Keccak implementation overview

Guido Bertoni¹ Joan Daemen¹ Michaël Peeters² Gilles Van Assche¹ Ronny Van Keer¹

http://keccak.noekeon.org/

Contents

| 1 | Gen | eral aspects | 7 |
|---|-----------------|---|----|
| | 1.1 | Specifications summary | 7 |
| | 1.2 | Bit and byte numbering conventions | 9 |
| | | 1.2.1 Some justification for our choice | 10 |
| | 1.3 | Operation count | 10 |
| | 1.4 | Memory | 12 |
| | 1.5 | Bit interleaving | 12 |
| | 1.6 | The lane complementing transform | 13 |
| 2 | Soft | ware | 15 |
| | 2.1 | PC and high-end platforms | 15 |
| | | 2.1.1 Using SIMD instructions | 16 |
| | | 2.1.2 SIMD instructions and tree hashing | 16 |
| | | 2.1.3 Batch or tree hashing on a graphics processing unit | 17 |
| | 2.2 | Small 32-bit platforms | 17 |
| | | 2.2.1 Implementation on a ARM Cortex-M3 | 17 |
| | 2.3 | Small 8-bit platforms | 18 |
| | | 2.3.1 Estimation on an Intel 8051 processor | 18 |
| | | 2.3.2 Implementation on a Atmel AVR processor | 19 |
| 3 | Har | dware | 21 |
| J | 3.1 | Introduction | 21 |
| | 3.2 | High-speed core | 21 |
| | 3.3 | Variants of the high-speed core | 23 |
| | 5.5 | 3.3.1 Keccak $[r = 1024, c = 576]$ | 23 |
| | | 3.3.2 Keccak $[r = 1024, t = 570]$ | 23 |
| | 3.4 | | 24 |
| | J. 4 | Low-area coprocessor | 26 |
| | | | 26 |
| | 3.5 | [| 26 |
| | 3.3 | FPGA implementations | 26 |
| | | 0 1 | |
| | | 3.5.2 Low-area coprocessor | 26 |
| 4 | | ection against side-channel attacks | 29 |
| | 4.1 | Introduction | 29 |
| | 4.2 | Power analysis | 30 |
| | | 4.2.1 Different types of countermeasures | 31 |
| | | 4.2.2 Secret sharing | 32 |
| | | Software implementation using two-share masking | 32 |

| 4.4 | Hardy | ware using three-share masking | 33 |
|-----|-------|-------------------------------------|----|
| | 4.4.1 | One-cycle round architecture | 34 |
| | 4.4.2 | Three-cycle round architecture | 34 |
| | 4.4.3 | Synthesis results | 38 |
| 4.5 | Comp | outing in parallel or sequentially? | 38 |

Introduction

This document gives an overview of the implementation aspects of Keccak, in software and hardware, with or without protection against side-channel attacks.

Acknowledgments

We wish to thank (in no particular order) Joachim Strömbergson for useful comments on our FPGA implementation, Joppe Bos for reporting a bug in the optimized implementation, all people who contributed to implementations or benchmarks of Keccak in hardware or software, Virgile Landry Nguegnia Wandji for his work on DPA-resistant Keccak implementations and finally Joris Delclef and Jean-Louis Modave for kindly lending us fast hardware.

Chapter 1

General aspects

In this chapter, we first briefly review the Keccak specifications formally defined in [10]. We then detail the bit and byte numbering conventions. General statements about the operation count and memory usage of Keccak follow. Finally, we introduce two implementation techniques, namely bit interleaving and the lane complementing transform.

1.1 Specifications summary

This section offers a summary of the Keccak specifications using pseudocode, sufficient to understand its structure and building blocks. In no way should this introductory text be considered as a formal and reference description of Keccak. For the formal definition of Keccak, we refer to [10].

Any instance of the Keccak sponge function family makes use of one of the seven Keccak-f permutations, denoted Keccak-f[b], where $b \in \{25, 50, 100, 200, 400, 800, 1600\}$ is the width of the permutation. These Keccak-f permutations are iterated constructions consisting of a sequence of almost identical rounds. The number of rounds $n_{\rm r}$ depends on the permutation width, and is given by $n_{\rm r}=12+2\ell$, where $2^\ell=b/25$. This gives 24 rounds for Keccak-f[1600].

```
KECCAK-f[b](A)

for i in 0 \dots n_r - 1

A = \text{Round}[b](A, \text{RC}[i])

return A
```

A Keccak-f round consists of a sequence of invertible steps each operating on the state, organized as an array of 5×5 lanes, each of length $w \in \{1, 2, 4, 8, 16, 32, 64\}$ (b = 25w). When implemented on a 64-bit processor, a lane of Keccak-f[1600] can be represented as a 64-bit CPU word.

```
Round[b](A, RC)
  \theta step
     C[x] = A[x,0] \oplus A[x,1] \oplus A[x,2] \oplus A[x,3] \oplus A[x,4],
                                                                                         \forall x \text{ in } 0 \dots 4
     D[x] = C[x-1] \oplus ROT(C[x+1], 1),
                                                                                         \forall x \text{ in } 0 \dots 4
     A[x,y] = A[x,y] \oplus D[x],
                                                                                         \forall (x, y) \text{ in } (0...4, 0...4)

ho and \pi steps
     B[y, 2x + 3y] = ROT(A[x, y], r[x, y]),
                                                                                         \forall (x, y) \text{ in } (0 ... 4, 0 ... 4)
  \chi step
     A[x,y] = B[x,y] \oplus ((\text{NOT } B[x+1,y]) \text{ AND } B[x+2,y]),
                                                                                   \forall (x, y) \text{ in } (0...4, 0...4)
     A[0,0] = A[0,0] \oplus RC
  return A
```

Here the following conventions are in use. All the operations on the indices are done modulo 5. A denotes the complete permutation state array and A[x,y] denotes a particular lane in that state. B[x,y], C[x] and D[x] are intermediate variables. The symbol \oplus denotes the bitwise exclusive OR, NOT the bitwise complement and AND the bitwise AND operation. Finally, ROT(W,r) denotes the bitwise cyclic shift operation, moving bit at position i into position i + r (modulo the lane size).

The constants r[x, y] are the cyclic shift offsets and are specified in the following table.

| | x = 3 | x = 4 | x = 0 | x = 1 | x = 2 |
|-------|-------|-------|-------|-------|-------|
| y=2 | 25 | 39 | 3 | 10 | 43 |
| y = 1 | 55 | 20 | 36 | 44 | 6 |
| y = 0 | 28 | 27 | 0 | 1 | 62 |
| y = 4 | 56 | 14 | 18 | 2 | 61 |
| y = 3 | 21 | 8 | 41 | 45 | 15 |

The constants RC[i] are the round constants. The following table specifies their values in hexadecimal notation for lane size 64. For smaller sizes they must be truncated.

| RC[0] | 0x00000000000000001 | RC[12] | 0x000000008000808B |
|--------|---------------------|--------|---------------------|
| RC[1] | 0x0000000000008082 | RC[13] | 0x800000000000008B |
| RC[2] | 0x800000000000808A | RC[14] | 0x8000000000008089 |
| RC[3] | 0×8000000080008000 | RC[15] | 0x8000000000008003 |
| RC[4] | 0x000000000000808B | RC[16] | 0x8000000000008002 |
| RC[5] | 0x0000000080000001 | RC[17] | 0×80000000000000080 |
| RC[6] | 0x8000000080008081 | RC[18] | 0×000000000000800A |
| RC[7] | 0x8000000000008009 | RC[19] | 0×800000008000000A |
| RC[8] | 0x000000000000008A | RC[20] | 0x8000000080008081 |
| RC[9] | 0x0000000000000088 | RC[21] | 0×8000000000008080 |
| RC[10] | 0x0000000080008009 | RC[22] | 0x0000000080000001 |
| RC[11] | 0×000000008000000A | RC[23] | 0x8000000080008008 |
| | | | |

We obtain the Keccak[r,c] sponge function, with parameters capacity c and bitrate r, if we apply the sponge construction to Keccak-f[r+c] and perform specific padding on the message input. The following pseudocode is restricted to the case of messages that span a whole number of bytes and where the bitrate r is a multiple of the lane size.

```
Keccak[r,c](M)
  PADDING
  P = M||0x01||0x00||...||0x00
  P = P \oplus 0x00||...||0x00||0x80
  Initialization
                                                               \forall (x,y) \text{ in } (0...4,0...4)
  S[x,y]=0,
  ABSORBING PHASE
  for every block P_i in P
    S[x,y] = S[x,y] \oplus P_i[x+5y],
                                                               \forall (x,y) such that x + 5y < r/w
    S = \text{Keccak-} f[r+c](S)
  Squeezing Phase
  Z = \text{empty string}
  while output is requested
    Z = Z||S[x,y],
                                                               \forall (x,y) such that x + 5y < r/w
    S = \text{Keccak-} f[r+c](S)
  return Z
```

Here S denotes the state as an array of lanes. The padded message P is organised as an array of blocks P_i , themselves organized as arrays of lanes. The || operator denotes byte string concatenation.

1.2 Bit and byte numbering conventions

Internally, the state of Keccak-f[b] is organized in three dimensions and its bits are identified with coordinates $x, y \in \mathbb{Z}_5$ and $z \in \mathbb{Z}_w$, for b = 25w and $w \in \{1, 2, 4, ..., 64\}$. Externally, the sponge and duplex constructions require to have a linear numbering of the bits from 0 to b-1. Here, the bit index i = z + w(5y + x) externally corresponds to the coordinates (x, y, z) internally.

A lane (i.e., bits with the same coordinates (x, y)) contains bits with w consecutive indices. If an input or output block is organized as lanes, the outer part of the state spans $\lceil \frac{r}{w} \rceil$ lanes. In a typical software implementation, the bits in a lane are packed together in a w-bit CPU word. Thus, this allows to organize the input and output in terms of CPU words. In addition, one can implement the operation ρ as a set of CPU word rotations (if the processor supports it).

When it comes to expressing the bit positions within a CPU word or a byte, or the byte positions within a CPU word, we take the following conventions.

1. The bits within a byte or within a CPU word are numbered from zero onwards, with the bit number i being the coefficient of 2^i when a byte or a word needs to be represented as an integer.

2. The bytes within a CPU word are numbered from zero onwards, with the byte at offset i being the coefficient of 256^i when a words needs to be represented as an integer ("little-endian" convention).

The consequence of the first convention is that, as in ρ , the operation consisting of moving bits of coordinates (x, y, z) to $(x, y, z + \delta \mod w)$ becomes a rotation "to the left" by δ positions if the lane is in a CPU word.

The consequence of the two conventions jointly is that the bytes consisting an input block do not have to be shuffled before being XORed into the state when the state is represented as an array of lanes on a little-endian processor. Similarly, the bytes consisting an output block can be taken as is from the state in the same representation.

Note that the SHA-3 API defined by NIST [28] follows the opposite convention for bit numbering and this is the reason behind the formal bit reordering described in [11].

An example of the formal bit reordering can be found in the files KeccakSpongeIntermediateValues_*.txt in [8].

1.2.1 Some justification for our choice

In this subsection, we detail our choice of conventions concerning the mapping between the bits of the Keccak-f[b] permutation and their representation in terms of w-bit CPU words and in the SHA-3 API defined by NIST [28]. (This part can be safely skipped for readers not interested.)

For the ρ operation to be translated into rotation instructions in the processor, the numbering z must be either an increasing or a decreasing function of the bit numbering in the processor's conventions. So, up to a constant offset, either z=0 is the most significant bit (MSB) and z=w-1 is the least significant bit (LSB), or vice-versa.

The input bits of the hash function come organized as a sequence of bytes. Within each block, the message bit $i=i_{\rm bit}+8i_{\rm byte}$ is going to be XORed with the state bit i. To avoid reordering bits or bytes and to allow a word-wise XORing, the message bit numbering should follow the same convention as the state bit numbering. In particular, if z=0 indicates the MSB (resp. LSB), $i_{\rm byte}=0$ should indicate the most (resp. least) significant byte within a word.

Hence, the choice is bound to either follow the little endian or the big endian convention. We found numbering the bits (or bytes) with increasing powers of 2 (or 256) a bit easier to express than to follow a decreasing rule. Furthermore, in its call for SHA-3, NIST defined a reference platform that happens to be little endian [27]. So we decided to follow the conventions detailed above.

The convention in the Update function of NIST's API is different, and this is the reason for applying the formal bit reordering described in [11]. It formalizes the chosen translation between the two conventions, while having a tiny impact on the implementation. In practice, only the bits of the last byte (when incomplete) of the input message need to be shifted.

1.3 Operation count

For Keccak, the bulk of the processing goes into the Keccak-f permutation and the XOR of the message blocks into the state. For an input message of l bits, the number of blocks to process, or in other words, the number of calls to Keccak-f, is given by:

$$\left\lceil \frac{l+2}{r} \right\rceil$$
.

| | r | С | Relative performance |
|---|-----|------|----------------------|
| | 576 | 1024 | ÷1.778 |
| | 832 | 768 | ÷1.231 |
| 1 | 024 | 576 | 1 |
| 1 | 088 | 512 | ×1.063 |
| 1 | 152 | 448 | ×1.125 |
| 1 | 216 | 384 | ×1.188 |
| 1 | 280 | 320 | ×1.250 |
| 1 | 344 | 256 | ×1.312 |
| 1 | 408 | 192 | ×1.375 |

Table 1.1: Relative performance of Keccak-f[r + c = 1600] with respect to Keccak[].

For an output length n smaller than or equal to the bitrate, the squeezing phase does not imply any additional processing. However, if more output bits are needed, the additional number of calls to Keccak-f for an n-bit output is $\left\lceil \frac{n}{r} \right\rceil - 1$.

When evaluating Keccak, the processing time is dominantly spent in the evaluation of Keccak-f. In good approximation, the throughput of Keccak for long messages is therefore proportional to r for a given permutation width b. For instances with r+c=1600, we will often write performance figures for the default bitrate r=1024. To estimate the performance for another bitrate, Table 1.1 provides the performance relative to the default bitrate. This is valid for long messages; for short messages, the processing time is determined by the required number of calls to Keccak-f (e.g., one when $l \le r-2$).

On a platform supporting operations on w-bit words, each lane can be mapped to such a word. If the platform supports only smaller, m-bit words, each lane has to be mapped to b/m such words. There are different kinds of mappings. As long as the same mapping is applied to all lanes, each bitwise Boolean operation on a lane is translated as b/m instructions on CPU words. The most straightforward mapping is to take as CPU word i the lane bits with $z = mi \dots m(i+1) - 1$. In that case, the b-bit rotations need to be implemented using a number of shifts and bitwise Boolean instructions. Another possible mapping, that translates the b-bit rotations into a series of m-bit CPU word rotation instructions, is introduced in Section 1.5.

If we use a b-bit platform, the evaluation of Keccak-f[b] uses in terms of lane operations

- 76n_r XORs,
- $25n_r$ ANDs and $25n_r$ NOTs, and
- $29n_{\rm r}$ *b*-bit rotations.

For Keccak-f[1600] specifically, this becomes

- 1824 XORs,
- 600 ANDs and 600 NOTs, and
- 696 64-bit rotations.

Some instruction sets propose an instruction that combines the AND and NOT operations in one operation as in χ , i.e., $c \leftarrow (\text{NOT } a)$ AND b. If such an instruction is not available on a CPU, almost 80% of the NOT operations can be removed by applying a *lane complementing transform* as explained in Section 1.6, turning a subset of the AND operations into OR operations.

In [9], we provide a simple implementation (called Simple) that maps a lane to a CPU word of 8, 16, 32 or 64 bits, and hence is suitable for Keccak-f[200], Keccak-f[400], Keccak-f[800] and Keccak-f[1600].

1.4 Memory

In terms of memory usage, Keccak has no feedforward loop, as opposed to many other constructions, and the message block can be directly XORed into the state. This can benefit applications for which the message is formatted on the fly or does not need to be kept after being hashed. This also applies where the hashing API has to implement a message queue. In general a message queue must be allocated, which can be avoided for sponge functions or similar.

The amount of working memory is limited to the state, the round number and some extra working memory for θ and χ . Five w-bit words of extra working memory allow the implementation of θ to compute the XOR of the sheets, while they can hold the five lanes of a plane when χ is being computed.

Examples of implementations that use a limited amount of memory can be found in the Compact and Compact8 implementations in [9]. These implementations also provide an API with a message queue such as Init, Update and Final, for which no extra memory needs to be allocated.

1.5 Bit interleaving

The technique of bit interleaving consists in coding an w-bit lane as an array of s = w/m CPU words of m bits each, with word i containing the lane bits with $z \equiv i \pmod{s}$. This can be applied to any version of Keccak-f to any CPU with word length m that divides its lane length w.

For readability, we now treat the concrete case of 64-bit lanes and 32-bit CPU words. This can be easily generalized. A 64-bit lane is coded as two 32-bit words, one containing the lane bits with even z-coordinate and the other those with odd z-coordinates. More exactly, a lane L[z] = a[x][y][z] is mapped onto words U and V with U[i] = L[2i] and V[i] = L[2i+1]. If all the lanes of the state are coded this way, the bitwise Boolean operations can be simply implemented with bitwise Boolean instructions operating on the words. The main benefit is that the lane translations in ρ and θ can now be implemented with 32-bit word rotations. A translation of L with an even offset 2τ corresponds to the translation of the two corresponding words with offset τ . A translation of L with an odd offset L corresponds to L corresponds to L with an odd offset L corresponds to L with an odd offset L corresponds to L corresponds to L with an odd offset L corresponds to L corresponds to L with an odd offset L corresponds to L corresponds

The bit-interleaving representation can be used in all of Keccak-f where the input and output of Keccak-f assume this representation. This implies that during the absorbing the input blocks must be presented in bit-interleaving representation and during the squeezing the output blocks are made available in bit-interleaving representation. When implementing Keccak in strict compliance to the specifications [10], the input blocks must be transformed to the bit-interleaving representation and the output blocks must be transformed back to the standard representation. However, one can imagine applications that require a secure

sponge function but no interoperability with respect to the exact coding. In that case one may present the input in interleaved form and use the output in this form too. The resulting function will only differ from Keccak by the order of bits in its input and output.

In hashing applications, the output is usually kept short and some overhead is caused by applying the bit-interleaving transform to the input. Such a transform can be implemented using shifts and bitwise operations. For implementing Keccak-f[1600] on a 32-bit CPU, we must distribute the 64 bits of a lane to two 32-bit words. On some platforms, look-up tables can speed up this process, although the gain is not always significant.

Several examples of implementations making use of bit interleaving can be found in [9], notably the Reference32BI and Simple32BI implementations. Also, the Optimized32 implementation can be set to use look-up tables by defining the UseInterleaveTables symbol in KeccakF-1600-opt32-settings.h. Intermediate values to help debug a bit-interleaved implementation can be found in the file KeccakPermutationIntermediateValues32BI.txt in [8]. Finally, KeccakTools can generate code that make use of bit interleaving for any lane size and any (smaller) CPU word size [7].

1.6 The lane complementing transform

The mapping χ applied to the 5 lanes in a plane requires 5 XORs, 5 AND and 5 NOT operations. The number of NOT operations can be reduced to 1 by representing certain lanes by their complement. In this section we explain how this can be done.

For the sake of clarity we denote the XOR operation by \oplus , the AND operation by \wedge , the OR operation by \vee and the NOT operation by $\oplus 1$. Assume that the lane with x=2 is represented its bitwise complement $\overline{a[2]}$. The equation for the bits of A[0] can be transformed using the law of De Morgan ($\overline{a} \wedge \overline{b} = \overline{a \vee b}$):

$$A[0] = a[0] \oplus (a[1] \oplus 1) \wedge (\overline{a[2]} \oplus 1) = \overline{a[0]} \oplus (a[1] \vee \overline{a[2]}).$$

The equation for the bits of A[1] now becomes $A[1] = a[1] \oplus (\overline{a[2]} \land a[3])$. This results in the cancellation of two NOT operations and A[0] being represented by its complement. Similarly, representing a[4] by its complement cancels two more NOT operations. We have

In the computation of the bits of A[4] the NOT operation cannot be avoided without introducing NOT operations in other places. We do however have two options:

$$\overline{A[4]} = \overline{a[4]} \oplus ((a[0] \oplus 1) \land a[1]), \text{ or } A[4] = \overline{a[4]} \oplus (a[0] \lor (a[1] \oplus 1)).$$

Hence one can choose between computing $\overline{A[4]}$ and A[4]. In each of the two cases a NOT operation must be performed on either a[0] or on a[1]. These can be used to compute A[0] rather than $\overline{A[0]}$ or $\overline{A[1]}$ rather than A[1], respectively, adding another degree of freedom for the implementer. In the output some lanes are represented by their complement and the implementer can choose from 4 output patterns. In short, representing lanes a[2] and a[4] by their complement reduces the number of NOT operations from 5 to 1 and replaces 2 or 3

AND operations by OR operations. It is easy to see that complementing any pair of lanes a[i] and $a[(i+2) \mod 5]$ will result in the same reduction. Moreover, this is also the case when complementing all lanes except a[i] and $a[(i+2) \mod 5]$. This results in 10 possible input patterns in total.

Clearly, this can be applied to all 5 planes of the state, each with its own input and output patterns. We apply a complementing pattern (or *mask*) p at the input of χ and choose for each plane an output pattern resulting in P. This output mask P propagates through the linear steps θ , ρ , π (and ι) as a symmetric difference pattern, to result in yet another (symmetric) mask $p' = \pi(\rho(\theta(P)))$ at the input of χ of the next round. We have looked for couples of masks (p,P) that are *round-invariant*, i.e., with $p = \pi(\rho(\theta(P)))$, and found one that complements the lanes in the following θ (x, y) positions at the output of χ or input of θ :

$$P: \{(1,0), (2,0), (3,1), (2,2), (2,3), (0,4)\}.$$

A round-invariant mask P can be applied at the input of Keccak-f. The benefit is that in all rounds 80% of the NOT operations are cancelled. The output of Keccak-f can be corrected by applying the same mask P. The overhead of this method comes from applying the masks at the input and output of Keccak-f. This overhead can be reduced by redefining Keccak-f as operating on the masked state. In that case, P must be applied to the root state 0^b and during the squeezing phase some lanes (e.g., 4 when r = 16w) must be complemented prior to being presented at the output.

The Optimized32 and Optimized64 implementations can be set to use lane complementing by defining the UseBebigokimisa symbol in KeccakF-1600-opt*-settings.h [9]. KeccakTools can generate code that make use of lane complementing for any lane size [7].

Chapter 2

Software

In this chapter, we give an overview of software implementations. Please note that we do not provide extensive benchmark results of Keccak on various platforms. Instead, we refer to eBASH and XBX for detailed and up-to-date benchmarking data on a wide range of platforms [1, 33]. In addition, a summary of the data coming from eBASH and XBX can be found on our web page [6].

2.1 PC and high-end platforms

The platforms supporting 64-bit instructions are in general well-suited for Keccak-f[1600], as a lane can be mapped onto a CPU word. The x86-64 instruction set present in the CPU of most recent PCs is a widely-used example. This instruction set does not have a combined AND+NOT instruction, so lane complementing can be used to reduce the number of NOT instructions. At the time of writing, the permutation Keccak-f[1600] plus XORing 1024 bits takes slightly more than 1600 cycles on a typical x86-64-based machine, hence enabling to absorb long messages with Keccak[] at about 12.6 cycles/byte. Keccak[] performs very well on IA64-based machines too, at around 6-7 cycles/byte [6, 1].

In [9], we provide an implementation suitable for 64-bit CPUs called Optimized64. The code uses only plain C instructions, without assembly nor SIMD instructions. If needed, we have applied lane complementing to reduce the number of NOTs. The operations in the round function have been expanded in macros generated by KeccakTools [7]. We have tried to interleave lines that apply on different variables to enable pipelining, while grouping sets of lines that use a common precomputed value to avoid reloading the registers too often. The order of operations is centered on the evaluation of χ on each plane, preceded by the appropriate XORs for θ and rotations for ρ , and accumulating the parity of sheets for θ in the next round.

In KeccakF-1600-opt64-settings.h, the number of rounds unrolled can be set to any number dividing 24, and the use of lane complementing can be turned on or off.

- For x86-64, the fastest option is to unroll 24 rounds (#define Unrolling 24) and to use lane complementing (#define UseBebigokimisa).
- For IA64, unrolling 6 rounds and not using lane complementing give the fastest results.

2.1.1 Using SIMD instructions

Some platforms do not have plain 64-bit instructions but have single-instruction multiple-data (SIMD) units capable of processing 64-bit or 128-bit data units.

In the family of Intel, AMD and compatible CPUs, a majority supports SIMD instruction sets known as MMX, SSE and their successors. These include bitwise operations on 64-bit and 128-bit registers. Thanks to the symmetry of the operations in Keccak-f, the performance of Keccak can benefit from these instruction sets. For instance, the pandn instruction performs the AND NOT operation bitwise on 128-bit registers (or one register and one memory location), which can be used to implement χ . Such an instruction replaces four 64-bit instructions or eight 32-bit instructions. Similarly, the pxor instruction computes the XOR of two 128-bit registers (or one register and one memory location), replacing two 64-bit XORs or four 32-bit XORs.

While instructions on 128-bit registers work well for θ , π , χ and ι , the lanes are rotated by different amounts in ρ . Consequently, the rotations in ρ cannot fully benefit from the 128-bit registers. However, the rotations in θ are all of the same amount and can be combined efficiently.

If a modern 64-bit CPU such as an Intel Core 2 Duo or an AMD Athlon 64 is restricted to 32-bit instructions (e.g., if the installed operating system uses only a legacy 32-bit mode), using the 128-bit registers of SSE yields faster throughput than using plain 32-bit instructions on such platforms. In such a case, the evaluation of Keccak-f[1600] and XORing 1024 bits takes about 2500 cycles, hence enabling to absorb long messages with Keccak[] at about 20 cycles/byte.

Some older PCs have 32-bit CPUs and a MMX unit with 64-bit SIMD registers. In such a case, it is also interesting to use SIMD instructions. E.g., on an Intel Pentium 3, the evaluation of Keccak-f[1600] and XORing 1024 bits takes about 5200 cycles, hence enabling to absorb long messages with Keccak[] at about 41 cycles/byte. This represents a 40% speedup compared to an implementation using only 32-bit instructions.

In [9], we provide implementations for 64-bit MMX and 128-bit SSE instructions. This can be set in KeccakF-1600-opt64-settings.h, by defining either UseMMX or UseSSE.

2.1.2 SIMD instructions and tree hashing

Parallel evaluations of two instances of Keccak can also benefit from SIMD instructions, for example in the context of tree hashing. In particular, a tree hashing mode as defined in [3] using Keccak[] and parameters (G = LI, H = 1, D = 2, B = 64, C = c = 576) can directly take advantage of two independent sponge functions running in parallel, each taking 64 bits of input alternatively. The final node then combines their outputs.

We have implemented the Keccak-f[1600] permutation with SSE2 instructions using only 64 bits of the 128-bit registers. By definition of these instructions, the same operations are applied to the other 64 bits of the same registers. It is thus possible to evaluate two independent instances of Keccak-f[1600] in parallel on a single core of the CPU.

In this instance of the tree hashing mode, the message bits can be directly input into the 128-bit SSE2 registers without any data shuffling. Using leaf interleaving and a block size B of 64 bits, 64 bits of message are used alternatively by the first sponge function then by the second sponge function. This matches how the data are organized in the SSE2 registers, where 64 bits are used to compute one instance of KECCAK-f[1600] and the other 64 bits to compute the second instance.

On a PC with an Intel Core 2 Duo, we have measured the speed of the double evaluation of Keccak-f[1600] and the XOR of two blocks of 1024 bits. This takes about 2520 cycles, hence

enabling to absorb long messages below 10 cycles/byte. Furthermore, one may consider the case r=1088 and C=c=512, for which the claimed security level is 2^{256} . While losing the power of two for the rate, the final node needs to absorb only one block (DC < r) and the overhead remains reasonable: one extra evaluation of Keccak-f per message. This benefits also to long messages, for which the number of cycles per byte is further reduced to about 9 cycles/byte.

2.1.3 Batch or tree hashing on a graphics processing unit

Using a graphics processing unit (GPU) allows for a high number of instances of Keccak to be computed in parallel. This can be useful in the case of the batch evaluation of several hash functions or of several key streams, or in the case of tree hashing. We are aware of at least two implementations of Keccak on a graphics processing unit [21, 30].

2.2 Small 32-bit platforms

In this section, we consider the implementation of Keccak on a 32-bit platform without a SIMD unit.

A first question to ask ourselves—at least if this choice is possible due to interoperability constraints—is to consider using either Keccak-f[800] or Keccak-f[1600]. On the one hand, for a given security level—hence a given capacity c—a sponge function using Keccak-f[1600] allows processing 800 more input bits per call to the permutation than one using Keccak-f[800]. On the other hand, the memory footprint of Keccak-f[800] is smaller as its the state uses only 25 words of 32 bits, while this number raises to 50 for Keccak-f[1600]. A lower number of words usually means a lower number of load/store operations from/to memory.

If one goes for Keccak-f[1600], the best is to make use of the bit interleaving technique (see Section 1.5). This way, all the operations are done on 32-bit words, including the rotations. If one goes for Keccak-f[800], each lane is simply mapped to a CPU word.

For 32-bit platforms in general, good starting points are the following.

- The target Simple 32BI contains a simple yet optimized implementation of Keccak-f [1600] in C using bit interleaving (see file Keccak-simple 32BI.c).
- The target Simple can be used to instantiate a simple yet optimized implementation of Keccak-f[800] in C, mapping a lane to a CPU word (see file Keccak-simple.c and set cKeccakB to 800 in Keccak-simple-settings.h).

2.2.1 Implementation on a ARM Cortex-M3

The ARM Cortex-M3 is a 32-bit RISC core with 16 registers. Its instruction set contains a combined AND+NOT instruction called bic, so that lane complementing is not necessary. The 32-bit rotations are natively supported and such rotations can be combined with other operations thanks to its barrel shifter.

In [9], we provide an implementation in assembly of Keccak-f [1600] for ARM Cortex-M3; see the files $KeccakF-1600-arm^*$. It uses bit interleaving and takes advantage of the barrel shifter to combine some XORs and rotations in one instruction.

2.3 Small 8-bit platforms

In this section, we consider the implementation of Keccak on a small 8-bit platform. As for small 32-bit platforms, the first question to ask ourselves is which instance of Keccak-f to implement. Generally speaking, larger instances allow having larger rates and processing more input bits for a given capacity c. However, the required memory also increases.

The bit interleaving technique, described in Section 1.5, could be used to represent the 64-bit, 32-bit and 16-bit lanes of Keccak-f[1600], Keccak-f[800] and Keccak-f[400], respectively, with bytes. Byte-level rotations, however, are not often present on 8-bit platforms. Rotations that propagate a bit through the carry are usually natively supported and can be used to perform rotations over lanes of any size, by chaining such rotations through all the bytes of a lane. So, in practice, we did not find bit interleaving to be advantageous in this case, although it may become so on some specific platforms.

If a small capacity of c=160 bits or so is enough for a given application, one can use Keccak-f[200], which uses only 8-bit operations and can be made very compact, as the state only uses 25 bytes.

For small 8-bit platforms in general, the target Compact8 in [9] is a good starting point to implement $K_{\rm ECCAK}$ -f[1600] on an 8-bit processor. Also, the target Simple can be used to instantiate a simple yet optimized implementation of $K_{\rm ECCAK}$ -f[200], mapping a lane to a byte (see file Keccak-simple.c and set cKeccakB to 200 in Keccak-simple-settings.h), although this implementation was not optimized for size.

2.3.1 Estimation on an Intel 8051 processor

We have estimated the performance of Keccak[] on the Intel 8051 microprocessor. The 8051 is an 8-bit processor equipped with an 8-bit data bus and a 16-bit address bus. Initially the 8051 could only address 128 bytes of internal RAM memory, but later versions were released to allow accessing more internal RAM using a technique called Expanded RAM [15]. Nowadays many manufacturers propose variants of the original 8051 with various levels of improvements, including even more powerful addressing modes, security features and shorter instruction cycles. The variant we have selected for our estimation is the 80C51RA/RB/RC microcontroller from Intel [14], with 512 bytes of on-chip RAM, split in 3 segments: 128 bytes of low internal RAM (direct and indirect access), 128 bytes of high internal RAM (indirect access only), and finally 256 bytes of external RAM (indirect access only).

The first problem to solve when implementing Keccak on such a constrained platform like the 80C51RA/RB/RC is the memory mapping. The performance of memory accesses depends on which memory segment is addressed, and so a careful mapping must be done to ensure that most operations are done in the low internal segment. This is particularly difficult for θ for which in-place evaluation requires additional temporary storage. However by following a tight schedule of operations, it is possible to maintain the complete state of Keccak in internal RAM, hence maximizing the performance.

Another problem is the implementation of 64-bit rotations in ρ using 8-bit rotate operations. At first the 8051 does not offer specific instructions to optimize this step, and so rotations are merely done by iterating several times rotate-through-carry instructions (with the exception of 4-bits rotation which can be done by swapping the byte digits). However using efficient memory transfer instructions like XCH that exchanges the accumulator with a byte in memory, it is actually possible to reduce the average number of cycles for rotation to only 4.3 cycles/byte.

| Step | Performance |
|-------------------|---------------|
| θ | 34560 cycles |
| ρ | 20808 cycles |
| π | 0 cycles |
| χ | 30048 cycles |
| L | 384 cycles |
| Words re-ordering | 15000 cycles |
| Total | 100800 cycles |

Table 2.1: Performance estimates of the Keccak-f[1600] permutation on the 8XC51RA/RB/RC

The performance estimates for Keccak including the details for each step are given in Table 2.1, for 24 rounds. One cycle refers to the number of controller clock oscillation periods, which is 12 in the case of our selected variant. It must be noted that the figures are the result of a best-effort paper exercise. Figures for an actual implementation might vary, in particular if it uses specific manufacturer improvements available on the platform.

2.3.2 Implementation on a Atmel AVR processor

The AVR platform uses an 8-bit RISC processor with 32 single-byte registers.

In [9], we provide an implementation in assembly of Keccak-f[1600] for Atmel AVR; see the files Keccak-avr8*.

Chapter 3

Hardware

In this chapter we report on our hardware implementations of Keccak without protection against side-channel attacks. For an overview and links to third-party hardware implementations of Keccak we refer to [5].

3.1 Introduction

Keccak allows to trade off area for speed and vice versa. Different architectures reflect different trade-offs. The two architectures we have investigated and implemented reflect the two ends of the spectrum: a high-speed core and a low-area coprocessor. Thanks to the symmetry and simplicity of its round function, Keccak allows trading off area for speed and vice versa. Different architectures reflect different trade-offs.

We have coded our architectures in VHDL for implementation in ASIC and FPGA [4]. For more details on the VHDL code, refer to the readme.txt file in the VHDL directory.

In these efforts we have focused on two instances of Keccak:

Keccak[r = 1024, c = 576]: the instance of Keccak with default parameter values. It is built on top of Keccak-f[1600], the largest instance of the Keccak-f family.

Keccaκ[r = 40, c = 160]: the instance of Keccak with the smallest instance of the Keccak-f family such that the capacity still provides a security level that is sufficient for many applications. It makes use of Keccak-f[200].

It should be noted that during the design of Keccak particular effort has been put to facilitate the hardware implementation. The round function is based only on simple Boolean expressions and there is no need for adders or S-boxes with complex logic (typically used in many cryptographic primitives). Avoiding these complex sub-blocks allow having a very short critical path for reaching very high frequencies. Another beneficial aspect of Keccak is that, unless intentionally forced, a general architecture implementing Keccak-f and the sponge construction can easily support all variants (rates, capacities) and use cases (MAC, MGF, KDF, PRG) for a given lane size.

3.2 High-speed core

The core presented in this section operates in a stand-alone fashion. The input block is transferred to the core, and the core does not use other resources of the system for performing the




Figure 3.1: The high-speed core

computation. The CPU can program a *direct memory access* (DMA) for transferring chunks of the message to be hashed, and the CPU can be assigned to a different task while the core is computing the hash.

The architecture of the high-speed core design is depicted in Figure 3.1. It is based on the plain instantiation of the combinational logic for computing one Keccak-f round, and use it iteratively.

The core is composed of three main components: the round function, the state register and the input/output buffer. The use of the input/output buffer allows decoupling the core from a typical bus used in a system-on-chip (SoC).

In the absorbing phase, the I/O buffer allows the simultaneous transfer of the input through the bus and the computation of Keccak-f for the previous input block. Similarly, in the squeezing phase it allows the simultaneous transfer of the output through the bus and the computation of Keccak-f for the next output block.

These buses typically come in widths of 8, 16, 32, 64 or 128 bits. We have decided to fix its width to the lane size w of the underlying Keccak-f permutation. This limits the throughput of the sponge engine to w per cycle. This imposes only a restriction if r/w (i.e., the rate expressed in number of lanes) is larger than the number of rounds of the underlying Keccak-f.

In a first phase the high-speed core has been coded in VHDL. Test benches for Keccak-f and the hash function are provided together with C code allowing the generation of test vectors for the test benches. We were able to introduce the lane size as a parameter, allowing us to generate VHDL for all the lane sizes supported by Keccak.

These first VHDL implementations have been tested on different FPGAs by J. Strömbergson [31], highlighting some possible improvements and problems with the tools available from FPGA vendors. We have improved the VHDL code for solving the problems, and this

| Number of round instances | Size | Critical Path | Frequency | Throughput |
|---------------------------|------------|---------------|-----------|--------------|
| n=1 | 48 kgates | 1.9 ns | 526 MHz | 22.44 Gbit/s |
| n=2 | 67 kgates | 3.0 ns | 333 MHz | 28.44 Gbit/s |
| n=3 | 86 kgates | 4.1 ns | 244 MHZ | 31.22 Gbit/s |
| n=4 | 105 kgates | 5.2 ns | 192 MHz | 32.82 Gbit/s |
| n=6 | 143 kgates | 6.3 ns | 135 MHZ | 34.59 Gbit/s |

Table 3.1: Performance estimation of variants of the high speed core of Keccak[r = 1024, c = 576].

has resulted in better results in ASIC as well.

The core has been tested using ModelSim tools. In order to evaluate the silicon area and the clock frequency, the core has been synthesized using Synopsys Design Compiler and a 130 nm general purpose ST technology library, worst case 105°C.

3.3 Variants of the high-speed core

The high-speed core can be modified to optimize for different aspects. In many systems the clock frequency is fixed for the entire chip. So even if the hash core can reach a high frequency it has to be clocked at a lower frequency. In such a scenario Keccak allows instantiating two, three, four or even six rounds in combinatorial logic and compute them in one clock cycle.

An alternative for saving area is to XOR the lanes composing the input blocks directly into the state register and to extract the lanes composing the output blocks directly from it.

3.3.1 Keccak[
$$r = 1024, c = 576$$
]

In this instantiation the width of the bus is 64 bits. The bitrate of 1024 bits and the number of rounds of Keccak-f[1600] being 24 implies a maximum rate of 43 bits per cycle.

The critical path of the core is 1.9 nanoseconds, of which 1.1 nanoseconds in the combinatorial logic of the round function. This results in a maximum clock frequency of 526MHz and throughput of 22.4 Gbit/s. The area needed for having the core running at this frequency is 48 kgate, composed of 19 kgate for the round function, 9 kgate for the I/O buffer and 21 kgate for the state register and control logic.

An alternative without separate I/O buffer allows saving about 8 kgate and decreases the throughput to 12.8 Gbit/s at 500MHz.

Thanks to the low critical path in the combinational logic, it is possible to instantiate two or more rounds per clock cycle. For instance, implementing two rounds gives a critical path of 3 nanoseconds, allowing to run the core at 333MHz reaching a throughput of 28Gbit/s. Such a core will consume 1024 bits every 12 clock cycle, thus the bus width must grow too to keep up with the throughput per cycle. Note that contrary to many cryptographic algorithm, in Keccak the processing does not impose the bottleneck in term of hardware implementation. Table 3.1 summarizes the throughputs for different variants.

3.3.2 Keccak[r = 40, c = 160]

In this instantiation the width of the bus is 8 bits. The bitrate of 40 bits and the number of rounds of Keccak-f[200] being 18 implies a maximum rate of 2.2 bits per cycle.

The critical path of the core is 1.8 nanoseconds, of which 1.1 nanoseconds in the combinatorial logic of the round function. This results in a maximum clock frequency of 555MHz and throughput of 1.23 Gbit/s. The area needed for having the core running at this frequency is 6.5 kgate, composed of 3 kgate for the round function, 3.1 kgate for the state register and control logic and less than 400 gates for the I/O buffer.

An alternative without separate I/O buffer allows saving about 400 gate and decreases the throughput to 0.96 Gbit/s at 555MHz.

3.4 Low-area coprocessor

A different approach can be taken in the design of the hardware accelerator: the core can use the system memory instead of having all the storage capabilities internally. The state of Keccak will be stored in memory and the coprocessor is equipped with registers for storing only temporary variables.

This kind of coprocessor is suitable for smart cards or wireless sensor networks where area is particularly important since it determines the cost of the device and there is no rich operating system allowing to run different processes in parallel.

The architecture is depicted in figure 3.2 where memory buffer labeled with A is reserved for the state, and B is reserved for temporary values. For the width of the data bus for performing memory access different values can be taken. We consider it equal to the lane size as a first assumption, and discuss later the implications if a smaller width is taken.

Internally the coprocessor is divided in two parts, a finite state machine (FSM) and a data path. The data path is equipped with 3 registers for storing temporary values. The FSM computes the address to be read and set the control signals of the data path. The round is computed in different phases.

- First the sheet parities are computed, and the 5 lanes are stored in a dedicated area of the memory.
- The second phase consists in computing the θ transformation, reading all the lanes of the state, and computing the XOR with the corresponding sheet parities. After computing a lane in this way, it is rotated according to ρ and written to the position defined by π . Now the intermediate state is completely stored in the buffer B.
- The last step is to compute χ and add the round constant, ι , to the lane in position (0,0). For doing this the coprocessor reads 3 lanes of a plane from the intermediate state, computes χ and writes the result to the buffer A, reads another element of the intermediate value and writes the new χ , and so on for the 5 elements of the plane.

The computation of one round of the Keccak-*f* permutation takes 215 clock cycles. Out of these, 55 are *bubbles* where the core is computing internally and not transferring data to or from the memory.

In a variant with memory words half the lane size, the number of clock cycles doubles but only for the part relative to read and write, not for the bubbles. In such an implementation one round of Keccak-*f* requires 375 clock cycles.

The buffer A, where the input of the permutation is written and where the output of the permutation is written and the end of the computation has the size of the state (25 times the lane size), while the memory space for storing temporary values has the size of the state times 1.2.




Figure 3.2: The low area coprocessor

| Device | Logic | Registers | Max Freq. | Throughput |
|--------------------|--------------------|--------------|-----------|------------|
| Altera StratixIII | 4684 (38000) ALUTs | 2641 (38000) | 206 MHz | 8.7 Gbit/s |
| EP3SE50F484C2 | | | | |
| Altera Cyclone III | 5770 (10320) LEs | 2641 (10320) | 145 MHz | 6.1 Gbit/s |
| EP3C10F256C6 | | | | |
| Virtex 5 | 1330 (7200) slices | 2640 (28800) | 122 MHz | 5.2 Gbit/s |
| XC5VLX50FF324-3 | | | | |

Table 3.2: Performance estimation of the high speed core of Keccak[r = 1024, c = 576] on different FPGAs, and in brackets the resources available in the different cases.

The low-area coprocessor has been coded in VHDL and simulated using Modelsim. As the core depicted in Section 3.2, the coprocessor has been synthesized using ST technology at 130 nm.

3.4.1 Keccak[
$$r = 1024, c = 576$$
]

In this instantiation the computation of the Keccak-f permutation takes 5160 clock cycles. The coprocessor has a critical path of 1.5 nanoseconds and can run up to 666.7 MHz resulting in a throughput of 132 Mbit/s. The area needed for attaining this clock frequency is 6.5 kgate. If the core is synthesized for a clock frequency limited to 200MHz, the area requirement is reduced to 5 kgate and the corresponding throughput is 39 Mbit/s. In both cases the amount of area needed for the registers is about 1 kgate.

It is interesting to note that the low area coprocessor is capable of reaching higher frequencies than the high speed core.

3.4.2 Keccak
$$[r = 40, c = 160]$$

In this instantiation the computation of the Keccak-f permutation takes 3870 clock cycles. The coprocessor has a critical path of 1.4 nanoseconds and can run up to 714 MHz resulting in a throughput of of 6.87 Mbit/s. The area for attaining this clock frequency is 1.6 kgate, If the core is synthesized for a clock frequency limited to 200MHz (500MHz), the area requirement is reduced to 1.3 (1.4) kgate and the corresponding throughput is 1.9 (4.8) Mbit/s. In both cases the amount of area needed for the registers is in the order of 100 gates.

3.5 FPGA implementations

We have used Altera Quartus II Web Edition version 9 [13] and Xilinx ISE WebPACK version 11.1 [16] to evaluate VHDL with the tools for FPGA. These tools provide estimations of the amount of resources needed and the maximum clock frequency reached.

3.5.1 High-speed core

We report in Tables 3.2 and 3.3 the estimation of the completed place and route for the high-speed core for the large and small versions respectively.

3.5.2 Low-area coprocessor

In the case of FPGA, the estimations are reported in tables 3.4 and 3.5.

| Device | Logic | Registers | Max Freq. | Throughput |
|--------------------|-------------------|-------------|-----------|------------|
| Altera StratixIII | 594 (38000) ALUTs | 333 (38000) | 206 MHz | 412 Mbit/s |
| EP3SE50F484C2 | | | | |
| Altera Cyclone III | 732 (10320) LEs | 333 (10320) | 145 MHz | 290 Mbit/s |
| EP3C10F256C6 | | | | |
| Virtex 5 | 190 (7200) slices | 340 (28800) | 122 MHz | 244 Mbit/s |
| XC5VLX50FF324-3 | | | | |

Table 3.3: Performance estimation of the high speed core of Keccak[r = 40, c = 160] on different FPGAs, and in brackets the resources available in the different cases.

| Device | Logic | Registers | Max Freq. | Throughput |
|--------------------|-------------------|-------------|-----------|-------------|
| Altera StratixIII | 855 (38000) ALUTs | 242 (38000) | 359 MHz | 71.2 Mbit/s |
| EP3SE50F484C2 | | | | |
| Altera Cyclone III | 1570 (5136) LEs | 242 (5136) | 183 MHz | 36.3 Mbit/s |
| EP3C5F256C6 | | | | |
| Virtex 5 | 448 (7200) slices | 244 (28800) | 265 MHz | 52.5 Mbit/s |
| XC5VLX50FF324-3 | | | | |

Table 3.4: Performance estimation of the low area coprocessor of Keccak[r = 1024, c = 576] on different FPGAs, and in brackets the resources available in the different cases.

| Device | Logic | Registers | Max Freq. | Throughput |
|--------------------|-------------------|------------|-----------|------------|
| Altera StratixIII | 131 (38000) ALUTs | 32 (38000) | 359 MHz | 3.7 Mbit/s |
| EP3SE50F484C2 | | | | |
| Altera Cyclone III | 205 (5136) LEs | 30 (5136) | 183 MHz | 1.9 Mbit/s |
| EP3C5F256C6 | | | | |
| Virtex 5 | 62 (7200) slices | 30 (28800) | 265 MHz | 2.7 Mbit/s |
| XC5VLX50FF324-3 | | | | |

Table 3.5: Performance estimation of the low area coprocessor of Keccak[r = 40, c = 160] on different FPGAs, and in brackets the resources available in the different cases.

Chapter 4

Protection against side-channel attacks

If the input to Keccak includes secret data or keys, side channel attacks may pose a threat. In this chapter, we report on Keccak implementations that offer a high level of resistance against power analysis by using the technique of masking (secret sharing).

4.1 Introduction

Sponge functions, among which Keccak, can be used in a wide range of modes covering the full range of symmetric cryptography functions. We refer to [10, 3] for examples. This includes functions that take as argument a secret key. Some other functions do not take a secret key but take as input data that should remain secret such as pseudorandom sequence generators or commit-challenge-response protocols. If such functionality is desired on devices to which an adversary has some kind of physical or logical access, protection against side channel and fault attacks is appropriate [2].

Side channel and fault attacks are attacks that do not exploit an inherent weakness of an algorithm, but rather a characteristics of the implementation. For their security cryptographic primitives inevitably rely on the fact that an adversary does not have access to intermediate computation results. As a consequence, even partial knowledge of intermediate computation results can give a complete breakdown of security, e.g., by allowing computation of the key. However, actual implementations may leak information on these results via characteristics such as computation time, power consumption or electromagnetic radiation. Another condition for the security of cryptographic primitives is that they are executed without faults. An implementation that makes faults, or can be manipulated to make faults, may be completely insecure.

We here concentrate on countermeasures against power analysis and electromagnetic radiation that make use of the algebraic properties of the step functions of Keccak. In the remainder of this chapter we will speak only about power analysis implying also electromagnetic analysis. The main difference between power analysis and electromagnetic analysis is that in the latter the adversary can make more sophisticated measurements and that the physical and electronic countermeasures are different. The countermeasures at the algorithmic level are however the same for both.

As far as timing attacks are concerned, it is straightforward to implement K_{ECCAK} in such a way that its execution time is independent of the input it processes, both in software as in hardware. Moreover, K_{ECCAK} -f does not make use of large lookup tables so that cache timing attacks pose no problem. The interleaving method of Section 1.5 can be implemented with table lookups, which depend on the input message only. Of course, it can also be implemented without any table at all.

Protection against fault attacks can be achieved by countermeasures that are independent of the cryptographic primitive being protected: fault detection at software level, at hardware level and by performing computations multiple times and verifying that the results are equal. Particularly good sources of information on side channel attacks and countermeasures are the proceedings of the yearly Cryptographic Hardware and Embedded Systems (CHES) conferences (http://www.chesworkshop.org/) and the text book [29]. A good source of information on fault attacks and countermeasures are the proceedings of the yearly Fault Diagnosis and Tolerance in Cryptography (FDTC) workshops (http://conferenze.dei.polimi.it/FDTC10/index.html).

4.2 Power analysis

The general set-up of power (and electromagnetic) analysis is that the attacker gets one or more traces of the measured power consumption. If only a single trace suffices to mount an attack, one speaks about simple power analysis (SPA). However, the dependence of the signal in the variables being computed is typically small and obscured by noise. This can be compensated for by taking many traces, each one representing an execution of the cryptographic primitive with different input values. These many traces are then subject to statistical methods to retrieve the key information. These attacks are called differential power analysis (DPA) [22]. An important aspect in these attacks is that the traces must be *aligned*: they must be combined in the time-domain such that corresponding computation steps coincide between the different traces.

In DPA one distinguishes between *first order* DPA and *higher order* DPA. In first-order, the attacker is limited to considering single time offsets of the traces. In *m*-th order the attacker may incorporate up to *m* time offsets in the analysis. Higher-order attacks are in principle more powerful but also much harder to implement [29].

In correlation power analysis (CPA) [12] one exploits the fact that the power consumption may be correlated to the value of bits (or bitwise differences of bits) being processed at any given moment: there is a difference in the expected value of the power consumption. In short, one exploits this by taking many traces and partitioning them in two subsets: in one set, a particular bit, the *target bit*, is systematically equal to 0 and in the other it is equal to 1. Then one adds the traces in each of the two sets and subtracts the results giving the compound trace. If now at any given time the power consumption is correlated to the target bit, one sees a high value in the compound trace. One can use this to retrieve key information by taking a target bit that depends on part of the key and trying different partitions based on the hypothesis for that part of the key. If wrong key guesses result in a partition where the bits of intermediate results are more or less balanced, the compound trace of the correct key guess will stand out. Note that if the power consumption is correlated to bit values or differences, it is also correlated to the Hamming values of words or Hamming distances between words.

Later, more advanced ways to measure the distance between distributions were introduced. In particular, mutual information analysis (MIA) [19] is a generalization of CPA in the sense that instead of just exploiting the fact that different bit values may result in different expected power consumption values, it is able to exploit the difference between the distributions of the power consumption for a bit being 0 or 1 respectively. So in short, when the power consumption distributions of a bit equal to 0 or 1 have equal mean values but different shapes, CPA will not work while MIA may still be able to distinguish the two distributions.

4.2.1 Different types of countermeasures

In the light of power analysis attacks, one must attempt implementing the cryptographic primitives such that the effort (or cost) of the adversary for retrieving the key is too high for her to be interesting. An important countermeasure is implementing the cryptographic primitives such that the power consumption and electromagnetic radiation leak as little as possible on the secret keys or data. Countermeasures can be implemented at several levels:

Transistor level Logical gates and circuits are built in such a way that the information leakage is reduced;

Platform level The platform supports features such as irregular clocking (clock jitter), random insertion of dummy cycles and addition of noise to power consumption;

Program level The order of operations can be randomized or dummy instructions can be inserted randomly to make the alignment of traces more difficult;

Algorithmic level The operations of the cryptographic algorithm are computed in such a way that the information leakage is reduced;

Protocol level The protocol is designed such that it limits the number of computations an attacker can conduct with a given key.

As opposed to protection against cryptographic attacks, protection against side channel attacks is never expected to be absolute: a determined attacker with a massive amount of resources will sooner or later be able to break an implementation. The engineering challenge is to put in enough countermeasures such that the attack becomes too expensive to be interesting. Products that offer a high level of security typically implement countermeasures on multiple levels.

The countermeasures at transistor level are independent of the algorithm to be implemented. Examples are wave dynamic differential logic (WDDL) [32] or SecLib [20]. These types of logic are evolutions of the dual rail logic, where a bit is coded using two lines in such a way that all the logic gates consume the same amount of energy independently of the values. They imply dedicated hardware for cryptography which takes more area, requires dedicated industrialization processes and is in general more expensive. Moreover, while these countermeasures may significantly reduce the information leakage, there always remains some leakage.

The countermeasures at platform level are also independent of the algorithm to be implemented. By randomizing the execution timing, alignment of power traces is made more difficult. The addition of noise to the power consumption increases the required number of traces. The timing randomization is particularly efficient against higher-order DPA as there the signals must be aligned in multiple places and any misalignment severely limits the effectiveness of attack. The addition of noise is also very efficient against attacks that look for dependencies in higher moments of the distributions, such as MIA (see Section 4.5).

The countermeasures at program level are partially dependent on the algorithm to be implemented. Insertion of dummy instructions is possible in any algorithm while changing the order of operations may be easier for some algorithms than for others.

The countermeasures at protocol level are also independent of the algorithm. However, what can be done at this level depends on the requirements of the application and in many cases the possibilities are limited.

Finally, the countermeasures at algorithmic level depend on the basic operations used in the algorithm. This is the type of countermeasures where the choice of operations in the

cryptographic primitive is relevant. One of the countermeasures of this type is that of secret sharing (or masking) and Keccak is particularly well suited for it.

4.2.2 Secret sharing

Masking is a countermeasure that offers protection against DPA at the algorithmic level. It consists of representing variables processed by a cryptographic primitive by two or more shares (as in secret sharing) where the (usually bitwise) sum of the shares is equal to the *native* variable. Subsequently the program or circuit computes the cryptographic primitive using the shares in such a way that the processed variables are independent from the native variables. Whether this is possible depends on the details of the cryptographic primitive and the type of masking. In any case, to achieve independence for a native variable, all but one of its shares must be generated (pseudo-)randomly for each execution of the cryptographic primitive. Clearly, the generation of the shares, the masking operation of the input words and unmasking operation of output, usually considered out of scope of DPA attacks, must also be carefully implemented to limit information leakage. For masking to be effective, the adversary shall have as little information as possible on the value of the shares.

Taking two shares offers protection against first-order DPA under the condition that all processed bits and their joint behavior at any time are independent from native variables. Providing protection against m-th order DPA requires at least m+1 shares.

Computing a linear function λ on the shares of a variable is straightforward. If we represent a native variable x by its shares x_i with $x = \sum_i x_i$ we can compute the shares y_i of $y = \lambda(x)$ by simply applying λ on the individual shares: $y_i = \lambda(x_i)$. A function that consists of the addition of a constant can be performed by adding it to a single share. As all separate operations are performed on shares that are independent of native variables, this provides protection against first-order DPA.

Computing a nonlinear function on the shares assuring all variables processed are independent of native variables depends on the nonlinear function at hand. This is in general a non-trivial problem. We refer to [2] for some examples. In this chapter we limit ourselves to the nonlinear step mapping χ in the round function of Keccak-f:

$$x_i \leftarrow x_i + (x_{i+1} + 1)x_{i+2}$$
 (4.1)

4.3 Software implementation using two-share masking

For software implementations, we have studied the application of masking with two shares, denoted by a and b. The step mapping χ is very similar to nonlinear step mapping γ in BaseKing, the cipher that is the subject of [17] and hence the techniques shown there can be readily applied. We must now compute the shares of x at the lefthand side of Equation (4.1) in such a way that all intermediate variables are independent of native variables. This can be realized by implementing following equations:

$$a_{i} \leftarrow a_{i} + (a_{i+1} + 1)a_{i+2} + a_{i+1}b_{i+2} b_{i} \leftarrow b_{i} + (b_{i+1} + 1)b_{i+2} + b_{i+1}a_{i+2}.$$

$$(4.2)$$

To achieve independence from native variables, the order in which the operations are executed is important. If the expressions are evaluated left to right, it can be shown that all intermediate variables are independent from native variables. The computations of all terms except the rightmost one involve only variables of a single share, hence here independence from x is automatic. For the addition of the mixed term to the intermediate result of the

computation, the presence of a[i] (or b[i]) as a linear term in the intermediate variable results in independence.

At the algorithm level, the effect of the introduction of two shares in the computation of the Keccak-f round function is rather simple. The linear part λ can be executed on the two shares separately, roughly doubling the workload. In the nonlinear step mapping χ the computation of a state word according to Equation (4.1), taking a XOR, AND and NOT instruction, is replaced by the computation of the shares according to Equations (4.2), taking in total 4 XOR, 4 AND and 2 NOT instructions. The addition of round constants ι and addition of input blocks can be performed on one share only.

As the order of execution is important, it is not sufficient to write a program in C or some other high-level language and compile it. The compiler may optimize away the desired order. An option is to compile and inspect the machine code afterwards, but the method that provides the highest level of assurance is to program in assembly language. It is however not sufficient to check only the sequencing of instructions. In general, the operations on two shares of the same variable are preferably executed in registers *physically isolated* from each other. More particularly, the program shall be such that at any given moment there is no register (or bus) content or transition that is correlated to a native variable. For example, if the number of shares is two and a register containing x_0 is loaded with x_1 , the power consumption can depend on the number of bits switched (e.g., on the Hamming weight of $x_0 \oplus x_1$) and it is likely to leak information on x. This can be solved by setting the register to zero in between. Another example is the simultaneous processing of two shares of a variable in different parts of the CPU. As the power consumption at any given moment depends on all processing going on, it depends on both shares and and hence there may be a dependence on the native variable. Clearly, care must be taken when attempting to build side-channel resistant implementations. So if sufficient care is taken, this provide provable resistance against first-order DPA. Higher-order DPA is in principle still possible but as explained in [17, Appendix A] very sensitive to noise and clock jitter. On smart cards, the addition of noise, dummy cycles and clock jitter are typically supported by the platform.

4.4 Hardware using three-share masking

At first sight, masking with two shares may offer protection against first order CPA. However, in dedicated hardware the occurrence of *glitches* [24] may result in correlation between the power consumption and native variables. In a combinatorial circuit computation is typically not monotonous but intermediate signals may switch several times per clock cycle resulting in the (possible transient) appearance of native variables. Hence, due to glitches, two-share masking cannot provide provable protection against first order CPA.

A solution to this problem was proposed in [25]. We refer to [25, 26] for an in-depth treatment and limit ourselves here to the explanation of the basic concept. The simple but powerful idea is to take as many shares as needed such that in any computation at least one of the shares is not taken as input. In this way, all intermediate variables are guaranteed to be independent from native variables for the same reason that the one-time pad offers perfect secrecy. For linear functions two shares are sufficient to realize this. For nonlinear functions, the required number of shares depends on the particular function.

As can be read in [25, 26], for most cryptographic algorithms, the mere number of shares required makes the application of this technique very expensive. Fortunately, χ with its simple structure is well suited for applying this technique and only three shares are required. This suitability is quite unique among cryptographic primitives. If we denote the shares by

a, *b* and *c*, a possible computation of χ applied to three shares is given by:

$$a_{i} \leftarrow b_{i} + (b_{i+1} + 1)b_{i+2} + b_{i+1}c_{i+2} + c_{i+1}b_{i+2}$$

$$b_{i} \leftarrow c_{i} + (c_{i+1} + 1)c_{i+2} + c_{i+1}a_{i+2} + a_{i+1}c_{i+2}$$

$$c_{i} \leftarrow a_{i} + (a_{i+1} + 1)a_{i+2} + a_{i+1}b_{i+2} + b_{i+1}a_{i+2}.$$

$$(4.3)$$

Clearly, the computation of each share takes as input only components of the other two shares and it provides provable security against first order CPA, even in the presence of glitches.

The three lines of Equation (4.3) are equivalent and only differ in the input shares and output share. We denote the computation of the nonlinear step χ' resulting in a share given the two other shares by, e.g., $a \leftarrow \chi'(b,c)$.

In the following subsections we present two architectures that make use of three shares. Other architectures can be derived based on different partitioning or sequences of the computational steps composing the round. For instance a low area coprocessor, as those presented in Section 3.4, can be protected using the secret sharing techniques.

4.4.1 One-cycle round architecture

A first architecture computes one round in one clock cycle and is depicted in Figure 4.1.

Before processing, the three shares a, b and c are generated from a random source. As the initial state of Keccak should be set equal to zero implying a + b + c = 0, the shares a and b can be generated randomly and c computed as c = a + b. The hardware for generating the three shares is out of the scope of our study, we just consider them as input to the core.

The combinatorial logic implements the round function and input data block absorbing. It has a layered structure. In a first (linear) layer, the absorbing of the input data block, $D_{\rm in}$, is implemented by adding it to one of the shares and then λ is applied to the three shares by three separate combinatorial blocks. In a second layer, the nonlinear step mapping χ is computed on the output of the first layer according to Equations (4.3) by three separate combinatorial blocks implementing χ' . Each block takes as input two shares and generates one share. The blocks only differ in the presence of ι in the leftmost block as ι only needs to be added to a single share.

We can estimate the cost of the secret sharing technique in terms of silicon area by comparing this architecture with our unprotected one-cycle round architecture in Section 3.2, based on Figure 4.1 and Equations (4.3). The number of registers required for storing the state is three times larger than the unprotected version. The cost of the linear part is three times larger as well. Regarding the nonlinear part, we have three blocks χ' instead of one χ and the cost of every χ' is also larger than that of χ . While χ requires basically an AND gate, a NOT and a XOR for each bit of the state, χ' requires three AND gates, one NOT and three XOR for a single share. So roughly, the protected nonlinear part is expected to be nine times larger than the unprotected χ .

4.4.2 Three-cycle round architecture

A second architecture reduces the amount of silicon at the cost of performance and is depicted in Figure 4.2. Since the round logic is composed of three equal blocks for λ and three equal blocks for χ' , we instantiate only one block of each and try to use them as much as possible.

Instead of three registers this architecture requires four registers, some multiplexing and a careful schedule. The schedule has some similarity to pipelining techniques.

For explaining the schedule we refer to Table 4.1. We use $\lambda(R_0)$ to denote the application of λ to register R_0 , and $\chi'(R_1, R_2)$ to denote the application of χ' using registers R_1 and R_2




Figure 4.1: Protected architecture computing one round in one clock cycle.

| cycle | R_0 | R_1 | R ₂ | R_3 |
|-------|-------------------------|-------------------------|-------------------------|-------------------------|
| 0 | - | - | - | - |
| 1 | input, A | - | _ | - |
| 2 | input, B | $\lambda(A+D_{\rm in})$ | _ | - |
| 3 | input, C | $\lambda(B)$ | $\lambda(A+D_{\rm in})$ | - |
| 4 | $\chi'(R_2,R_1),C'$ | $\lambda(C)$ | $\lambda(B)$ | $\lambda(A+D_{\rm in})$ |
| 5 | $\chi'(R_2,R_1),A'$ | $\lambda(C')$ | $\lambda(C)$ | $\lambda(A+D_{\rm in})$ |
| 6 | $\chi'(R_2,R_3),B'$ | $\lambda(A')$ | $\lambda(C')$ | $\lambda(C)$ |
| 7 | $\chi'(R_2,R_1),B''$ | $\lambda(B')$ | $\lambda(A')$ | $\lambda(C')$ |
| 8 | $\chi'(R_2,R_1),C''$ | $\lambda(B'')$ | $\lambda(B')$ | $\lambda(C')$ |
| 9 | $\chi'(R_2, R_3), A''$ | $\lambda(C'')$ | $\lambda(B'')$ | $\lambda(B')$ |
| 10 | $\chi'(R_2, R_1), A'''$ | $\lambda(A'')$ | $\lambda(C'')$ | $\lambda(B'')$ |
| 11 | $\chi'(R_2,R_1),B'''$ | $\lambda(A''')$ | $\lambda(A'')$ | $\lambda(B'')$ |
| 12 | $\chi'(R_2,R_3),C'''$ | $\lambda(B''')$ | $\lambda(A''')$ | $\lambda(A'')$ |

Table 4.1: The content of the registers during the computation.

as inputs. The three initial values of the shares are indicated as A, B and C. The values after the first round by A', B' and C', after the second round by A'', etc. At clock cycle zero the registers do not yet contain any relevant data, as indicated with a -. Instead of loading the three shares in parallel, as done in the previous architecture, one share per clock cycle is loaded into register R_0 during the three initial clock cycle.

The content of R_1 is in general the result of λ applied to R_0 , with the D_{in} enabled and XORed at the input of $\lambda(R_0)$ before an initial Keccak-f round.

The content of R_0 is the result of χ' applied to either the couple R_2 and R_1 or R_2 and R_3 . While in the one-cycle round architecture ι was applied only to one of the three shares, here it is applied to all of the three shares. This does not change the result of the computation and simplifies the control logic.

The registers R_2 and R_3 are used for maintaining the values required for the computation of χ' . In the second clock cycle B is loaded into R_0 , and R_1 receives $\lambda(A)$. In the third clock cycle C is loaded into R_0 , R_1 receives $\lambda(A)$ and $\lambda(A)$ is moved from R_1 to R_2 . In the fourth clock cycle no more shares need to be loaded. Instead R_0 receives χ' applied to the content of R_1 and R_2 , which means $\lambda(A)$ and $\lambda(B)$. It follows that R_0 now contains the share C' after the first round. R_1 receives $\lambda(C)$, $\lambda(B)$ is moved from R_1 to R_2 and $\lambda(A)$ is moved from R_2 to R_3 . In the fifth clock cycle R_0 receives χ' applied to the content of R_1 and R_2 , being $\lambda(B)$ and $\lambda(C)$ and hence contains the share A after the first round. R_1 receives $\lambda(C')$, $\lambda(C)$ is moved from R_1 to R_2 , while $\lambda(A)$ remains in R_3 . Since shares A' and C' as input of the second round have already been computed, in the next clock cycle share B' must be computed, requiring $\lambda(A)$ and $\lambda(C)$, and they are in R_1 and R_3 respectively before clock cycle five. Thus we have described also what will be computed in the next clock cycle and this is basically the end of a round. Starting from the next clock cycle there will be a loop of three clock cycles where the alternation of data contained in the registers are those for computing a round.

Using this circuit, a Keccak computation takes 3 initial cycles plus 24 cycles per execution of Keccak-f.




Figure 4.2: Protected architecture computing one round in three clock cycles.

| Core | Size | Frequency | Throughput |
|-----------------------|------|-----------|------------|
| (r = 1024) | KGE | MHz | Gbit/s. |
| Unprotected one-cycle | 48 | 526 | 22.4 |
| One-cycle (fast) | 183 | 500 | 21.3 |
| One-cycle (compact) | 127 | 200 | 8.5 |
| Three-cycle (fast) | 115 | 714 | 10.1 |
| Three-cycle (medium) | 106 | 500 | 7.1 |
| Three-cycle (compact) | 95 | 200 | 2.8 |

Table 4.2: Performance and gate count of the different implementations

4.4.3 Synthesis results

We have implemented both architectures in VHDL and synthesized it for understanding the maximum frequency and silicon area demand. We have used the same technology library adopted for the unprotected implementation reported in Section 3.2, a 130 nm general purpose library from STMicroelectronics, and the Synopsys Design Compiler.

Table 4.2 summarizes the gate count and performance numbers of the different implementations, together with the numbers for the one-cycle round unprotected architecture described in Section 3.2. It gives the total gate count, the maximum frequency and the throughput assuming a bitrate of 1024 bits.

The gate count figures are in line with the estimations made in Section 4.4.1. All implementations have an I/O buffer of 9 Kilo gate equivalent (KGE) for connecting the core to a system bus. This allows to load the I/O buffer, 64-bit per clock cycle, simultaneously with the application of Keccak-f on the previous input block.

In the three-cycle round architecture the computation of one round is executed in 3 clock cycles and the initialization of the state requires also 3 clock cycles. Thanks to the very short critical path, the fast variant of the three-cycle round architecture can reach 714 MHz, still resulting in a very competitive speed of 10 Gbit/s. If we compare the one-cycle round and the three-cycle round architectures both running at 500Mhz, we can see that the three-cycle requires 40% less silicon area at a cost of a throughput reduction by a factor 3.

It is interesting to note that the strategy adopted in the design of the function allows implementing this countermeasure with a cost consisting only of silicon area with almost no penalty in terms of throughput: it is reduced only by 5%, from 22.4 Gbit/s. to 21.3 Gbit/s. when using a rate of 1024 bits.

4.5 Computing in parallel or sequentially?

The use of three shares computed at different times gives a provable resistance against first-order DPA, but not against DPA of higher order. Higher-order DPA involves taking measurements corresponding to the computation of the different shares, and this task is more complex due to its higher sensitivity to noise than first-order DPA [17, Appendix A].

In the architectures presented in the previous sections, several computations take place in parallel and the power consumption at a given time depends on all these computations. For this reason, these architectures are not provably resistant against first-order DPA. For instance, in [26], the authors analyzed an implementation with three shares working in parallel. Storing the three shares in a register causes a power consumption to depend on the three shares simultaneously. In the absence of noise, the distribution of the consumption

depends on the native variable being stored and hence is vulnerable to MIA. In this section, we explain that the introduction of noise has a much stronger impact on the feasibility of this attack for a two-share or three-share implementation than for an unmasked one, and that the qualitative difference is similar to the one between first-order and higher-order DPA. Furthermore, we argue that a masked implementation has per construction a higher algorithmic noise level than an unmasked one.

With three shares, a native bit equal to 0 (resp. 1) can be represented as 000, 011, 101 or 110 (resp. 001, 010, 101 or 111). If the three shares are processed simultaneously, the power consumption can leak the Hamming weight of the shares, which means 0 or 2 for a native bit 0, or 1 or 3 for a native bit 1. Clearly, these two distributions are different and can be distinguished.

We start the discussion by constructing a simple model where the power consumption is equal to the Hamming weight plus an additive Gaussian noise of standard deviation σ , expressed in the same unit as the impact of the Hamming weight on the power consumption. In this model, the distribution of the power consumption for three shares is as in Figure 4.3(c), with $\sigma=0.2$. Similarly, we can look at an unmasked implementation, where the power consumption is 0 or 1 plus the Gaussian noise, and at masking with two shares. In this last case, the Hamming weight is 0 or 2 for a native bit 0 (represented as 00 or 11) or 1 for a native bit 1 (represented as 01 or 10). The two cases can be found in Figures 4.3(a) and 4.3(b), respectively.




Figure 4.3: Distribution of the power consumption for a simple model. The solid line shows the distribution for a bit with native value 0 and the dashed line for a bit 1. Sub-figures (a), (b) and (c) show the case of one, two and three shares, respectively, with a noise level of $\sigma = 0.2$. Sub-figures (d), (e) and (f) follow the same model with $\sigma = 2$.

Following this model, we compute the number of samples that are needed to distinguish between the distribution for a native bit 0 and the distribution for a native bit 1. We follow the same reasoning as in [17, Appendix A]. The number z of samples needed to distinguish one distribution over the other is inversely proportional to the Kullback-Leibler divergence

between the two distributions [23], z = 1/D(f|g) with

$$D(f|g) = \int f(x)(\log(f(x)) - \log(g(x))dx.$$

In this model, the scaling of z as a function of σ is different for one, two or three shares. For one share, $z\sim 2\sigma^2$ samples are needed to distinguish a native bit 0 from a native bit 1 from unmasked values, whereas about z^2 samples are needed for the same noise level when two shares are used, and this number grows to about z^3 samples for three-share masking. Hence, the difference between the one-share and two-share and three-share implementation is qualitatively the same as the one between first-order and second-order and third-order DPA.

The real-world behavior is likely to differ from this simple model. Nevertheless, we expect a significantly higher sensitivity to noise for three shares than for one. Qualitatively, the three pairs of distributions are different. For one share, the mean is different for native bits 0 and 1. For two shares, the two distributions have the same mean but a different variance. For three shares, the two distributions have the same mean and variance; they differ only starting from their third moment. Figures 4.3(d), 4.3(e) and 4.3(f) illustrate this with the simple model and a higher noise level $\sigma = 2$.

So far, we have assumed that the three levels of masking are subject to the same noise level σ . However, the masking by itself introduces noise, as m-1 shares are randomly and independently chosen for every execution of the algorithm. The very dependence of the processing of a bit in the power consumption that an attacker exploits turns against her, as it becomes an additional source of noise due the randomization. For instance, in the one-cycle-one-round implementation of Keccak-f[1600] with three shares, the noise due to the Hamming weight of 3200 random bits must be compared to a small number of unmasked bit values that a differential power analysis attempts at recovering.

Bibliography

design/mcs51/er 51.htm.

[1] D. J. Bernstein and T. Lange (editors), eBACS: ECRYPT Benchmarking of cryptographic systems, http://bench.cr.yp.to, accessed 21 December 2010. [2] G. Bertoni, J. Daemen, M. Peeters, and G. Van Assche, Note on side-channel attacks and their countermeasures, Comment on the NIST Hash Competition Forum, May 2009, http: //keccak.noekeon.org/NoteSideChannelAttacks.pdf. [3] _____, Cryptographic sponge functions, January 2011, http://sponge.noekeon.org/. [4] _____, Hardware implementation of Keccak in VHDL, 2011, available from http:// keccak.noekeon.org/. [5] _____, Keccak hardware performance figures page, 2011, http://keccak.noekeon.org/hw_ performance.html. [6] ______, Keccak software performance figures page, 2011, http://keccak.noekeon.org/sw_ performance.html. [7] _____, KeccakTools software, January 2011, http://keccak.noekeon.org/. [8] _____, Known-answer and Monte Carlo test results, 2011, available from http://keccak. noekeon.org/. [9] _____, Reference and optimized implementations of Keccak, 2011, available from http: //keccak.noekeon.org/. [10] _____, The Keccak reference, January 2011, http://keccak.noekeon.org/. [11] _____, The Keccak SHA-3 submission, January 2011, http://keccak.noekeon.org/. [12] E. Brier, C. Clavier, and F. Olivier, Correlation power analysis with a leakage model, CHES (M. Joye and J.-J. Quisquater, eds.), Lecture Notes in Computer Science, vol. 3156, Springer, 2004, pp. 16-29. [13] Altera corporation, Quartus II web edition software, http://www.altera.com. [14] Intel Corporation, Intel 8XC51RA/RB/RC hardware description, http://www.intel.com/ design/mcs51/manuals/272668.htm. _____, Intel MCS 51/251 microcontrollers - expanded RAM, http://www.intel.com/

[16] Xilinx corporation, ISE WebPACK software, http://www.xilinx.com.

- [17] J. Daemen, M. Peeters, and G. Van Assche, *Bitslice ciphers and power analysis attacks*, Fast Software Encryption 2000 (B. Schneier, ed.), Lecture Notes in Computer Science, vol. 1978, Springer, 2000, pp. 134–149.
- [18] G. Gielen and J. Figueras (eds.), 2004 design, automation and test in Europe conference and exposition (DATE 2004), 16-20 February 2004, Paris, France, IEEE Computer Society, 2004.
- [19] B. Gierlichs, L. Batina, P. Tuyls, and B. Preneel, *Mutual information analysis*, CHES (E. Oswald and P. Rohatgi, eds.), Lecture Notes in Computer Science, vol. 5154, Springer, 2008, pp. 426–442.
- [20] S. Guilley, P. Hoogvorst, Y. Mathieu, R. Pacalet, and J. Provost, *CMOS structures suitable for secured hardware*, in Gielen and Figueras [18], pp. 1414–1415.
- [21] G. Hoffmann, Keccak implementation on GPU, 2010, http://www.cayrel.net/spip.php? article189.
- [22] P. C. Kocher, J. Jaffe, and B. Jun, *Differential power analysis*, Advances in Cryptology Crypto '99 (M. Wiener, ed.), Lecture Notes in Computer Science, vol. 1666, Springer, 1999, pp. 388–397.
- [23] S. Kullback and R. A. Leibler, *On information and sufficiency*, Ann. Math. Statist. **22** (1951), no. 1, 79–86, http://projecteuclid.org/euclid.aoms/1177729694.
- [24] S. Mangard, N. Pramstaller, and E. Oswald, *Successfully attacking masked AES hardware implementations*, CHES (J.R. Rao and B. Sunar, eds.), Lecture Notes in Computer Science, vol. 3659, Springer, 2005, pp. 157–171.
- [25] S. Nikova, V. Rijmen, and M. Schläffer, *Secure hardware implementation of nonlinear functions in the presence of glitches*, ICISC (P. J. Lee and J. H. Cheon, eds.), Lecture Notes in Computer Science, vol. 5461, Springer, 2008, pp. 218–234.
- [26] _____, Secure hardware implementation of nonlinear functions in the presence of glitches, Journal of Cryptology, 2010, to appear.
- [27] NIST, Announcing request for candidate algorithm nominations for a new cryptographic hash algorithm (SHA-3) family, Federal Register Notices **72** (2007), no. 212, 62212–62220, http://csrc.nist.gov/groups/ST/hash/index.html.
- [28] _____, ANSI C cryptographic API profile for SHA-3 candidate algorithm submissions, revision 5, February 2008, available from http://csrc.nist.gov/groups/ST/hash/sha-3/Submission_Reqs/crypto_API.html.
- [29] E. Oswald S. Mangard and T. Popp, *Power analysis attacks revealing the secrets of smart-cards*, Springer-Verlag, 2007.
- [30] G. Sevestre, Keccak tree hashing on GPU, using Nvidia Cuda API, 2010, http://sites.google.com/site/keccaktreegpu/.
- [31] J. Strömbergson, *Implementation of the Keccak hash function in FPGA devices*, http://www.strombergson.com/files/Keccak_in_FPGAs.pdf.
- [32] K. Tiri and I. Verbauwhede, *A logic level design methodology for a secure DPA resistant ASIC or FPGA implementation*, in Gielen and Figueras [18], pp. 246–251.

[33] C. Wenzel-Benner and J. Gräf, XBX: keeping two eyes on embedded hashing, http://xbx.das-labor.org/, accessed 21 December 2010.