

Note on KECCAK parameters and usage

The KECCAK sponge function family is characterized by three parameters: the bitrate r , the capacity c (where $r + c$ is the width of the underlying permutation) and the diversifier d . We propose in [5] four instances that can be taken as functions for the four (fixed) output lengths NIST requires for SHA-3 and a variable-output-length instance, denoted by KECCAK[], with default values for the parameters. Section 1 below recalls the KECCAK offering: its parameters, security claim and design strategy, and our proposal to NIST.

Whilst we are happy with our choice, there are other valid parameter choices that NIST or others may prefer. In this note we discuss our choice of parameters and other possible ways of using the KECCAK family.

With its arbitrary length, the output of KECCAK can be truncated at the length requested by the user. In Section 2 we discuss how using a single function has clear advantages and, if needed, simple ways to achieve diversification. The capacity c is the security parameter of KECCAK and the use of a single instance with fixed capacity puts a ceiling on the achievable security level. We explain in Section 2 why this limit for KECCAK[] is high enough not to be a problem.

In Section 3 we examine the issues of letting the user choose the capacity c while keeping $c + r = 1600$, which allows trading off claimed security for speed by increasing c and decreasing r , or vice versa.

In Section 4 we explain the choice of KECCAK- f [1600] in our standard proposal and discuss where the KECCAK- f permutations with other widths may be adequate.

A way to exploit parallelism is to apply tree hashing. This is especially relevant on modern CPUs with their multiple cores and SIMD architecture within a core. In Section 5 we explain that a tree hashing mode calling KECCAK as a compression function can take advantage of both.

Finally, we address the question of migration possibilities to a more secure version, should KECCAK be chosen as a standard and a weakness be discovered later. We propose in Section 6 two techniques based on input pre-processing with very limited impact on implementations.

1 The KECCAK offering

As defined in [5], KECCAK is a family of sponge functions with members KECCAK[r, c, d] characterized by three parameters:

- bitrate r ,
- capacity c and
- diversifier d .

The sum $r + c$ determines the width of the KECCAK- f permutation used in the sponge construction and is restricted to values in $\{25, 50, 100, 200, 400, 800, 1600\}$. The diversifier value satisfies $0 \leq d < 256$.

The sponge construction uses $r + c$ bits of state, of which r are updated with message bits between each application of KECCAK- f during the absorbing phase and output during the squeezing phase. The remaining c bits are not directly affected by message bits, nor are they taken as output.

The purpose of the diversifier is to provide diversification, i.e., two instances of KECCAK with two different values of d behave as two independent hash functions (even with same values of r and c). See Section 2.2 for a discussion.

1.1 The security claim and design strategy

KECCAK allows one to choose its security parameter c independently from the output length. We express our security claim for KECCAK in [5] as a *flat sponge claim* [2]. This type of claim implies that the expected complexity of any attack should be the same as for a random oracle, up to $2^{c_{\text{claim}}/2}$. The value c_{claim} is called the *claimed capacity* and fully determines the claimed security level of the variable-output-length function.

The design philosophy underlying KECCAK is the *hermetic sponge strategy*: adopting the sponge construction using a permutation that should not have structural distinguishers [6, Chapter 4]. In this approach, we can make a flat sponge claim with claimed capacity c_{claim} equal to the parameter c in the construction and trade in claimed security level for speed by increasing c and decreasing r accordingly.

Additional information on the security claim and design strategy is given in [4].

1.2 Our proposal for SHA-3

In [14], NIST requires the candidate algorithms to support at least four different output lengths $n \in \{224, 256, 384, 512\}$ with associated security levels. Hence, we have defined four fixed-output-length variants (where $\lfloor \cdot \rfloor_n$ indicates truncation to the first n bits):

- $n = 224$: $\lfloor \text{KECCAK}[r = 1156, c = 448, d = 28] \rfloor_{224}$
- $n = 256$: $\lfloor \text{KECCAK}[r = 1088, c = 512, d = 32] \rfloor_{256}$
- $n = 384$: $\lfloor \text{KECCAK}[r = 832, c = 768, d = 48] \rfloor_{384}$
- $n = 512$: $\lfloor \text{KECCAK}[r = 576, c = 1024, d = 64] \rfloor_{512}$

The capacity values were chosen to meet the requirement that (second) preimage resistance should be 2^n (with n the output length). The different diversifier values $d = n/8$ address a requirement expressed by NIST on the hash forum mailing list [13, 23-Jun-2008], that a hash function with a given output length should not be the prefix of another one with larger output length.

In addition, we proposed $\text{KECCAK}[\]$ (with default parameters), where the user may truncate the output at the desired output length. The default bitrate $r = 1024$ is a power of two to ease data alignment and the resulting capacity is $c = 1600 - 1024 = 576$. The default value for the diversifier d is 0.

2 Letting the user choose the output length

In many use cases of hash functions the output length is determined by the application. This is the case for key derivation functions and several important public key signature and key establishment schemes, for instance the widely used RSA padding schemes [11, 12]. In those cases, either the output must be truncated or an additional construction called mask generating function (MGF) must be applied to provide longer outputs [11, 12].

Consider a protocol to be designed with the requirement of a specific digest length ℓ . When using a hash function family that consists of a set of instances with different output lengths and ℓ is not among them, one must first choose an instance and either truncate or specify an MGF construction. When using a variable-output-length hash function, no such choice must be made and it suffices to truncate the output to the desired length. The advantage of a variable-output-length hash function becomes even more important if a protocol or application requires digests whose length is a parameter of the protocol.

2.1 What about the security level?

Traditionally, hash function users expect a security level that matches its output length: $2^{n/2}$ for collision-resistance and 2^n for (second) preimage resistance. As stated in Section 1.1, a variable-output-length hash function with a claimed capacity c_{claim} shall resist to any attack with complexity below $2^{c_{\text{claim}}/2}$, but nothing is claimed above this level. Hence, the value $2^{c_{\text{claim}}/2}$ acts as a ceiling for the security level.

This ceiling poses no problem if high enough. For instance, the ceiling is at 2^{288} in the case of KECCAK[], as it has capacity $c = 576$. Consider an application where we need a 512-bit output. Traditionally, a (second) preimage resistance level of 2^{512} would be expected, while for KECCAK[] with output truncated to 512 bits a security level of *only* 2^{288} is claimed. However, the difference between these two security levels is purely philosophical with no practical implications whatsoever. By translating these computation complexities into physical quantities such as time or energy, both are simply out of reach and will remain so in the foreseeable future [8].

2.2 What about diversification?

A single function for all output lengths may pose problems when a scheme requires that different output lengths are generated with different hash function instances. Diversification is actually a requirement that may arise for other aspects than different output lengths. A scheme or protocol may require different hash function instances even if their output lengths are the same. In KECCAK[] the diversifier is fixed to 0 and as such does not appear to address this requirement. However, diversification can be established at very small cost using a well-established technique called *domain separation*. Domain separation is an efficient means to construct different function instances from a single underlying function. If the underlying function is secure, the derived functions can be considered as independent functions.

One can implement domain separation by appending or prepending different constants to the input for each of the function instances: $f_i(M) = \text{KECCAK}(M||C_i)$ or $f_i(M) = \text{KECCAK}(C_i||M)$. As a concrete example, one can use a convention based on namespaces such as KECCAKNS for diversification [6, Section 6.3]. The use of the diversifier d is actually a built-in way to achieve domain separation.

3 Letting the user choose the capacity

For standardization, one option is to impose a small set of (or just a single instance of) parameter values. Another option is to allow the user to freely choose them. We consider in particular the case where a user can freely¹ choose the capacity of KECCAK with $r = 1600 - c$ so that the width of KECCAK- f is fixed. In this section, we describe the advantages and disadvantages of this option.

As explained in [4], the hermetic sponge strategy allows the user to trade in speed for claimed security, or vice versa, by choosing the capacity. Relative performance estimates for various (r, c) pairs are listed in Table 1.

If the user decides to lower the capacity to $c = 256$, providing a claimed security level equivalent to that of AES-128, the performance will be 31% greater than for the default value $c = 576$. If the user wants an output truncated to 512 bits to provide the traditionally expected (second) preimage resistance of 2^{512} by setting the capacity to $c = 1024$, she can do this at the cost of a performance decreased by 78%.

A variable capacity can also result in important efficiency gain in applications dealing with (mostly) short messages. Consider for example an application with messages that are exactly 1024 bits long. The padding will extend these messages

¹We limit the choice to multiples of 8 to avoid intra-byte bit shuffling.

r	c	Relative performance
576	1024	$\div 1.778$
832	768	$\div 1.231$
1024	576	1
1088	512	$\times 1.063$
1152	448	$\times 1.125$
1216	384	$\times 1.188$
1280	320	$\times 1.250$
1344	256	$\times 1.312$
1408	192	$\times 1.375$

Table 1: Relative performance of $\text{KECCAK}[r, c]$ with respect to $\text{KECCAK}[\]$.

by 32 bits resulting in a two-block message and hence applying $\text{KECCAK}[\]$ results in two calls to $\text{KECCAK-}f$. If we decrease the capacity by 32 bits to 544 (still providing an astronomical security level), a padded message fits in a single block and only one call to $\text{KECCAK-}f$ must be made.

In [8] we provide a simple application to help determine the capacity value and output length given required security levels for collision-resistance and (second) pre-image resistance.

3.1 What about the indifferenciability?

The sponge indifferenciability proof of [3] assumes the capacity is fixed and does not prove indifferenciability of a set of sponge functions calling the same underlying function with different capacity values. However, for the padding function used in KECCAK , we have proven an indifferenciability theorem in [6, Section 3.1.2] for the case of variable capacity and diversifier values. We refer to that section for a more in-depth explanation.

3.2 What about the implementation cost?

An argument against tunable parameters in a standard is that it makes implementations more expensive, as they usually have to support all parameter values to fully implement the standard. However, for KECCAK , the main implementation cost is for the $\text{KECCAK-}f[1600]$ permutation that is the same for all capacity values. The additional cost of the variable capacity value consists of the required support for the configurable bitrate r determining the length of the message blocks to be XORed into the state and of the coding of the bitrate in the padding. The cost of supporting a variable capacity value with a fixed state width is therefore quite limited.

3.3 What about the burden of choice for the user?

Another argument against tuneable parameters in a standard is that it puts the burden of choice on the hash function user, typically a designer of a protocol or scheme. In particular, the choice of the capacity value determines a ceiling to the security level that the sponge function provides and one could argue that the user usually does not have the responsibility or the expertise to make that choice. In our opinion, the security claim of KECCAK is easy to understand and the user can be guided in the choice of the capacity by some simple recommendations. For example, one could fix a maximum capacity value c_{\max} and recommend taking a capacity equal to twice the output length for output lengths below $c_{\max}/2$ bits and a capacity equal to c_{\max} bits for higher output lengths.

4 Parameters of the KECCAK- f permutation

All KECCAK members we propose for standardization make use of the same permutation: KECCAK- f [1600]. A single implementation of this permutation supports all the proposed variants, hence reducing cost, for instance, in hardware implementations. Furthermore, the choice of KECCAK- f [1600] favors 64-bit CPUs and yet remains efficient on 32-bit (and smaller) processors.

Software implementations of KECCAK- f use bitwise Boolean operations and (cyclic) shifts on CPU words. A typical implementation maps each lane to a CPU word, resulting in the state of KECCAK represented in 25 words of 64 bits each. The choice of the lane size therefore favors CPUs with the corresponding word size. Specifically, the implementation of KECCAK- f [1600] on a 64-bit CPU can exploit 64-bit wide Boolean operations and 64-bit rotations.

Because of the bit-oriented design of KECCAK- f , other approaches are possible. For instance, KECCAK- f [1600] can be efficiently implemented on a 32-bit CPU by using the *bit interleaving technique* [6, Section 7.2.2]. Here the odd and even bits of each lane are split, and the state of KECCAK- f [1600] is represented as 50 words of 32 bits. Rotations are then performed as cyclic shifts on 32-bit words, making them efficient on a 32-bit processor. There is a cost associated to the conversion of the input message into this representation, but this cost remains small compared to the evaluation of the permutation itself (see [6, Section 7.2.2] for the performance penalty). Note that the use of, for example, modular addition would have prevented the bit interleaving technique.

Some families of hash functions make use of two distinct compression functions, one oriented to 32-bit words and one to 64-bit words, in order to provide different output lengths and/or security levels. A full implementation on a given platform of such a family includes two separate compression functions, and hence at least one of the two will have a word length different from that of the CPU. In contrast, all KECCAK members we propose for standardization can be implemented with a single permutation KECCAK- f [1600] that thanks to bit interleaving can work with either 25 words on a 64-bit CPU or 50 words on a 32-bit CPU.

In terms of memory footprint, KECCAK- f [1600] requires 200 bytes of RAM for the state and some working memory [6, Section 7.2]. The sponge construction allows implementations to XOR the message block into the state directly, relieving the application from dedicating a memory area for it. This optimization applies where the hashing API is composed of functions such as `Init`, `Update` and `Final`. In general a message queue must be allocated, which can be avoided for sponge functions or similar.

The choice of width 1600 allows for a high bitrate even for high capacity values. For instance, KECCAK can process 800 more input bits per evaluation of KECCAK- f [1600] than of KECCAK- f [800] when c is fixed. However, the designer of an application on a memory-constrained device may opt for a smaller state size by using an alternate set of parameters. KECCAK[$r = 288, c = 512$] for instance uses 100 bytes of RAM. And if 256 bits of capacity are enough for such an application, KECCAK[$r = 144, c = 256$] uses only 50 bytes. Similar ideas apply to hardware implementations, where KECCAK- f [800] and KECCAK- f [400] can be seen as compact alternatives. Using a smaller width has a price, though, as it requires to support another KECCAK- f permutation. This may be acceptable if such an application is exceptional or operates in a rather closed system, freeing the standard from supporting anything else other than KECCAK- f [1600].

5 Tree hashing

A way to exploit parallelism is to use tree hashing [7]. This technique can exploit SIMD architectures, multiple cores, or both. Like most hash functions, KECCAK can benefit from this technique. We do not propose tree hashing in the specifications because a sound and well-defined tree hashing construction can work above the mode of operation and so using an unmodified instance of KECCAK. The drawbacks of this technique, though, are the larger memory footprint and the extra fixed processing cost, which can be significant for smaller messages.

In the light of two recent papers [10, 7], a sound tree hashing mode can be easily built as an application on top of existing hash functions and does not have to be embedded in the mode of operation. We define in [6, Section 6.4] an example of such a tree hashing application called KECCAKTREE.

Tree hashing can not only benefit from multiple cores, they can also exploit SIMD architectures on a single core. For instance, a specific instance of KECCAKTREE can reach about 9 cycles/byte (single core) on NIST's reference platform using SSE2 instructions [6, Section 7.3.3]. Further improvements may be obtained, in the future, with larger SIMD registers, and of course by moving to a multiple core implementation.

This technique is not useful for short messages, however, as there is a fixed additional cost corresponding to the processing of a couple of extra blocks (the number depending on the chosen parameters). Also, the memory footprint increases with the number of KECCAK- f permutations that can be evaluated in parallel.

On platforms with less parallelism, KECCAKTREE can only partially exploit the parallelism available in the chosen tree structure or can even be implemented sequentially (and is thus not significantly slower than KECCAK itself for long messages). Except for the memory footprint and for short messages, it can be advantageous to use a tree enabling a high level of parallelism and let the target platform organize the computation to take advantage of this parallelism or less.

Finally, it is worth noting that the arbitrarily-long output length of KECCAK comes in handy for tree hashing. Referring to [6, Section 6.4] and [7] for the technical explanations, the intermediate hashing nodes need to produce at least c bits of output, while the four fixed-output-length variants output only $n = c/2$ bits. This is another reason for proposing an arbitrary output length.

6 On the safety margin

In this section, we explain how the safety margin in KECCAK can be increased or decreased simply by changing the number of rounds in KECCAK- f and explain why we think the nominal number of rounds provide a high safety margin. Finally, we describe two techniques to build a *safe mode* into KECCAK implementations at little additional cost, which one could migrate to in the hypothetical case that a weakness in KECCAK is found.

6.1 Changing the number of rounds

The number of rounds of the KECCAK- f permutations is defined and fixed in [5] and reflects the trade-off between performance and safety margin made in the design. Nevertheless, the specifications make it easy to define KECCAK with an increased or decreased number of rounds. With the exception of the addition of a round constant, the rounds are identical. As the round constants are defined for any number of rounds, it is sufficient to modify the total number of rounds in the specifications.

So, someone who would like to use KECCAK but does not feel comfortable with its safety margin can simply adopt a version with more rounds. Someone who feels that KECCAK has an excessive safety margin can adopt a version with fewer rounds.

6.2 The safety margin with the nominal number of rounds

As reflected in our estimates for the safety margin of KECCAK against different types of attack in [6, Section 5.4], we think KECCAK- f has about twice as many rounds as strictly required for KECCAK to stand up to its security claim, for any choice of the capacity. The high number of rounds in KECCAK- f [1600] is due to our adoption of the hermetic sponge strategy [4] and our wish to keep a safety margin against all distinguishers, irrespective of their strength and applicability to KECCAK itself, e.g., see [9].

In September 2009 we have increased the number of rounds from 18 to 24 in KECCAK- f [1600]. We took this decision after the publication of a valid but non-threatening 16-round structural distinguisher in [1]. We refer to [9] for a treatment of this.

6.3 Migration path in the presence of a deployed standard

We expect a hash standard to be ubiquitous both in software and dedicated hardware implementations. If a weakness is discovered that has a real-world security impact, it is beneficial to have an affordable migration path towards a version without this weakness. On the NIST SHA-3 mailing list Ron Rivest [13, 2-Aug-2009] and other researchers proposed having a security parameter (e.g., the number of rounds) to be determined by the user. Disadvantages of this approach were discussed and the most important ones are the increased implementation cost due to the additional parameter, the burden of having to choose the security parameter value by the hash function user and the risk of denial-of-service attacks. Moreover, the support of a smooth choice for the security parameter may actually introduce new weaknesses, as observed by Stefan Lucks in his message to the NIST SHA-3 mailing list [13, 3-Aug-2009].

In the most lightweight version of this approach the security parameter would have only two values: one nominal value and one high-security value (e.g., tripling the number of rounds). In case of emergency, it would then be possible to migrate to the high-security value. We describe here two methods for migrating to a more secure version that applies to KECCAK without impact on the hash function implementation itself.

Both methods we propose consist of an input pre-processing step. In all use cases the input to a sponge function is a bitstring, typically made of message bits and possible key bits. After padding, the input consists of a sequence of r -bit blocks. Before presenting it to the sponge construction, this input can then be expanded by inserting bytes with fixed values in certain places. Depending on where these bytes are inserted, this has an effect similar to reducing the rate of the sponge function or multiplying the number of rounds of the underlying permutation.

The first option is to reduce the effective bitrate from r bits to $r - \delta$ bits by inserting after every input block of $r - \delta$ bits a block of δ bits equal to zero. This reduces the number of bits an attacker can exploit from r to $r - \delta$. Note that with this approach the hermetic sponge strategy is abandoned as the effective capacity is increased while the claimed capacity stays fixed.

The second option is to multiply the effective number of rounds of the underlying permutation by a factor α by inserting after every input block of r bits $\alpha - 1$ blocks of r bits with fixed and well-defined values. The α applications of the underlying permutation interleaved with the application of the fixed blocks, can then be seen as a single permutation with α as many rounds as the original one and with the

fixed-value blocks as round constants. As it is generally expected that increasing the number of rounds increases the safety margin with respect to almost all attacks, this provides a migration path to a security fix in case of a hypothetical security weakness. In this case the hermetic sponge strategy can be maintained as the single permutation with α as many rounds is assumed to have no structural distinguishers.

Both methods have the advantage of leaving KECCAK- f untouched, which limits the cost of migrating should the need occur.

The KECCAK Team, February 2010

Guido Bertoni, Joan Daemen, Michaël Peeters and Gilles Van Assche

References

- [1] J.-P. Aumasson and W. Meier, *Zero-sum distinguishers for reduced Keccak-f and for the core functions of Luffa and Hamsi*, Available online, 2009, <http://131002.net/data/papers/AM09.pdf>.
- [2] G. Bertoni, J. Daemen, M. Peeters, and G. Van Assche, *Sponge functions*, ECRYPT Hash Workshop 2007, May 2007, also available as public comment to NIST from http://www.csrc.nist.gov/pki/HashWorkshop/Public_Comments/2007_May.html.
- [3] ———, *On the indistinguishability of the sponge construction*, Advances in Cryptology – Eurocrypt 2008 (N. P. Smart, ed.), Lecture Notes in Computer Science, vol. 4965, Springer, 2008, <http://sponge.noekeon.org/>, pp. 181–197.
- [4] ———, *Cryptographic sponges*, 2009, <http://sponge.noekeon.org/>.
- [5] ———, *KECCAK specifications, version 2*, NIST SHA-3 Submission, September 2009, <http://keccak.noekeon.org/>.
- [6] ———, *KECCAK sponge function family main document*, NIST SHA-3 Submission (updated), September 2009, <http://keccak.noekeon.org/>.
- [7] ———, *Sufficient conditions for sound tree and sequential hashing modes*, Cryptology ePrint Archive, Report 2009/210, 2009, <http://eprint.iacr.org/>.
- [8] ———, *Tune KECCAK to your requirements*, 2009, <http://keccak.noekeon.org/tune.html>.
- [9] ———, *Note on zero-sum distinguishers of KECCAK-f*, Comment on the NIST Hash Competition, January 2010, <http://keccak.noekeon.org/NoteZeroSum.pdf>.
- [10] Y. Dodis, L. Reyzin, R. Rivest, and E. Shen, *Indistinguishability of permutation-based compression functions and tree-based modes of operation, with applications to MD6*, Fast Software Encryption (O. Dunkelman, ed.), Lecture Notes in Computer Science, vol. 5665, Springer, 2009, pp. 104–121.
- [11] IEEE, *P1363-2000, standard specifications for public key cryptography*, 2000.
- [12] RSA Laboratories, *PKCS # 1 v2.1 RSA Cryptography Standard*, 2002.
- [13] NIST, *Mailing list on NIST's cryptographic hash workshops and hash algorithm competition*, http://csrc.nist.gov/groups/ST/hash/email_list.html.
- [14] ———, *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>.