Suit up!—Made-to-Measure Hardware Implementations of ASCON

Hannes Groß, Erich Wenger, Christoph Dobraunig and Christoph Ehrenhöfer

Graz University of Technology

Institute for Applied Information Processing and Communications (IAIK)

Inffeldgasse 16a, 8010 Graz, Austria

Abstract—**Having ciphers that provide confidentiality and authenticity, that are fast in software and efficient in hardware, these are the goals of the CAESAR authenticated encryption competition. In this paper, the promising CAESAR candidate ASCON is implemented in hardware and optimized for different typical applications to fully explore ASCON's design space. Thus, we are able to present hardware implementations of Ascon suitable for RFID tags, Wireless Sensor Nodes, Embedded Systems, and applications that need maximum performance. For instance, we show that an ASCON implementation with a single unrolled round transformation is only 7 kGE large, but can process up to 5.5 Gbit/sec of data (0.75 cycles/byte), which is already enough to encrypt a Gigabit Ethernet connection. Besides, ASCON is not only fast and small, it can also be easily protected against DPA attacks. A threshold implementation of ASCON just requires about 8 kGE of chip area, which is only 3.1 times larger than the unprotected low-area optimized implementation.**

I. Introduction

Symmetric cryptography has a rich history of competitions to find good and secure cryptographic primitives. Winners of such competitions like Rijndael—now known as the Advanced Encryption Standard (AES)—serve as base of our modern information and communication systems. Currently, in a competition called CAESAR [1], over 45 different authenticated encryption (AE) schemes battle for entry in the final portfolio to be tomorrows information security working horses. The final portfolio will be announced at the end of 2017 and till then, the participating primitives have to survive an annual selection process.

Cryptanalysis is not enough to settle the battle for the entry in the next rounds. It is likely—as in competitions before—that we get to some point during CAESAR, where every remaining authenticated encryption scheme, has no shown cryptographic weaknesses and equivalent cryptographic properties as the other candidates. Thus, cryptanalysis results of the pure algorithms should be supported by implementation results at this point regarding the selection process. Therefore, it is useful to compare other properties of the different algorithms like speed and code size of software implementations, area and throughput of hardware implementations, or the costs of protection mechanism against implementation attacks like side-channel analysis.

Providing hardware implementations of CAESAR candidates is an important task to base the decisions during CAESAR on quantifiable numbers and not on intuition. Thus, preliminary hardware implementations have been included in many submissions to the CAESAR competition like for NORX [2] and ICEPOLE [3]. However, these preliminary implementations are often only optimized for one specific design criteria (e.g. speed), and do not include any mechanisms to counteract side channel attacks. Moreover, not all CAESAR candidates already provide hardware implementation results. Therefore, the design space evaluation and comparison of the CAESAR candidates is an ongoing process [1].

In this work, we provide the first hardware implementations for the CAESAR candidate ASCON. We maximize the impact of our contribution by not focusing on only one criteria. Instead, we investigate the requirements of typical authenticated encryption applications and check the suitability of ASCON for all of them. In particular, we present implementations of ASCON suitable for RFID tags, Wireless Sensor Nodes, Embedded Systems, and applications, which need a maximum of performance. In addition, we provide protected implementations of ASCON for applications, where side-channel analysis is a threat. The resulting ASCON implementations are compared with all the hardware implementations of other CAESAR competitors available to us. This comparison shows that ASCON provides the flexibility to meet the requirements of many AE applications and is even able to compete with ciphers designed for one specific use case. To the best of our knowledge, this is the first work where such a vast amount of design spaces is evaluated and compared to other competitors regarding CAESAR.

This paper is structured as follows: Section II derives the requirements of today's most demanding applications for authenticated encryption. The ASCON algorithm itself is then briefly presented in Section III. In Section IV, the different variants of ASCON hardware designs are discussed and their underlying design principles are explained in more detail. Furthermore, it is shown how the ASCON implementations can be easily protected against first-order differential power analysis (DPA) attacks. Finally, the results of the hardware implementations are discussed in Section V, and compared to other CAESAR candidates and other authenticated encryption schemes.

II. Target Applications of Authenticated Encryption Schemes

The authenticated encryption candidates from the CAESAR competition will be implemented in software and in hardware. Although general-purpose microprocessors are constantly improving, certain demanding applications necessitate dedicated hardware designs. Software implementation would simply not perform sufficiently.

In the following, a summary of typical applications is given to elaborate their respective requirements on authenticated encryption implementations.

A common concern of symmetric cryptography is **performance**. Especially servers and encrypted/authenticated high-speed network links potentially require hardware accelerators to cope with the protected network traffic. Prominent examples for hardware accelerators are Intel's AES [4] and upcoming SHA [5] instructions. The most important characteristic of those hardware accelerators is *throughput*. However, other applications have different requirements.

One target application of authenticated encryption are passively powered **Radio Frequency Identification** (RFID) tags. These tags are wirelessly supplied by the reader field and have extremely high cost constraints in order to fulfill the requirements for mass production. The vital characteristic of these tags is the maximum read range. The limiting factor for the read range of passive RFID tags is their power consumption. Therefore, CAESAR candidates suitable for RFID should consume *very low power* and should be optimized for *low area*.

Similarly to RFID tags, **Wireless Sensor Nodes** constitute an important topic for the community. When wireless sensor nodes are deployed in the field, a single battery should ideally last for a life-cycle. Therefore, it is crucial to have CAESAR candidates, which can be optimized for *low energy* consumption. Current wireless sensor nodes come with low-energy microprocessors. Therefore, a low-energy cipher design must come with a bus interface to be usable for the microprocessor.

While servers, RFID tags, and wireless sensor nodes are very specific applications, most currently deployed systems nowadays are **Embedded Systems** that come with a microprocessor and a lot of peripherals. These microprocessors are the default solution for many consumer and industrial applications. Manufacturers add small but performant memory-mapped cryptographic coprocessors in order to fulfill versatile security tasks. On the one hand these coprocessors should be as small as possible, but on the other hand they should also be very fast which is normally a contradiction. Therefore, they can be best characterized in terms of *throughput per area*—in the following also referred to as *efficiency*.

TABLE I.   OPTIMIZATION GOALS FOR DIFFERENT AUTHENTICATED ENCRYPTION APPLICATIONS

Application	Optimization Goal	Interface
High performance computing	high throughput	custom
RFID tags	low power and low area	custom
Wireless sensor nodes	low energy	memory-mapped
Embedded systems	high throughput per area	memory-mapped

Table I summarizes the optimization goals for different applications measured in traditional hardware metrics. An additional design dimension is **implementation security**. In fact, there is a whole class of practical side-channel attacks which modern devices need to be prepared for. Protecting implementations against physical attacks is not trivial. Since Kocher et al.'s [6] paper about differential power analysis, the research community struggles with the secure implementation of AES. To avoid such challenges, it is an important requirement for all CAESAR candidates to be easily protectable against implementation attacks.

In the following, it is shown that ASCON is extremely flexible and can be optimized according to the needs of all of the above mentioned applications. Furthermore, it is demonstrated that ASCON can be efficiently protected against power analysis attacks.


Fig. 1.   The encryption of ASCON-128 (taken from [7])

III.   ASCON

ASCON [7] has a sponge-like mode of operation as depicted in Figure 1. Its state size, the permutation p and mode of operation are chosen in a way that allows compact hardware implementations, while still providing high throughput. ASCON comes in two different versions, namely ASCON-128 and ASCON-96 with 128 and 96 bit security level, respectively. This paper focuses on ASCON-128, while the modifications for the faster ASCON-96—essentially a doubling of the rate—are straight-forward.

A. Mode of Operation

ASCON has a state size of 320 bits (consisting of five 64-bit words $x_0, \ldots, x_4$) that are updated in four phases: *Initialization*, *Processing of Associated Data*, *Processing of Plaintext/Ciphertext*, and *Finalization*. All phases use the same permutation function p that is applied twelve times in the *Initialization* and *Finalization* phase. The lighter variant of p with six rounds is used for processing the data and ensures high performance. The data is handled in 64-bit blocks.

The *Initialization* phase takes the secret key K (128 bits) and the public nonce N (128 bits). This nonce has to be fresh for every encryption and must not be used twice. If the nonce is used twice or multiple times, then the confidentiality is jeopardized.

After the *Initialization* phase the optional associated data A_i is processed. Associated data is information, which does not need to be confidential, but must not be altered by an attacker. Each block A_i is added to the secret state. If there is no associated data to process, the whole step can be omitted.

In the *Encryption* phase, each plaintext block P_i is xored with the secret state to produce one ciphertext block C_i. Six consecutive round transformations p are executed for each of the 64-bit data blocks.

After the generation of the ciphertext, the *Finalization* starts. The output of the *Finalization* is the 128-bit tag T. With the help of this tag, modifications of the ciphertext and the associated data can be detected during decryption (validation). Decryption is very similar to encryption. Just the part, where the ciphertext is processed instead of the plaintext differs slightly. Thus, no inverse of the permutation is needed for decryption. So, both encryption and decryption can be implemented with just a slight overhead compared to encryption only.

B. Permutation

ASCON-128 uses two permutations p^6 and p^{12}. The two permutations are the 6 and 12 iterative execution of the round transformation p. The round transformation p consists of a constant addition to x_2, followed by an application of a substitution layer, and a linear layer.

The substitution layer is the parallel application of 64 5-bit S-boxes. The S-boxes used for ASCON are an affine transformation of the χ mapping of Keccak [8]. This affine mapping improves some cryptographic properties of the Keccak's χ mapping, while still leaving the core of the S-box and therefore the algebraic degree of 2 intact. Moreover, the ASCON S-box can be implemented, using only a few logical operations, which are highly parallelizeable (Figure 2).

$$x_4 \oplus (x_4 \ggg 7) \oplus (x_4 \ggg 41) \to x_4$$
$$x_3 \oplus (x_3 \ggg 10) \oplus (x_3 \ggg 17) \to x_3$$
$$x_2 \oplus (x_2 \ggg 1) \oplus (x_2 \ggg 6) \to x_2$$
$$x_1 \oplus (x_1 \ggg 61) \oplus (x_1 \ggg 39) \to x_1$$
$$x_0 \oplus (x_0 \ggg 19) \oplus (x_0 \ggg 28) \to x_0$$

Fig. 2. Substitution layer with 5-bit S-box (left) and linear layer (right)

The linear layer consists of five applications of the function $\Sigma_{l_i,r_i}(x_i) = x_i \oplus (x_i \ggg l_i) \oplus (x_i \ggg r_i)$ to each 64-bit word of the state $(x_0, \ldots, x_4)$. The Σ function is similar to the one used in SHA-2 [9], except that other rotation values (l_i, r_i) are used. The rotation values (l_i, r_i) are different for every 64-bit word in one round and are: $(19, 28)$ for x_0, $(61, 39)$ for x_1, $(1, 6)$ for x_2, $(10, 17)$ for x_3, and $(7, 41)$ for x_4.

C. Hardware Security Properties of ASCON

In order to protect ASCON against side-channel attacks it is important to reflect on the properties of ASCON that make the life of an attacker hard. ASCON uses a mode of operation which is based on MonkeyDuplex [10]. In contrast to MonkeyDuplex, ASCON uses a keyed *Initialization* and *Finalization*. This has the effect that a state recovery during the processing of data neither leads to the recovery of the secret key, nor allows universal forgeries.

Therefore, Power-analysis attacks on the data processing phase may be applied in order to recover the internal states, but do not allow the attacker to recover the key. In addition, power-analysis attacks on the *Finalization* are hard, since the attacker would have to attack both the key and many unknown state bits that have no influence on the emitted tag. In fact, the most vulnerable phase is the *Initialization*. For ASCON-128, three out of five input bits of the first S-boxes are publicly known and the other two bits belong to the secret key.

To prevent side-channel analysis, ASCON allows to apply well known protection mechanisms. The ASCON S-box is an affine transformation of the Keccak S-box. Hence, it is possible to reuse essential parts of techniques for side channel protection like the threshold implementation of Keccak from Bilgin et al. [11] as discussed in Section IV-D. Therefore, ASCON can be protected against first-order DPA attacks by using only three shares and four random bits per round.

IV. HARDWARE DESIGNS

ASCON allows to be optimized for many practical applications, where both confidentiality and authenticity are required. These include all applications mentioned in Section II with their different optimization goals. In the following, three implementation variants with different design goals are introduced.

A. High Throughput Design (ASCON-fast)

Due to the low complexity of ASCON's round transformation, it is possible to fully unroll a complete round transformation and still achieve high frequencies. As it turns out, the round transformations are so hardware-friendly that even multiple rounds can be computed in a single clock cycle. The ASCON-fast variants aim at a maximal data throughput with a minimum of processing delay. Therefore, at least one round transformation is performed in every clock cycle and no pipelining stages are used. Each ASCON-fast variant uses a different number of unrolled round transformations. The datapath of ASCON-fast as it is shown in Figure 3 mainly consists of the unrolled round transformations (five 64-bit state registers, 64 parallel S-boxes, and the linear diffusion layer). Only a few additional multiplexers and XOR-gates are needed to connect the unrolled round transformation with the data bus and the key register.
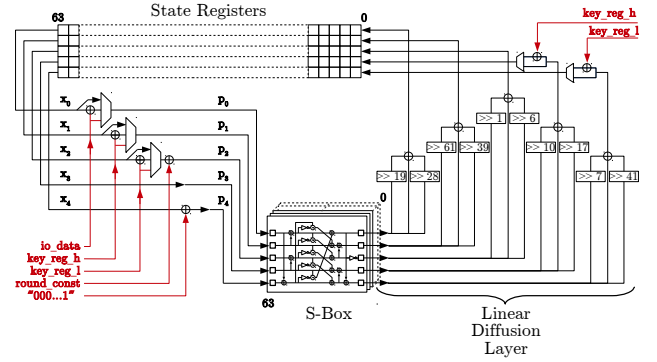

Fig. 3. Datapath of ASCON-fast (one round transformation per cycle)

B. 64-bit Datapath Design (ASCON-64-bit)

The design idea behind the ASCON-64-bit implementation is based on the inherent 64-bit structure of ASCON. Instead of a concrete implementation of the S-box and the linear diffusion layer, this design uses an arithmetic logic unit (ALU) comparable to a microcontroller design. Consequently, the controlpath works similar to a sequential program code that is executed by the datapath in Figure 4.

Besides the five state registers, there exist also two temporary registers, which—together with the inputs from the controlpath—form the input operands of the ALU. The ALU itself consists of an iterative barrel-shifter unit, three logic operations, and a data-storage unit that takes the 64-bit bus data input and stores it either in the high or the low part of the selected operand. On the output of the ALU the result of the operation is selected that is then applied to the destination register. During the execution, the S-box and linear layer are iteratively calculated using the operations of the ALU. Thus, one round operation takes 59 clock cycles.
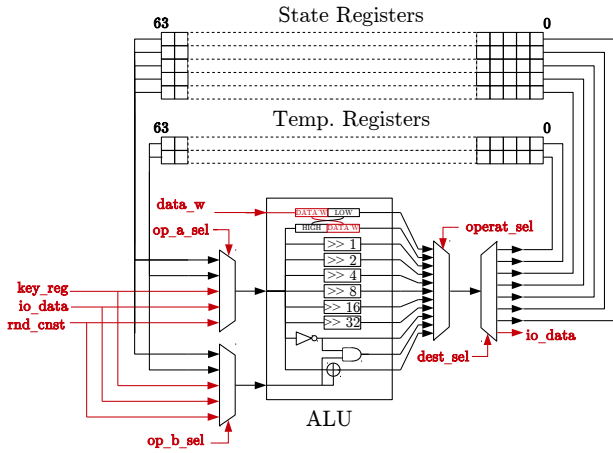
Fig. 4. Datapath of the *64-bit* variant of ASCON

C. Low Area Design (ASCON-x-low-area)

The datapath of the so-called ASCON-x-low-area variant (see Figure 5) uses a radical low-area approach, which can be summarized as "one bit operation per cycle". The state consists of five clock gated shift registers with independent shift-enable inputs. For the S-box calculation, all state registers are activated and shifted bit-slice-wise through the single S-box instance. The result is stored in the least-significant bits of the state. Accordingly, the whole S-box layer operation consumes 64 cycles. The subsequent linear diffusion layer is split up into five interleaved subiterations in which each state register is updated individually. As a single state bit depends on two other bits of the same state row, the linear layer cannot be calculated without temporarily storing either the results or the state row itself, respectively. Thus, another (temporary) shift register is needed that in one iteration holds the result of the current linear layer operation and in the next iteration is used to write the result back. Once the first subiteration of the linear layer is finished, the calculation of the next state row and the write back operation can be done in parallel. This uncompromising low-area approach results in 512 clock cycles per round transformation.
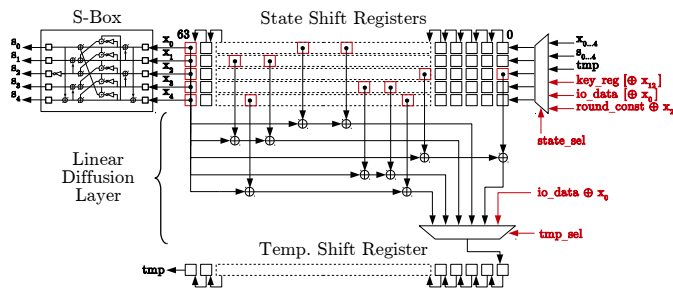

Fig. 5. Datapath of the *x-low-area* variant of ASCON

D. Protecting ASCON against Side Channel Analysis Attacks

In practice it is not sufficient to have fast, small, low-power, or low-energy designs. In addition, those designs need to be protected against the serious threat of implementation attacks like power-analysis attacks. In 2006, Nikova et al. [12] introduced the concept of a threshold implementation (TI) scheme that is provably resistant against first-order DPA attacks. The basic principle behind TI is that the calculations are never performed directly on the security critical data but indirectly by applying modified transformations on so-called shares (named A, B and C in Figure 6).
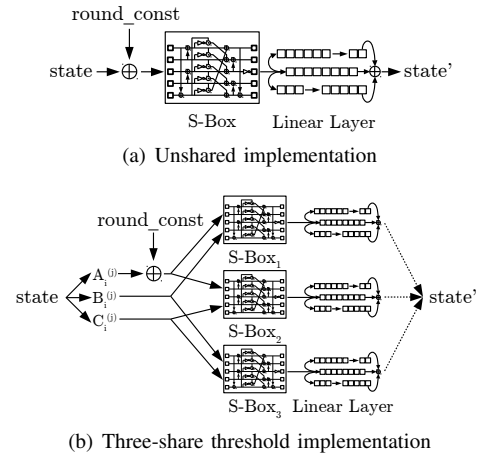

(a) Unshared implementation

(b) Three-share threshold implementation

Fig. 6. The principle of secret sharing explained on one round of ASCON

Linear transformations, like the addition of the round constant or ASCON's linear layer, can be performed on each share separately. Non-linear operations, like S-box lookups, have to be carefully transformed to remain valid. To withstand first-order DPA attacks, the S-box component functions have to fulfill the correctness, non-completeness, and uniformity properties. A sharing of an S-box is said to be correct if the sum of the resulting output shares equals the result of the S-box function applied to the sum of the input shares (see Figure 6.b). Also the non-completeness property is represented in Figure 6.b by having three S-box component functions which are each independent of at least one input share. The last property which is required for a valid TI is the so-called uniformity which results for an S-box implementation in the equivalent requirement for each component function to be invertible.

One of the best explored algorithm regarding threshold implementations is the new SHA-3 standard Keccak [8]. In 2014, Bilgin et al. [11] not only showed how to share the Keccak's S-box function χ very efficiently with three shares, but also how to reduce the high amount of random bits required to recover the uniformity property. ASCON uses an affine transformation of Keccak's χ function. Therefore, the current findings of Bilgin et al. and all future findings on the Keccak S-box can be directly applied to ASCON. Since, the sharing of the χ function is already well explained by Bilgin et al., the details of ASCON's shared S-box are summarized in Appendix A for the interested reader.

In order to investigate the additional costs for protecting ASCON against first-order DPA attacks, three-share TI versions of the *fast* and the *x-low-area* variants are implemented. ASCON-fast-TI is assumed to be used in an embedded system scenario as a cryptographic co-processor that is controlled by a microcontroller. The microcontroller is responsible for performing the initial sharing of the state and for distributing its source of randomness. The necessary random bits (four fresh random bits per round transformation) are delivered over

the 64-bit AMBA APB bus interface in advance of each round transformation and stored in an internal register.

Ascon-x-low-TI, on the other hand, is designed to be used within a microcontroller-free design, e.g., a commercial RFID tag. Therefore, a direct access to a random number generator is assumed to be available, like it is used in EPC Gen2 UHF tags. Because the Ascon-x-low-TI implementation only performs one S-box operation per cycle, the output of the random number generator is directly used and does not need to be stored in advance.

In the next section we evaluate how the implemented Ascon designs perform compared to each other and compared to related work.

V. Results

All Ascon designs are implemented in VHDL and evaluated using a Cadence-based ASIC design-flow. The VHDL code is open-source and is available online.[1] This will facilitate future comparisons by fellow researchers using different ASIC or FPGA design flows. For the following results, a 90 nm UMC standard performance low-K library from Faraday is used with a global clock of 1 MHz and a 1 V power supply. The designs are compiled with the Cadence Encounter RTL compiler version v08.10-s28_1 and routed with Cadence NanoRoute v08.10-s155.

TABLE II. CHARACTERISTICS OF THE ASCON-128 HARDWARE IMPLEMENTATIONS

Design	Chip Area w/o interface	Area w/ interface	Throughput		Power at 1 MHz	Energy
	[kGE]	[kGE]	[cycles/byte]	[Mbps]	[μW]	[μJ/byte]
Unprotected Implementations						
Ascon-fast						
1 round	7.08	7.95	0.75	5,524	43	33
2 rounds	10.61	11.48	0.38	8,425	72	27
3 rounds	14.26	15.13	0.25	10,407	102	25
6 rounds	24.93	25.80	**0.13**	**13,218**	184	**23**
Ascon-64-bit	4.99	5.86	44.25	72	32	1,397
Ascon-x-low-area	**2.57**	**3.75**	384.00	14	**15**	5,706
Threshold Implementations						
Ascon-fast-TI						
1 round	28.61	30.42	0.75	3,774	183	137
2 rounds	47.46	49.13	0.38	6,289	315	118
3 rounds	66.45	68.27	0.25	7,143	447	112
6 rounds	123.52	125.19	**0.13**	**9,018**	830	**104**
Ascon-x-low-TI	**7.97**	**9.19**	384.00	15	**45**	17,234

The results of all practical evaluations are collected in Table II. Some implementations can process up to 8 bytes of data in a single clock cycle (0.125 cycles per byte) and other implementations are as small as 2.57 kGE. Especially the characteristics of Ascon-fast with one unrolled round are impressive. This implementation needs 7.08 kGE (7.95 kGE with key register and 64-bit bus interface), reaches a maximum clock frequency of 517 MHz, and can therefore process up to 5.5 Gbit per second. This means that the most straightforward design is easily sufficient to encrypt a gigabit Ethernet connection on the fly. At 100 MHz, the design only needs 529 μW (38 μW static leakage and 4.9 μW/MHz dynamic power) and is therefore also suitable for mobile applications. As it also provides the best performance per throughput, it is perfectly suitable for *embedded systems*.

If higher throughput is required, Ascon-fast with six unrolled rounds can process more than 13 Gbit/sec at 206 MHz.

TABLE III. CHARACTERISTICS OF RELATED IMPLEMENTATIONS

Design	Chip Area	Throughput	Power	Technology
	[kGE]	[Mbps]	[μW/MHz]	
AES-CCM [13]	3.77	57	5.12	STM 65 nm
AES-OCB2 [13]	5.92	113	8.11	STM 65 nm
AES-ALE [13]	2.70	244	10.55	STM 65 nm
Minalpher [14] low-area	2.81	369	—	
SILC [15] V1	15.70	764	—	
AES-OCB [16]	22.55	854	—	TSMC 90 nm
SILC [15] V2	23.10	2,635	—	
Scream ED [17] 1 Round	6.23	4,577	—	STM 65 nm
Keccak MonkeyDupl. [18]	5.90	4,900	42	
Scream ED [17] 2 Round	8.31	5,190	—	STM 65 nm
Minalpher [14] high-speed	14.32	6,104	—	
Norx [2]	59.00	10,000	—	UMC 180 nm
ICEPOLE [3]	—	41,364	—	Xilinx Virtex 6

This is more than sufficient even for 10 gigabit network connections. For *RFID* applications, where size as well as power matter, Ascon-x-low is only 2.57 kGE large and requires as little as 15 μW for a 1 MHz clock source. In an for RFID tags more suitable, power saving 130 nm low-leakage UMC technology, the power consumption is reduced to 4.1 μW.

Energy, the most critical characteristic of *wireless sensor nodes*, is the product of power and runtime. The Ascon-fast implementations require the least amount of energy because of their low runtimes. Even though six unrolled rounds give the best energy results it probably would be more reasonable to use Ascon-fast with one unrolled round for *wireless sensor nodes*.

A. Related Work

A fair comparison of hardware designs is a difficult task. Different designers make diverging assumptions about, e.g., key registers or bus interfaces. Additionally, results are highly dependent on the used manufacturing technology, the used toolchain (e.g., Cadence, Synopsis, Menthor, Xilinx, or Altera), and the external operating conditions (e.g., power supply voltage or ambient temperature). Therefore, the following comparison with related work has to be interpreted carefully.

In Table III, the hardware results of several authenticated encryption designs are listed. There are standardized AES-based implementations [13], [16] and CAESAR candidates based on sponge constructions [2], [3], [18] and block ciphers [17], [15], [14] depicted. Figure 7 visually combines Tables II and III. The horizontal axis shows the throughput and the vertical axis depicts the area footprint. The dashed equi-efficiency lines indicate a constant throughput per area ratio.

It seems that Ascon provides, together with Keccak MonkeyDuplex [18], excellent performance per area. All AES implementations are a magnitude slower than even the slowest Ascon-fast implementation. Only Norx [2] and ICEPOLE [3] achieve similar or higher performance. However, Norx is more than twice as large (59 kGE) as the largest Ascon design (26 kGE). There are no ASIC results available for ICEPOLE so far. In terms of size, Ascon needs six times fewer registers to store the state and the processed data (1280 + 1024 vs. 320 + 64) than ICEPOLE. ICEPOLE S-boxes also have an algebraic degree of four and are thus potentially harder to share

than the Ascon S-boxes which probably results in a higher overhead for protected implementations.

B. First-Order DPA Resistant Threshold Implementations

Compared to the unprotected versions, the state size intuitively triples for the TI variants Ascon-fast-TI and Ascon-x-low-TI because of the three state shares. Furthermore, also the logic inside the datapath is implemented in triplicate, but the S-box sharing results in a increased resource consumption because of the more complex component functions. This increase in the resource requirements affects the Ascon-fast-TI more than the Ascon-x-low-TI because the latter only uses one S-box instance. As a result, Ascon-fast-TI is 4.0 times larger than Ascon-fast. Nevertheless, it still offers a throughput of 3.77 Gbit/sec. Ascon-x-low-TI, on the other hand, has a comparably low overhead factor of only 3.1.

The authors of this paper are not aware of the existence of any other side-channel protected CAESAR candidates so far to compare Ascon to. Due to its prevalence, a lot of effort of the hardware security community was spent on the question on how to protect AES against DPA attacks. In 2011, Moradi et al. [19] (later improved by Bilgin et al.[20]) implemented a low-area version of AES (encryption only), which is protected by the TI masking scheme. As it turned out, the overhead factor for protecting their AES core against first-order DPA attacks is about 4.6 (2.4 kGE vs. 11 kGE). The major reason for this implementation overhead is the S-box, which is very difficult to protect compared to the S-box of Ascon.

VI. Conclusions

In this work we exhaustively evaluated that Ascon is ready for all challenges that today's authenticated encryption applications pose. We demonstrated how to implement high-throughput designs that can process up to 13.2 gigabit of data per second. A design designated for RFID applications requires as little as 2.6 kGE and 15 μW of power at 1 MHz. Ascon-fast also fulfills all requirements of wireless sensor nodes in particular and embedded systems in general. It requires only 7.1 kGE and 33 μJ of energy per byte and can easily encrypt and authenticate the traffic of an one-gigabit Ethernet connection. As additional contribution, we demonstrated the ease of protecting Ascon against power-analysis attacks. Ascon-fast-TI is a power-analysis protected threshold design that is still able to process 3.6 gigabit of data per second. The comparison with other authenticated-encryption designs showed that Ascon is ready for all challenges the CAESAR competition might hold. And there is still room for improvement: a pipelined Ascon design can process more data at higher frequencies. Such a design can easily be derived from the released Apache-licensed source codes.[2]

[2]http://ascon.iaik.tugraz.at/implementation.html

REFERENCES

[1] D. Bernstein, "CAESAR: Competition for Authenticated Encryption: Security, Applicability, and Robustness," Available online at http://competitions.cr.yp.to/caesar.html, 2014.

[2] J. Aumasson, P. Jovanovic, and S. Neves, "NORX," Submission to the CAESAR competition: http://competitions.cr.yp.to/round1/norxv1.pdf, 2014.

[3] P. Morawiecki, K. Gaj, E. Homsirikamol, K. Matusiewicz, J. Pieprzyk, M. Rogawski, M. Srebrny, and M. Wójcik, "ICEPOLE: High-Speed, Hardware-Oriented Authenticated Encryption," in *CHES 2014 - 16th International Workshop, Busan, South Korea, September 23-26, 2014. Proceedings*, ser. Lecture Notes in Computer Science, L. Batina and M. Robshaw, Eds., vol. 8731. Springer, 2014, pp. 392–413. [Online]. Available: http://dx.doi.org/10.1007/978-3-662-44709-3_22

[4] Intel Corporation, "Intel Advanced Encryption Standard Instructions (AES-NI)," March 2011, available online: http://software.intel.com/en-us/articles/intel-advanced-encryption-standard-instructions-aes-ni/.

[5] ——, "New Instructions Supporting the Secure Hash Algorithm on Intel Architecture Processors," July 2013, available online: https://software.intel.com/en-us/articles/intel-sha-extensions.

[6] P. C. Kocher, J. Jaffe, and B. Jun, "Differential Power Analysis," in *Proceedings of the 19th Annual International Cryptology Conference on Advances in Cryptology*, ser. CRYPTO '99. London, UK, UK: Springer-Verlag, 1999, pp. 388–397. [Online]. Available: http://dl.acm.org/citation.cfm?id=646764.703989

[7] C. Dobraunig, M. Eichlseder, F. Mendel, and M. Schläffer, "Ascon," Submission to the CAESAR competition: http://ascon.iaik.tugraz.at, 2014.

[8] G. Bertoni, J. Daemen, M. Peeters, and G. Van Assche, "Keccak Specifications," Submission to NIST (Round 3), 2011. [Online]. Available: http://keccak.noekeon.org

[9] National Institute of Standards and Technology, "FIPS PUB 180-4: Secure Hash Standard," Federal Information Processing Standards Publication 180-4, U.S. Department of Commerce, March 2012, available online: http://www.itl.nist.gov. [Online]. Available: http://csrc.nist.gov/publications/fips/fips180-4/fips-180-4.pdf

[10] J. Daemen, "Permutation-based Encryption, Authentication and Authenticated Encryption," DIAC – Directions in Authenticated Ciphers, July 2012.

[11] B. Bilgin, J. Daemen, V. Nikov, S. Nikova, V. Rijmen, and G. Van Assche, "Efficient and First-Order DPA Resistant Implementations of Keccak," in *Smart Card Research and Advanced Applications*, ser. Lecture Notes in Computer Science, A. Francillon and P. Rohatgi, Eds. Springer International Publishing, 2014, pp. 187–199. [Online]. Available: http://dx.doi.org/10.1007/978-3-319-08302-5_13

[12] S. Nikova, C. Rechberger, and V. Rijmen, "Threshold Implementations Against Side-channel Attacks and Glitches," in *Proceedings of the 8th International Conference on Information and Communications Security*, ser. ICICS'06. Berlin, Heidelberg: Springer-Verlag, 2006, pp. 529–545. [Online]. Available: http://dx.doi.org/10.1007/11935308_38

[13] A. Bogdanov, F. Mendel, F. Regazzoni, V. Rijmen, and E. Tischhauser, "ALE: AES-Based Lightweight Authenticated Encryption," in *Fast Software Encryption - 20th International Workshop, FSE 2013, Singapore, March 11-13, 2013. Revised Selected Papers*, ser. Lecture Notes in Computer Science, S. Moriai, Ed., vol. 8424. Springer, 2013, pp. 447–466. [Online]. Available: http://dx.doi.org/10.1007/978-3-662-43933-3_23

[14] Y. Sasaki, Y. Todo, K. Aoki, Y. Naito, T. Sugawara, Y. Murakami, M. Matsui, and S. Hirose, "Minalpher," 8 2014, dIAC 2014: Directions in Authenticated Ciphers, Santa Barbara, USA [Accessed: 2014 09 30]. [Online]. Available: http://2014.diac.cr.yp.to/slides/matsui-minalpher.pdf

[15] T. Iwata, K. Minematsu, J. Guo, S. Morioka, and E. Kobayashi, "SILC: SImple Lightweight CFB," 8 2014, dIAC 2014: Directions in Authenticated Ciphers, Santa Barbara, USA [Accessed: 2014 09 30]. [Online]. Available: http://2014.diac.cr.yp.to/slides/iwata-silc.pdf

[16] M. Parelkar, *Authenticated Encryption in Hardware*. George Mason University, 2005.

[17] V. Grosso, G. Leurent, F. Standaert, K. Varici, F. Durvaux, L. Gaspar, and S. Kerckhof, "CAESAR candidate SCREAM
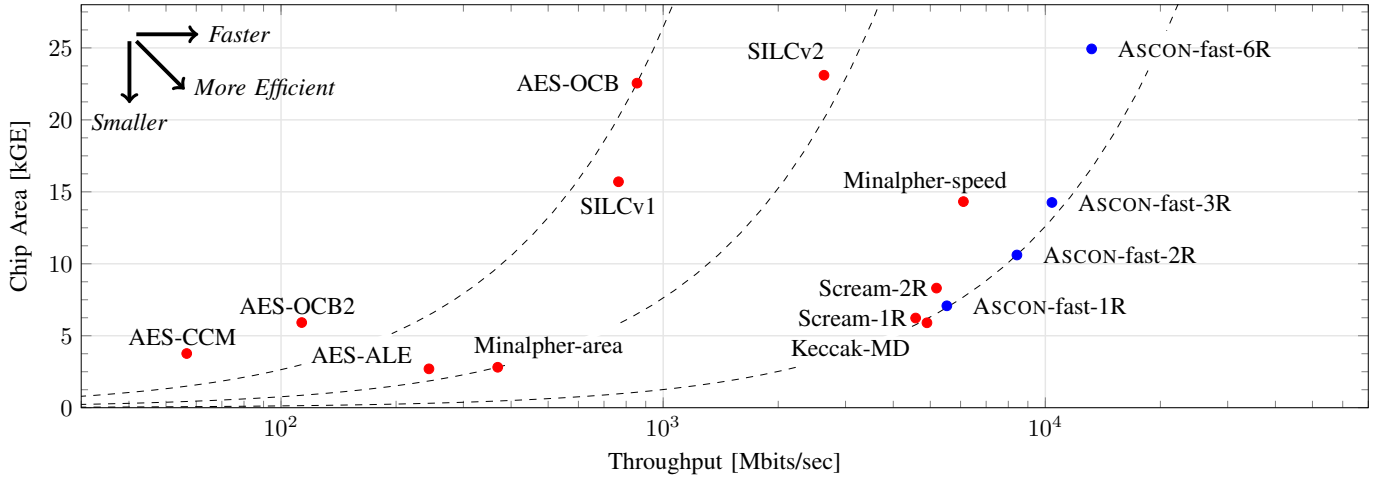
Fig. 7. A comparison between related work and ASCON considering the throughput and consumed area, and the resulting efficiency

Side-Channel Resistant Authenticated Encryption with Masking," 8 2014, dIAC 2014: Directions in Authenticated Ciphers, Santa Barbara, USA [Accessed: 2014 09 30]. [Online]. Available: http://2014.diac.cr.yp.to/slides/leurent-scream.pdf

[18] T. Yalçin and E. B. Kavun, "On the Implementation Aspects of Sponge-based Authenticated Encryption for Pervasive Devices," in *Proceedings of the 11th International Conference on Smart Card Research and Advanced Applications*, ser. CARDIS'12. Berlin, Heidelberg: Springer-Verlag, 2013, pp. 141–157. [Online]. Available: http://dx.doi.org/10.1007/978-3-642-37288-9_10

[19] A. Moradi, A. Poschmann, S. Ling, C. Paar, and H. Wang, "Pushing the Limits: A Very Compact and a Threshold Implementation of AES," in *Proceedings of the 30th Annual International Conference on Theory and Applications of Cryptographic Techniques: Advances in Cryptology*, ser. EUROCRYPT'11. Berlin, Heidelberg: Springer-Verlag, 2011, pp. 69–88. [Online]. Available: http://dl.acm.org/citation.cfm?id=2008684.2008693

[20] B. Bilgin, B. Gierlichs, S. Nikova, V. Nikov, and V. Rijmen, "A More Efficient AES Threshold Implementation," in *Progress in Cryptology - AFRICACRYPT 2014 - 7th International Conference on Cryptology in Africa, Marrakesh, Morocco, May 28-30, 2014. Proceedings*, ser. Lecture Notes in Computer Science, D. Pointcheval and D. Vergnaud, Eds., vol. 8469. Springer, 2014, pp. 267–284.

APPENDIX

The efficient implementation of a non-linear function like the S-box of ASCON is not trivial. Quite nicely, ASCON is using an affine transformation of the χ function used in Keccak [8], the new SHA-3 standard. In 2014, Bilgin et al. [11] showed a shared implementation of the Keccak S-box based on three shares with direct-sharing property. This property results in a very compact shared representation of the S-box χ with low hardware costs.

In the following, the same notation as in Bilgin et al.'s paper is used on ASCON. The state shares are represented by the capital letters A, B and C, with the bit index i (modulo 5) inside one of the 64 5-bit state slices with index j.

$$
\begin{aligned}
A'_i \leftarrow \chi'_i(B,C) &\triangleq B_i + (B_{i+1}+1)\,B_{i+2} \\
&\quad + B_{i+1}C_{i+2} + B_{i+2}C_{i+1} \\
B'_i \leftarrow \chi'_i(C,A) &\triangleq C_i + (C_{i+1}+1)\,C_{i+2} \\
&\quad + C_{i+1}A_{i+2} + C_{i+2}A_{i+1} \\
C'_i \leftarrow \chi'_i(A,B) &\triangleq A_i + (A_{i+1}+1)\,A_{i+2} \\
&\quad + A_{i+1}B_{i+2} + A_{i+2}B_{i+1}
\end{aligned}
\tag{1}
$$

The shared S-box implementation in Equation 1 fulfills the correctness and the non-completeness properties of TI implementations, but the component functions are not invertible and therefore also not uniform. However, because of the similarities to the Keccak S-box, the findings of Bilgin et al. can also be applied to the ASCON S-box. A straight-forward way to recover the uniformity is the insertion of additional random bits in form of the virtual variables S and P in Equation 2. Even though this ensures the uniformity of the component functions, this would require 640 fresh high-quality random bits per cipher round which is not quite practical!

$$
\begin{aligned}
A'_i &\leftarrow \chi'_i(B,C) + P_i + S_i \\
B'_i &\leftarrow \chi'_i(C,A) + P_i \\
C'_i &\leftarrow \chi'_i(A,B) + S_i
\end{aligned}
\tag{2}
$$

Bilgin et al. showed that when only three out of the five S-box output bits are considered, the joint uniformity is already satisfied. Consequently, only the two remaining state slice bits need to be repaired by using Equation 2. Thus, for bit slice positions $i \in \{0,1,2\}$ Equation 1 is sufficient.

The requirement for fresh random bits can be decreased further when the relation between the individual state slices is considered. Since, each sharing of one state slice is independent of the other shared slices, these independent shares can be used to recover the uniformity property of a different state slice sharing. It only needs to be ensured that the independence between the state slices is not violated. Equation 3 can then be applied to all bits of the state slices where the bit index $i \in \{3,4\}$ except the one slice with the state slice index $j = 0$

which requires fresh randomness in order to not violate the independence requirement. The total amount of fresh random bits per cipher round is then reduced to only four bits. The rest of the bits required for the virtual variables is taken from independent bit slices.

$$A'^{(j)}_i \leftarrow \chi'_i(B^{(j)}, C^{(j)}) + A^{(j-1)}_i + B^{(j-1)}_i$$
$$B'^{(j)}_i \leftarrow \chi'_i(C^{(j)}, A^{(j)}) + A^{(j-1)}_i \qquad (3)$$
$$C'^{(j)}_i \leftarrow \chi'_i(A^{(j)}, B^{(j)}) + B^{(j-1)}_i$$