles with no additional hardware costs. Further, since MixColumnsSerial is performed on column 0, also a vertical shifting direction is required for this column. Consequently, columns 0 and $d - 1$ consist of flip-flop cells with two inputs (6 GE), while columns 1 to $d - 2$ consist of flip-flop cells with only one input (4.67 GE). The overall gate count for this module is $s \cdot d \cdot ((d - 2) \cdot 4.67 + 2 \cdot 6)$ GE and for all flavors it occupies the majority of the area required (between 65 and 77.5%).

MCS calculates the last row of A_t in one clock cycle. The result is stored in the **State** module, that is in the last row of column 0, which has been shifted upwards at the same time. Consequently, after d clock cycles the MixColumnsSerial operation is applied to an entire column. Then the whole state array is rotated by one position to the left and the next column is processed. In total $d \cdot (d + 1)$ clock cycles are required to perform MCS. As an example of the hardware efficiency of MCS we depict A_{100} in the upper and its sub-components in the lower right part of Figure 4. Using our library, for a multiplication by 2, 4 and 8, we need 2.67 GE,

4.67 GE, and 7 GE when using the irreducible polynomial $x^4 + x + 1$, respectively. Therefore the choice of the coefficients has only a minor impact on the overall gate count, as the majority is required to sum up the intermediate results. For example, in the case of A_{100} , 56 out of 75.33 GE are required for the XOR sum. The gate counts for the other matrices are: 80 GE, 99 GE, 145 GE, and 144 GE for A_{144} , A_{196} , A_{256} , and A_{288} , respectively.

AC performs the AddConstant operation by XORing the sum of the round constant RC with the current internal constant IC . Furthermore, since AC is only applied to the first column, the input to the XNOR gate is gated with a NAND gate. Instead of using an AND gate in combination with an XOR gate, our approach allows to reduce the area required from $s \cdot 6.67$ to $s \cdot 6$ GE.

SC performs the SubCells operation and consists of a single instantiation of the corresponding Sbox. For $s = 4$ we used an optimized Boolean representation of the PRESENT Sbox, which only requires 22.33 GE and for $s = 8$ we used Canright's representation of the AES Sbox [24] which requires 233 GE. It takes $d \cdot d$ clock cycles to perform AddConstant and SubCells on the whole state.

Controller uses a Finite State Machine (FSM) to generate all control signals required. Furthermore, also the round constants and the internal constants are generated within this module, as their values are used for the transition conditions of the FSM. The FSM consists of one idle state, one state for the combined execution of AC and SC, $d - 1$ states for ShR and two states for MCS (one for processing one column and another one to rotate the whole state to the left). Naturally, its gate count varies depending on d : 197 GE, 210 GE, 235 GE, and 254 GE for $d = 5, 6, 7, 8$, respectively.

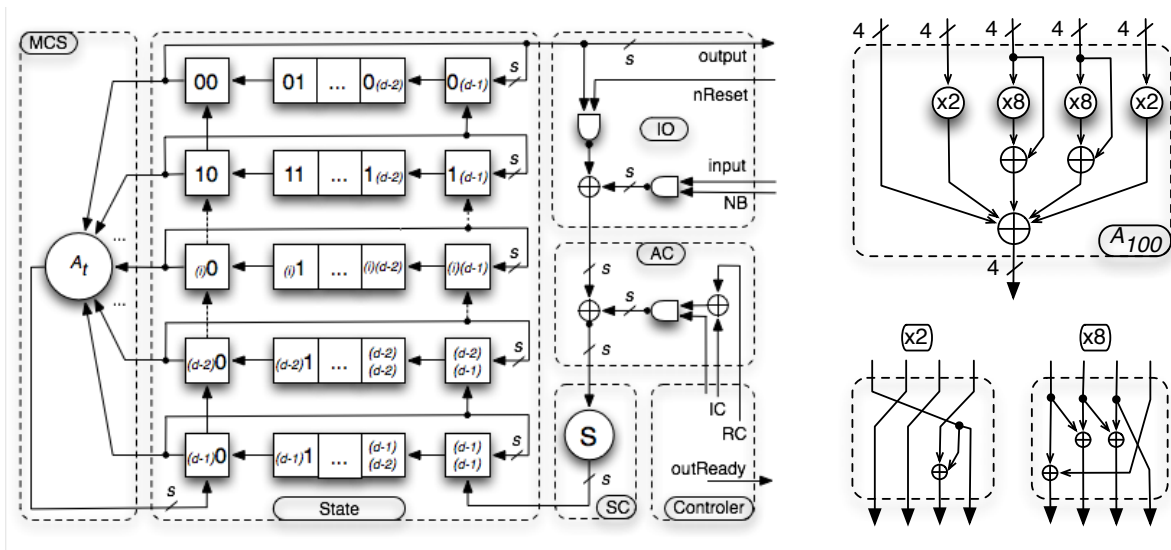


Fig. 4. Serial hardware architecture of PHOTON (left). As an example for its component A_t we also depict A_{100} with its sub-components (right).

5.3 Hardware results and comparison

We assume the message to be padded correctly and the IVs to be loaded at the beginning of the operation. Then it requires $d \cdot d + (d - 1) + d \cdot (d + 1)$ clock cycles to perform one round of the permutation P , resulting in a total latency of $12 \cdot (2 \cdot d \cdot (d + 1) - 1)$ clock cycles. Table 5 compares our results to previous works, sorted after preimage and collision resistance levels. Area requirements are provided in GE, while the latency is given in clock cycles for only the internal permutation P (or the internal block-cipher E), and the whole hash function H . Further metrics are Throughput in kbps and a Figure of Merit (FOM) proposed by [4]. In order to have a comparison for a best case scenario and a real-world application, we provide the latter two metrics for ‘long’ messages (omitting any padding influences) and for 96-bit messages, where we do take padding into account. In particular this means that a 96-bit message is padded with ‘1’ and as many ‘0’s as required. Furthermore Merkle-Damgård constructions need additional 64 bits to encode the message length.

The parameters n , c , r and r' stand for the hash output size, the capacity, the input bitrate and the output bitrate respectively. Finally, the column “Pre” gives the claimed preimage resistance security and “Col” the claimed collision resistance security.

Table 5. Overview of parameters, security level, and performance of several lightweight hash functions. Throughput and FOM figures have been derived at a clock frequency of 100 KHz. We marked by a * the preimage resistances of PHOTON-128/16/16, PHOTON-160/36/36, PHOTON-224/32/32 and PHOTON-256/32/32 in order to indicate that these PHOTON variants can achieve equal preimage resistance compared to its competitors by simply adding one more squeezing round. This will increase the hash output size n by r' bits and slightly reduce the throughput for small messages, while the area and the long message performances will remain the same.

Name	Ref.	Parameters				Security		Performance							
		n	c	r	r'	Pre	Col	Area [GE]	Latency [clk]		Throughput [kbps]		FOM [nb/clk/GE ²]		
									P/E	H	long	96-bit	long	96-bit	
64-bit preimage resistance															
SQUASH	[73]	64	x	x	x	64	0	2646	31800	31800	0.2	0.15	0.29	0.14	
DM-PRESENT-80	[17]	64	64	80	x	64	32	1600	547	547	14.63	5.85	57.13	19.04	
DM-PRESENT-80	[17]	64	64	80	x	64	32	2213	33	33	242.42	96.67	495.01	165.00	
DM-PRESENT-128	[17]	64	64	128	x	64	32	1886	559	559	22.90	8.59	64.37	32.19	
DM-PRESENT-128	[17]	64	64	128	x	64	32	2530	33	33	387.88	145.45	605.98	302.99	
KECCAK-f[200]	[40]	64	128	72	72	64	32	2520	900	900	8.00	5.33	12.6	8.4	
PHOTON-80/20/16		80	80	20	16	64	40	865	708	3540	2.82	1.51	37.73	20.12	
PHOTON-80/20/16		80	80	20	16	64	40	1168	132	660	15.15	8.08	111.13	59.27	
64-bit collision resistance															
U-QUARK	[4]	136	128	8	8	128	64	1379	544	9248	1.47	0.61	7.73	3.20	
U-QUARK	[4]	136	128	8	8	128	64	2392	68	1156	11.76	4.87	20.56	8.51	
H-PRESENT-128	[17]	128	128	64	x	128	64	2330	559	559	11.45	5.72	21.09	10.54	
H-PRESENT-128	[17]	128	128	64	x	128	64	4256	32	32	200.00	100.00	110.41	55.21	
ARMADILLO2-B	[6]	128	128	64	x	128	64	4353	256	256	25.00	12.50	13.19	6.60	
ARMADILLO2-B	[6]	128	128	64	x	128	64	6025	64	64	100.00	50.00	27.55	13.77	
KECCAK-f[400]	[40]	128	256	144	144	128	64	5090	1000	1000	14.40	9.60	5.56	3.71	
PHOTON-128/16/16		128	128	16	16	112*	64	1122	996	7968	1.61	0.69	12.78	5.48	
PHOTON-128/16/16		128	128	16	16	112*	64	1708	156	1248	10.26	4.4	35.15	15.06	
80-bit collision resistance															
D-QUARK	[4]	176	160	16	16	160	80	1702	704	7744	2.27	0.80	7.85	2.77	
D-QUARK	[4]	176	160	16	16	160	80	2819	88	968	18.18	6.42	22.88	8.08	
ARMADILLO2-C	[6]	160	160	80	x	160	80	5406	320	320	25.00	10.00	8.55	3.42	
ARMADILLO2-C	[6]	160	160	80	x	160	80	7492	80	80	100.00	40.00	17.82	7.13	
SHA-1	[56]	160	160	512	x	160	80	5527	344	344	148.84	27.91	48.72	9.14	
PHOTON-160/36/36		160	160	36	36	124*	80	1396	1332	6660	2.70	1.03	13.87	5.28	
PHOTON-160/36/36		160	160	36	36	124*	80	2117	180	900	20	7.62	44.64	17.01	
112-bit collision resistance															
S-QUARK	[4]	256	224	32	32	224	112	2296	1024	8192	3.13	0.85	5.93	1.62	
S-QUARK	[4]	256	224	32	32	224	112	4640	64	512	50.00	13.64	23.22	6.33	
PHOTON-224/32/32		224	224	32	32	192*	112	1736	1716	12012	1.86	0.56	6.19	1.86	
PHOTON-224/32/32		224	224	32	32	192*	112	2786	204	1428	15.69	4.71	20.21	6.06	
128-bit collision resistance															
ARMADILLO2-E	[6]	256	256	128	x	256	128	8653	512	512	25.00	9.38	3.34	1.25	
ARMADILLO2-E	[6]	256	256	128	x	256	128	11914	128	128	100.00	37.50	7.05	2.64	
SHA-2	[30]	256	256	512	x	256	128	10868	1128	1128	45.39	8.51	3.84	0.72	
PHOTON-256/32/32		256	256	32	32	224*	128	2177	996	7968	3.21	0.88	6.78	1.85	
PHOTON-256/32/32		256	256	32	32	224*	128	4362	156	1248	20.51	5.59	10.78	2.94	

As can be seen, our proposals compete well in terms of area requirements, since they are 18% to 75% smaller compared to previous proposals with a similar preimage/collision resistance level. For a smaller area, the throughput of PHOTON variants is comparable to the QUARK proposals¹⁰. Alternatively, for a similar area, PHOTON variants are much faster than the QUARK proposals. This can be observed in the Figure of Merit column of the results Table. One could argue that the throughput of two proposals can not be compared because the security margin is not taken in account. However, we would like to emphasize that the security margin is very hard to measure as it greatly depends on the simplicity of the scheme, the amount of work spent by the cryptanalysts, etc. Unlike most of the lightweight hash functions proposed, in the case of PHOTON, we chose very simple to analyse internal permutations, thus directly leveraging the extensive analysis work already known for AES-like permutations. While 8 rounds over 12 of the internal permutations of PHOTON can be distinguished from a random permutation, we provide strong arguments that this is very unlikely to be much improved.

We did not include power figures in Table 5 for several reasons. First, the power consumption strongly depends on the technology used and cannot be compared between different technologies in a fair manner. Furthermore, simulated power figures strongly depend on the simulation method used, and the effort spent. Instead, we just briefly list the simulated power figures for our proposals here: 1.59, 2.29, 2.74, 4.01, and 4.55 μ W for serialized implementation of PHOTON-80/20/16, PHOTON-128/16/16, PHOTON-160/36/36, PHOTON-224/32/32, and PHOTON-256/32/32, respectively. The d -parallel implementations require 2.7, 3.45, 4.35, 6.5, and 8.38 μ W, respectively. This let us conclude that all PHOTON flavors seem to be suitable for ultra-constrained devices, such as passive RFID tags, which was one of our initial design goals.

5.4 Software implementation

We give in Table 6 our software implementation performances for the PHOTON variants. The processor used for the benchmarks is an Intel(R) Core(TM) i7 CPU Q 720 clocked at 1.60GHz. For comparison purposes, we also benchmarked the speed of an AES permutation (without the key schedule) and a modified version of it with a serially computable MDS matrix instead (the 4×4 matrix A given in Section 2.2). As expected, the table-based implementations reach the same speed for both versions. We also benchmarked other lightweight hash function designs. QUARK reference code [4], very likely to be optimizable, runs at 8k, 30k and 22k cycles per byte for U-QUARK, D-QUARK and S-QUARK, respectively. The optimized PRESENT code [16] runs at 90 cycles per byte, hence the estimate speed for DM-PRESENT-80, DM-PRESENT-128 and H-PRESENT-128 are 72, 45 and 90 cycles per byte, respectively.

Table 6. Software performances in cycles per byte of the PHOTON variants for long messages.

PHOTON-80/20/16	PHOTON-128/16/16	PHOTON-160/36/36	PHOTON-224/32/32	PHOTON-256/32/32
95 c/B	156 c/B	116 c/B	227 c/B	157 c/B

6 Conclusion

We proposed PHOTON, the most lightweight hash-function family known so far, very close to the theoretical optimum. Our proposal is based on the well known AES design strategy, but we introduced a new mixing layer building method that is perfectly fit for small area scenarios. This allows us to directly leverage the extensive work done on AES and AES-like hash functions so as to provide good confidence in the security of our scheme. Finally, PHOTON is not only the smallest hash function, but it also achieves excellent area/throughput trade-offs and we obtained very acceptable performances with simple software implementations.

¹⁰ We synthesized the publicly available VHDL source code of U-QUARK using the same tool chain and ASIC library as for our proposals. The post-synthesis figures for U-QUARK are slightly higher than the previously published ones, *i.e.* 1400 GE instead of 1379 GE, which indicates that PHOTONs smaller footprint is not caused by a different tool chain. However, for comparison we took the previously available figures, which is in favour of QUARK.

Acknowledgement

The authors would like to thank the anonymous referees for their helpful comments. Also, we are very grateful to Dag Arne Osvik and AlpCode for providing an optimized Boolean representation of the PRESENT Sbox, to Jean-Philippe Aumasson for providing his cube testers source code and to Christina Boura for her help with zero-sum distinguishers.

References

1. Masayuki Abe, editor. *Advances in Cryptology - ASIACRYPT 2010 - 16th International Conference on the Theory and Application of Cryptology and Information Security, Singapore, December 5-9, 2010. Proceedings*, volume 6477 of *LNCS*. Springer, 2010.
2. Elena Andreeva, Bart Mennink, and Bart Preneel. The Parazoa Family: Generalizing the Sponge Hash Functions. Cryptology ePrint Archive, Report 2011/028, 2011.
3. Jean-Philippe Aumasson, Itai Dinur, Willi Meier, and Adi Shamir. Cube Testers and Key Recovery Attacks on Reduced-Round MD6 and Trivium. In Orr Dunkelman, editor, *FSE*, volume 5665 of *LNCS*, pages 1–22. Springer, 2009.
4. Jean-Philippe Aumasson, Luca Henzen, Willi Meier, and María Naya-Plasencia. Quark: A Lightweight Hash. In Mangard and Standaert [47], pages 1–15. <http://131002.net/quark/>.
5. Gildas Avoine and Philippe Oechslin. A Scalable and Provably Secure Hash-Based RFID Protocol. In *PerCom Workshops*, pages 110–114. IEEE Computer Society, 2005.
6. Stéphane Badel, Nilay Dagtekin, Jorge Nakahara, Khaled Ouafi, Nicolas Reffé, Pouyan Sepehrdad, Petr Susil, and Serge Vaudenay. ARMADILLO: A Multi-purpose Cryptographic Primitive Dedicated to Hardware. In Mangard and Standaert [47], pages 398–412.
7. Guido Bertoni, Joan Daemen, Michael Peeters, and Gilles Van Assche. Sponge functions. Ecrypt Hash Workshop 2007, May 2007.
8. Guido Bertoni, Joan Daemen, Michael Peeters, and Gilles Van Assche. On the Indifferentiability of the Sponge Construction. In Paterson [59], pages 181–197.
9. Guido Bertoni, Joan Daemen, Michael Peeters, and Gilles Van Assche. Keccak specifications. Submission to NIST (Round 2), 2009.
10. Guido Bertoni, Joan Daemen, Michael Peeters, and Gilles Van Assche. Sponge-Based Pseudo-Random Number Generators. In Mangard and Standaert [47], pages 33–47.
11. Guido Bertoni, Joan Daemen, Michael Peeters, and Gilles Van Assche. On the security of the keyed sponge construction. In G. Leander and S.S. Thomsen, editors, *SKEW*, 2011.
12. Alex Biryukov and Dmitry Khovratovich. Related-Key Cryptanalysis of the Full AES-192 and AES-256. In Matsui [48], pages 1–18.
13. Alex Biryukov, Dmitry Khovratovich, and Ivica Nikolic. Distinguisher and Related-Key Attack on the Full AES-256. In Shai Halevi, editor, *CRYPTO*, volume 5677 of *LNCS*, pages 231–249. Springer, 2009.
14. Alex Biryukov and David Wagner. Slide Attacks. In Lars R. Knudsen, editor, *FSE*, volume 1636 of *LNCS*, pages 245–259. Springer, 1999.
15. Cline Blondeau, Mara Naya-Plasencia, Marion Videau, and Erik Zenner. Cryptanalysis of ARMADILLO2. Cryptology ePrint Archive, Report 2011/160, 2011.
16. Andrey Bogdanov, Lars R. Knudsen, Gregor Leander, Christof Paar, Axel Poschmann, Matthew J. B. Robshaw, Yannick Seurin, and C. Vikkelsoe. PRESENT: An Ultra-Lightweight Block Cipher. In Paillier and Verbauwhede [57], pages 450–466. <http://lightweightcrypto.org/present/>.
17. Andrey Bogdanov, Gregor Leander, Christof Paar, Axel Poschmann, Matthew J. B. Robshaw, and Yannick Seurin. Hash Functions and RFID Tags: Mind the Gap. In Elisabeth Oswald and Pankaj Rohatgi, editors, *CHES*, volume 5154 of *LNCS*, pages 283–299. Springer, 2008.
18. Wieb Bosma, John Cannon, and Catherine Playoust. The Magma algebra system. I. The user language. *J. Symbolic Comput.*, 24(3-4):235–265, 1997. Computational algebra and number theory (London, 1993).
19. Christina Boura. Private discussions, May 2011.
20. Christina Boura, Anne Canteaut, and Christophe De Cannière. Higher-order differential properties of Keccak and Luffa. In Antoine Joux, editor, *FSE*, LNCS. Springer, 2011 to appear.
21. Gilles Brassard, editor. *Advances in Cryptology - CRYPTO '89, 9th Annual International Cryptology Conference, Santa Barbara, California, USA, August 20-24, 1989, Proceedings*, volume 435 of *LNCS*. Springer, 1990.
22. Christophe De Cannière, Orr Dunkelman, and Miroslav Knezevic. KATAN and KTANTAN - A Family of Small and Efficient Hardware-Oriented Block Ciphers. In Christophe Clavier and Kris Gaj, editors, *CHES*, volume 5747 of *LNCS*, pages 272–288. Springer, 2009.
23. Christophe De Cannière and Bart Preneel. Trivium. In Robshaw and Billet [63], pages 244–266.

24. David Canright. A Very Compact S-Box for AES. In Josyula R. Rao and Berk Sunar, editors, *CHES*, volume 3659 of *LNCS*, pages 441–455. Springer, 2005. The HDL specification is available at the author’s official webpage <http://faculty.nps.edu/drcanrig/pub/index.html>.
25. Florent Chabaud and Antoine Joux. Differential Collisions in SHA-0. In Hugo Krawczyk, editor, *CRYPTO*, volume 1462 of *LNCS*, pages 56–71. Springer, 1998.
26. Joan Daemen and Vincent Rijmen. AES Proposal: Rijndael. NIST AES proposal, 1998.
27. Joan Daemen and Vincent Rijmen. *The Design of Rijndael: AES - The Advanced Encryption Standard*. Springer, 2002.
28. Ivan Damgård. A Design Principle for Hash Functions. In Brassard [21], pages 416–427.
29. Orr Dunkelman, editor. *Fast Software Encryption, 16th International Workshop, FSE 2009, Leuven, Belgium, February 22-25, 2009, Revised Selected Papers*, volume 5665 of *LNCS*. Springer, 2009.
30. Martin Feldhofer and Christian Rechberger. A Case Against Currently Used Hash Functions in RFID Protocols. In Robert Meersman, Zahir Tari, and Pilar Herrero, editors, *OTM Workshops (1)*, volume 4277 of *LNCS*, pages 372–381. Springer, 2006.
31. Thomas Fuhr and Thomas Peyrin. Cryptanalysis of RadioGatún. In Dunkelman [29], pages 122–138.
32. Henri Gilbert and Thomas Peyrin. Super-Sbox Cryptanalysis: Improved Attacks for AES-Like Permutations. In Hong and Iwata [38], pages 365–383.
33. Tim Good and Mohammed Benaissa. ASIC Hardware Performance. In Robshaw and Billet [63], pages 267–293.
34. Michael Gorski, Stefan Lucks, and Thomas Peyrin. Slide Attacks on a Class of Hash Functions. In Josef Pieprzyk, editor, *ASIACRYPT*, volume 5350 of *LNCS*, pages 143–160. Springer, 2008.
35. Martin Hell, Thomas Johansson, Alexander Maximov, and Willi Meier. The Grain Family of Stream Ciphers. In Robshaw and Billet [63], pages 179–190.
36. Dirk Henrici, Joachim Götze, and Paul Müller. A Hash-based Pseudonymization Infrastructure for RFID Systems. In *SecPerU*, pages 22–27. IEEE Computer Society, 2006.
37. Shoichi Hirose. Some Plausible Constructions of Double-Block-Length Hash Functions. In Matthew J. B. Robshaw, editor, *FSE*, volume 4047 of *LNCS*, pages 210–225. Springer, 2006.
38. 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 *LNCS*. Springer, 2010.
39. Ari Juels and Stephen A. Weis. Authenticating Pervasive Devices with Human Protocols. In Shoup [67], pages 293–308.
40. Elif Bilge Kavun and Tolga Yalcin. A Lightweight Implementation of Keccak Hash Function for Radio-Frequency Identification Applications. In Siddika Berna Ors Yalcin, editor, *RFIDSec*, volume 6370 of *LNCS*, pages 258–269. Springer, 2010.
41. Dmitry Khovratovich and Ivica Nikolic. Rotational Cryptanalysis of ARX. In Hong and Iwata [38], pages 333–346.
42. Dmitry Khovratovich, Ivica Nikolic, and Christian Rechberger. Rotational Rebound Attacks on Reduced Skein. In Abe [1], pages 1–19.
43. Lars R. Knudsen and Vincent Rijmen. Known-Key Distinguishers for Some Block Ciphers. In Kurosawa [44], pages 315–324.
44. Kaoru Kurosawa, editor. *Advances in Cryptology - ASIACRYPT 2007, 13th International Conference on the Theory and Application of Cryptology and Information Security, Kuching, Malaysia, December 2-6, 2007, Proceedings*, volume 4833 of *LNCS*. Springer, 2007.
45. Mario Lamberger, Florian Mendel, Christian Rechberger, Vincent Rijmen, and Martin Schläffer. Rebound Distinguishers: Results on the Full Whirlpool Compression Function. In Matsui [48], pages 126–143.
46. Su-Mi Lee, Young Ju Hwang, Dong Hoon Lee, and Jong In Lim. Efficient Authentication for Low-Cost RFID Systems. In Osvaldo Gervasi, Marina L. Gavrilova, Vipin Kumar, Antonio Laganà, Heow Pueh Lee, Youngsong Mun, David Taniar, and Chih Jeng Kenneth Tan, editors, *ICCSA (1)*, volume 3480 of *LNCS*, pages 619–627. Springer, 2005.
47. Stefan Mangard and François-Xavier Standaert, editors. *Cryptographic Hardware and Embedded Systems, CHES 2010, 12th International Workshop, Santa Barbara, CA, USA, August 17-20, 2010. Proceedings*, volume 6225 of *LNCS*. Springer, 2010.
48. Mitsuru Matsui, editor. *Advances in Cryptology - ASIACRYPT 2009, 15th International Conference on the Theory and Application of Cryptology and Information Security, Tokyo, Japan, December 6-10, 2009. Proceedings*, volume 5912 of *LNCS*. Springer, 2009.
49. Florian Mendel, Thomas Peyrin, Christian Rechberger, and Martin Schläffer. Improved Cryptanalysis of the Reduced Grøstl Compression Function, ECHO Permutation and AES Block Cipher. In Michael J. Jacobson Jr., Vincent Rijmen, and Reihaneh Safavi-Naini, editors, *Selected Areas in Cryptography*, volume 5867 of *LNCS*, pages 16–35. Springer, 2009.
50. Florian Mendel, Christian Rechberger, Martin Schläffer, and Søren S. Thomsen. The Rebound Attack: Cryptanalysis of Reduced Whirlpool and Grøstl. In Dunkelman [29], pages 260–276.
51. Ralph C. Merkle. One Way Hash Functions and DES. In Brassard [21], pages 428–446.
52. Amir Moradi, Axel Poschmann, San Ling, Chrstof Paar, and Huaxiong Wang. Pushing the Limits: A Very Compact and a Threshold Implementation of the AES. In Paterson [59].

53. National Institute of Standards and Technology. FIPS 180-1: Secure Hash Standard. <http://csrc.nist.gov>, April 1995.
54. National Institute of Standards and Technology. FIPS 180-2: Secure Hash Standard. <http://csrc.nist.gov>, August 2002.
55. National Institute of Standards and Technology. Announcing Request for Candidate Algorithm Nominations for a New Cryptographic Hash Algorithm (SHA-3) Family. *Federal Register*, 27(212):62212–62220, November 2007. Available: http://csrc.nist.gov/groups/ST/hash/documents/FR_Notice_Nov07.pdf (2008/10/17).
56. M. O’Neill. Low-Cost SHA-1 Hash Function Architecture for RFID Tags. In Sandra Dominikus and Manfred Aigner, editors, *RFIDSec*, 2008. Available via <http://events.iaik.tugraz.at/RFIDSec08/Papers/>.
57. Pascal Paillier and Ingrid Verbauwhede, editors. *Cryptographic Hardware and Embedded Systems - CHES 2007, 9th International Workshop, Vienna, Austria, September 10-13, 2007, Proceedings*, volume 4727 of *LNCS*. Springer, 2007.
58. Sangwoo Park, Soo Hak Sung, Sangjin Lee, and Jongin Lim. Improving the Upper Bound on the Maximum Differential and the Maximum Linear Hull Probability for SPN Structures and AES. In Thomas Johansson, editor, *FSE*, volume 2887 of *LNCS*, pages 247–260. Springer, 2003.
59. 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 *LNCS*. Springer, 2011.
60. Thomas Peyrin. Cryptanalysis of Grindahl. In Kurosawa [44], pages 551–567.
61. Thomas Peyrin. Improved Differential Attacks for ECHO and Grøstl. In Tal Rabin, editor, *CRYPTO*, volume 6223 of *LNCS*, pages 370–392. Springer, 2010.
62. Thomas Peyrin, Henri Gilbert, Frédéric Muller, and Matthew J. B. Robshaw. Combining Compression Functions and Block Cipher-Based Hash Functions. In Xuejia Lai and Kefei Chen, editors, *ASIACRYPT*, volume 4284 of *LNCS*, pages 315–331. Springer, 2006.
63. Matthew J. B. Robshaw and Olivier Billet, editors. *New Stream Cipher Designs - The eSTREAM Finalists*, volume 4986 of *LNCS*. Springer, 2008.
64. Carsten Rolfes, Axel Poschmann, Gregor Leander, and Christof Paar. Ultra-Lightweight Implementations for Smart Devices - Security for 1000 Gate Equivalents. In Gilles Grimaud and François-Xavier Standaert, editors, *CARDIS*, volume 5189 of *LNCS*, pages 89–103. Springer, 2008.
65. Yu Sasaki, Yang Li, Lei Wang, Kazuo Sakiyama, and Kazuo Ohta. Non-full-active Super-Sbox Analysis: Applications to ECHO and Grøstl. In Abe [1], pages 38–55.
66. Adi Shamir. SQUASH - A New MAC with Provable Security Properties for Highly Constrained Devices Such as RFID Tags. In Kaisa Nyberg, editor, *FSE*, volume 5086 of *LNCS*, pages 144–157. Springer, 2008.
67. 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 *LNCS*. Springer, 2005.
68. Virtual Silicon Inc. 0.18 μm VIP Standard Cell Library Tape Out Ready, Part Number: UMCL18G212T3, Process: UMC Logic 0.18 μm Generic II Technology: 0.18 μm , July 2004.
69. Xiaoyun Wang, Yiqun Lisa Yin, and Hongbo Yu. Finding Collisions in the Full SHA-1. In Shoup [67], pages 17–36.
70. Xiaoyun Wang and Hongbo Yu. How to Break MD5 and Other Hash Functions. In *EUROCRYPT*, pages 19–35, 2005.
71. Xiaoyun Wang, Hongbo Yu, and Yiqun Lisa Yin. Efficient Collision Search Attacks on SHA-0. In Shoup [67], pages 1–16.
72. Hirotaka Yoshida, Dai Watanabe, Katsuyuki Okeya, Jun Kitahara, Hongjun Wu, Özgül Küçük, and Bart Preneel. MAME: A Compression Function with Reduced Hardware Requirements. In Paillier and Verbauwhede [57], pages 148–165.
73. Serge Zhilyaev. Evaluating a new MAC for current and next generation RFID. Master’s thesis, University of Massachusetts Amherst, 2010, available via <http://scholarworks.umass.edu/cgi/viewcontent.cgi?article=1477&context=theses>.

A Process for Choosing Optimal Extended Sponge Framework Parameters

We describe here the process to follow in order to pick the appropriate extended sponge framework parameters while maximizing the throughput of the PHOTON variant created. The process takes as input:

- COL: the minimum collision resistance required in \log_2
- PRE: the minimum preimage resistance required in \log_2
- SPRE: the minimum second-preimage resistance required in \log_2
- a, k, N : no better attack than exhaustive search when used as a MAC of at least N -bit with a k -bit key and limited to a queries in \log_2

and outputs the parameters t , n , c , r and r' .

We denote $c_{min} = \max\{2 \cdot \text{COL}, 2 \cdot \text{SPRE}, \text{PRE}, k + a + 1\}$ the minimal capacity and $n_{min} = \max\{2 \cdot \text{COL}, \text{SPRE}, \text{PRE}, N\}$ the minimal output size. First, among the five available PHOTON permutations P_t , choose the smallest one such that $t > \max\{c_{min}, N\}$ (otherwise the security requirements can not be fulfilled). Note that allowing a bigger t will lead to faster versions, but bigger area. We set $c = c_{min}$ and $r = t - c$ in order to maximize the speed by reducing the capacity to the minimum allowable. Finally, if $c_{min} \geq 2 \cdot \text{PRE}$ we set $n = r' = n_{min}$, otherwise we can choose any value n and r' verifying $t \geq n \geq n_{min}$ and $n - r' \geq \text{PRE}$ which is a trade-off between small message speed and hash output size.

As an example, using this process with inputs $\text{COL} = 0$, $\text{PRE} = 64$, $\text{SPRE} = 0$, $a = 15$, $k = 64$, $N = 64$, we obtain the smallest PHOTON variant proposed, PHOTON-80/20/16 with $t = 100$, $n = 80$, $c = 80$, $r = 20$ and $r' = 16$. Note that for each permutation P_t , we advice not to build PHOTON variants with much higher maximal complexity security claims than the ones proposed in this article (namely 2^{64} , 2^{112} , 2^{124} , 2^{192} and 2^{224} respectively).

B PHOTON Sboxes

For cells of $s = 4$ bits, the SubCells layer uses the PRESENT Sbox (in hexadecimal display):

$$\text{SBOX}_{\text{PRE}} = [0xc, 0x5, 0x6, 0xb, 0x9, 0x0, 0xa, 0xd, 0x3, 0xe, 0xf, 0x8, 0x4, 0x7, 0x1, 0x2].$$

For cells of $s = 8$ bits, the SubCells layer uses the AES Sbox (in hexadecimal display):

$$\begin{aligned} \text{SBOX}_{\text{AES}} = \\ & [0xca, 0x82, 0xc9, 0x7d, 0xfa, 0x59, 0x47, 0xf0, 0xad, 0xd4, 0xa2, 0xaf, 0x9c, 0xa4, 0x72, 0xc0, \\ & 0xb7, 0xfd, 0x93, 0x26, 0x36, 0x3f, 0xf7, 0xcc, 0x34, 0xa5, 0xe5, 0xf1, 0x71, 0xd8, 0x31, 0x15, \\ & 0x04, 0xc7, 0x23, 0xc3, 0x18, 0x96, 0x05, 0x9a, 0x07, 0x12, 0x80, 0xe2, 0xeb, 0x27, 0xb2, 0x75, \\ & 0x09, 0x83, 0x2c, 0x1a, 0x1b, 0x6e, 0x5a, 0xa0, 0x52, 0x3b, 0xd6, 0xb3, 0x29, 0xe3, 0x2f, 0x84, \\ & 0x53, 0xd1, 0x00, 0xed, 0x20, 0xfc, 0xb1, 0x5b, 0x6a, 0xcb, 0xbe, 0x39, 0x4a, 0x4c, 0x58, 0xcf, \\ & 0xd0, 0xef, 0xaa, 0xfb, 0x43, 0x4d, 0x33, 0x85, 0x45, 0xf9, 0x02, 0x7f, 0x50, 0x3c, 0x9f, 0xa8, \\ & 0x51, 0xa3, 0x40, 0x8f, 0x92, 0x9d, 0x38, 0xf5, 0xbc, 0xb6, 0xda, 0x21, 0x10, 0xff, 0xf3, 0xd2, \\ & 0xcd, 0x0c, 0x13, 0xec, 0x5f, 0x97, 0x44, 0x17, 0xc4, 0xa7, 0x7e, 0x3d, 0x64, 0x5d, 0x19, 0x73, \\ & 0x60, 0x81, 0x4f, 0xdc, 0x22, 0x2a, 0x90, 0x88, 0x46, 0xee, 0xb8, 0x14, 0xde, 0x5e, 0x0b, 0xdb, \\ & 0xe0, 0x32, 0x3a, 0x0a, 0x49, 0x06, 0x24, 0x5c, 0xc2, 0xd3, 0xac, 0x62, 0x91, 0x95, 0xe4, 0x79, \\ & 0xe7, 0xc8, 0x37, 0x6d, 0x8d, 0xd5, 0x4e, 0xa9, 0x6c, 0x56, 0xf4, 0xea, 0x65, 0x7a, 0xae, 0x08, \\ & 0xba, 0x78, 0x25, 0x2e, 0x1c, 0xa6, 0xb4, 0xc6, 0xe8, 0xdd, 0x74, 0x1f, 0x4b, 0xbd, 0x8b, 0x8a, \\ & 0x70, 0x3e, 0xb5, 0x66, 0x48, 0x03, 0xf6, 0x0e, 0x61, 0x35, 0x57, 0xb9, 0x86, 0xc1, 0x1d, 0x9e, \\ & 0xe1, 0xf8, 0x98, 0x11, 0x69, 0xd9, 0x8e, 0x94, 0x9b, 0x1e, 0x87, 0xe9, 0xce, 0x55, 0x28, 0xdf, \\ & 0x8c, 0xa1, 0x89, 0x0d, 0xbf, 0xe6, 0x42, 0x68, 0x41, 0x99, 0x2d, 0x0f, 0xb0, 0x54, 0xbb, 0x16]. \end{aligned}$$

C PHOTON Mixing Matrices

We use x^4+x+1 as the irreducible polynomial for multiplication in $\text{GF}(2^4)$, and this applies to all permutations with 4-bit cells, e.g., A_{100} , A_{144} , A_{196} , A_{256} . $x^8 + x^4 + x^3 + x + 1$, as in AES, is for multiplication in $\text{GF}(2^8)$ as used in A_{288} .

$$(A_{100})^5 = \begin{pmatrix} 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 1 & 2 & 9 & 9 & 2 \end{pmatrix}^5 = \begin{pmatrix} 1 & 2 & 9 & 9 & 2 \\ 2 & 5 & 3 & 8 & 13 \\ 13 & 11 & 10 & 12 & 1 \\ 1 & 15 & 2 & 3 & 14 \\ 14 & 14 & 8 & 5 & 12 \end{pmatrix}$$

$$\begin{aligned}
(A_{144})^6 &= \begin{pmatrix} 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ 1 & 2 & 8 & 5 & 8 & 2 \end{pmatrix}^6 = \begin{pmatrix} 1 & 2 & 8 & 5 & 8 & 2 \\ 2 & 5 & 1 & 2 & 6 & 12 \\ 12 & 9 & 15 & 8 & 8 & 13 \\ 13 & 5 & 11 & 3 & 10 & 1 \\ 1 & 15 & 13 & 14 & 11 & 8 \\ 8 & 2 & 3 & 3 & 2 & 8 \end{pmatrix} \\
(A_{196})^7 &= \begin{pmatrix} 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 1 & 4 & 6 & 1 & 1 & 6 & 4 \end{pmatrix}^7 = \begin{pmatrix} 1 & 4 & 6 & 1 & 1 & 6 & 4 \\ 4 & 2 & 15 & 2 & 5 & 10 & 5 \\ 5 & 3 & 15 & 10 & 7 & 8 & 13 \\ 13 & 4 & 11 & 2 & 7 & 15 & 9 \\ 9 & 15 & 7 & 2 & 11 & 4 & 13 \\ 13 & 8 & 7 & 10 & 15 & 3 & 5 \\ 5 & 10 & 5 & 2 & 15 & 2 & 4 \end{pmatrix} \\
(A_{256})^8 &= \begin{pmatrix} 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 2 & 4 & 2 & 11 & 2 & 8 & 5 & 6 \end{pmatrix}^8 = \begin{pmatrix} 2 & 4 & 2 & 11 & 2 & 8 & 5 & 6 \\ 12 & 9 & 8 & 13 & 7 & 7 & 5 & 2 \\ 4 & 4 & 13 & 13 & 9 & 4 & 13 & 9 \\ 1 & 6 & 5 & 1 & 12 & 13 & 15 & 14 \\ 15 & 12 & 9 & 13 & 14 & 5 & 14 & 13 \\ 9 & 14 & 5 & 15 & 4 & 12 & 9 & 6 \\ 12 & 2 & 2 & 10 & 3 & 1 & 1 & 14 \\ 15 & 1 & 13 & 10 & 5 & 10 & 2 & 3 \end{pmatrix} \\
(A_{288})^6 &= \begin{pmatrix} 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ 2 & 3 & 1 & 2 & 1 & 4 \end{pmatrix}^6 = \begin{pmatrix} 2 & 3 & 1 & 2 & 1 & 4 \\ 8 & 14 & 7 & 9 & 6 & 17 \\ 34 & 59 & 31 & 37 & 24 & 66 \\ 132 & 228 & 121 & 155 & 103 & 11 \\ 22 & 153 & 239 & 111 & 144 & 75 \\ 150 & 203 & 210 & 121 & 36 & 167 \end{pmatrix}
\end{aligned}$$

D PHOTON Constants

The round constants have been generated by a 4-bit linear feedback shift register with maximum cycle length, while the internal constants $IC_d(i)$ have been generated by shift registers with a cycle length of d . Please note that this design decision was made with hardware implementations in mind. Since one needs these counters for the control logic, the generation of the round constant and the internal constants is basically for free. For all variants we used shift registers with $l = 3$ bits, except for $d = 8$, where we used $l = 4$. Let us denote the internal state of the shift register with $X_r = (x_{l-1}, \dots, x_1, x_0)$, where $x_i = \{0, 1\}$, and $X_0 = (0, \dots, 0, 0)$. Then in each update iteration the new content of the shift register is given by $X_{r+1} = (x_{l-2}, \dots, x_0, FB(X_r))$, where $FB(X_r)$ is the feedback function. For the round constants we chose $FB(X_r) = x_3 \text{ XNOR } x_2$, while our choices for the feedback functions for the internal constants are shown in Table 7. Tables 8-11 display constants for all square sizes, round numbers and row positions.

Table 7. Feedback functions for internal constants generation.

d	5	6	7	8
$FB(X_r)$	x_2 NOR x_1	NOT x_2	x_2 XNOR x_0	NOT x_3
$IC_d(\cdot)$	[0, 1, 3, 6, 4]	[0, 1, 3, 7, 6, 4]	[0, 1, 2, 5, 3, 6, 4]	[0, 1, 3, 7, 15, 14, 12, 8]

Table 8. Constants for $d = 5$.

round row	1	2	3	4	5	6	7	8	9	10	11	12
0	1	3	7	14	13	11	6	12	9	2	5	10
1	0	2	6	15	12	10	7	13	8	3	4	11
2	2	0	4	13	14	8	5	15	10	1	6	9
3	7	5	1	8	11	13	0	10	15	4	3	12
4	5	7	3	10	9	15	2	8	13	6	1	14

Table 9. Constants for $d = 6$.

round row	1	2	3	4	5	6	7	8	9	10	11	12
0	1	3	7	14	13	11	6	12	9	2	5	10
1	0	2	6	15	12	10	7	13	8	3	4	11
2	2	0	4	13	14	8	5	15	10	1	6	9
3	6	4	0	9	10	12	1	11	14	5	2	13
4	7	5	1	8	11	13	0	10	15	4	3	12
5	5	7	3	10	9	15	2	8	13	6	1	14

Table 10. Constants for $d = 7$.

round row	1	2	3	4	5	6	7	8	9	10	11	12
0	1	3	7	14	13	11	6	12	9	2	5	10
1	0	2	6	15	12	10	7	13	8	3	4	11
2	3	1	5	12	15	9	4	14	11	0	7	8
3	4	6	2	11	8	14	3	9	12	7	0	15
4	2	0	4	13	14	8	5	15	10	1	6	9
5	7	5	1	8	11	13	0	10	15	4	3	12
6	5	7	3	10	9	15	2	8	13	6	1	14

Table 11. Constants for $d = 8$.

round row	1	2	3	4	5	6	7	8	9	10	11	12
0	1	3	7	14	13	11	6	12	9	2	5	10
1	0	2	6	15	12	10	7	13	8	3	4	11
2	2	0	4	13	14	8	5	15	10	1	6	9
3	6	4	0	9	10	12	1	11	14	5	2	13
4	14	12	8	1	2	4	9	3	6	13	10	5
5	15	13	9	0	3	5	8	2	7	12	11	4
6	13	15	11	2	1	7	10	0	5	14	9	6
7	9	11	15	6	5	3	14	4	1	10	13	2

E Test Vectors

Below are test vectors for all discussed flavours of PHOTON. The absorbing and squeezing position of the state array is underlined.

	IV	m	$P(m)$
PHOTON-80/20/16	<u>0 0 0 0 0</u> 0 0 0 0 0 0 0 0 0 0 0 0 0 0 1 4 1 4 1 0	0 0 0 0 0	<u>3 3 D 5 F</u> 6 2 9 B 9 5 C 4 8 1 6 5 C E 7 B 7 7 0 C
PHOTON-128/16/16	<u>0 0 0 0 0 0</u> 0 2 0 1 0 1 0	0 0 0 0	<u>9 5 F C 3 C</u> E 2 2 A 2 A 6 3 2 D 6 F E B 4 E 0 B 6 2 5 9 2 D 8 D 0 3 2 9
PHOTON-160/36/36	<u>0 0 0 0 0 0 0</u> <u>0 0</u> 0 2 8 2 4 2 4	0 0 0 0 0 0 0 0 0	<u>1 F 0 D 4 A 1</u> <u>D D</u> 0 A 3 1 D E C F 5 B 6 9 B 6 6 E 0 C 8 F 6 4 4 C E E E 9 0 2 0 F 4 3 A 9 D E 7 4
PHOTON-224/32/32	<u>0 0 0 0 0 0 0 0</u> 0 3 8 2 0 2 0	0 0 0 0 0 0 0 0	<u>1 7 3 0 4 2 4 2</u> 9 C F 2 6 E 1 0 8 D 3 D 9 C F 9 0 0 E 2 7 B D C C 6 2 9 B 3 D 1 A F 4 1 F 1 C B 7 4 8 3 F C C 0 8 9 1 6 B 8 2 C
PHOTON-256/32/32	<u>00 00 00 00 00 00</u> 00 40 20 20	00 00 00 00	<u>4D BD 90 36 1C B5</u> E0 9E 5C 38 A9 C9 E9 D5 66 08 CF 52 CB 6B C8 8B 93 16 E8 C2 C0 69 25 F7 18 CC 62 9C AE 79