KANGAROOTWELVE: Fast Hashing Based on KECCAK- p

Guido Bertoni³, Joan Daemen^{1,2}, Michaël Peeters¹, Gilles Van Assche^{1(✉)},
Ronny Van Keer¹, and Benoît Viguier²

¹ STMicroelectronics, Diegem, Belgium
`gilles.vanassche@st.com`

² Radboud University, Nijmegen, The Netherlands

³ Security Pattern, Brescia, Italy

Abstract. We present KANGAROOTWELVE, a fast and secure arbitrary output-length hash function aiming at a higher speed than the FIPS 202's SHA-3 and SHAKE functions. While sharing many features with SHAKE128, like the cryptographic primitive, the sponge construction, the eXtensible Output Function (XOF) and the 128-bit security strength, KANGAROOTWELVE offers two major improvements over its standard counterpart. First it has a built-in parallel mode that efficiently exploits multi-core or SIMD instruction parallelism for long messages, without impacting the performance for short messages. Second, relying on the cryptanalysis results on KECCAK over the past ten years, we tuned its permutation to require twice less computation effort while still offering a comfortable safety margin. By combining these two changes KANGAROOTWELVE consumes less than 0.55 cycles/byte for long messages on the latest Intel®'s SkylakeX architectures. The generic security of KANGAROOTWELVE is guaranteed by the use of SAKURA encoding for the tree hashing and of the sponge construction for the compression function.

Keywords: Symmetric cryptography · Hash function · Tree hashing
KECCAK · Software performance

1 Introduction

Most cryptography involves careful trade-offs between performance and security. The performance of a cryptographic function can be objectively measured, although it can yield a wide spectrum of figures depending on the variety of hardware and software platforms that the users may be interested in. Out of these, performance on widespread processors is easily measurable and naturally becomes the most visible feature. Security on the other hand cannot be measured. The best one can do is to obtain security assurance by relying on public scrutiny by skilled cryptanalysts. This is a scarce resource and the gaining of

insight requires time and reflection. With the growing emphasis on provable security reduction of modes, the fact that the security of the underlying primitives is still based on public scrutiny should not be overlooked.

In this paper we present the hash function KANGAROOTWELVE, or more exactly an *eXtendable Output Function* (XOF). KANGAROOTWELVE makes use of a tree hash mode with SAKURA encoding [9,30] and the sponge construction [7], both proven secure. Its underlying permutation is a member of the KECCAK- $p[1600, n_T]$ family, differing from that of KECCAK only in the number of rounds. Since its publication in 2008, the round function of KECCAK was never tweaked [6]. Moreover, as for most symmetric cryptographic primitives, third-party cryptanalysis has been applied to reduced-round versions of KECCAK. Hence KANGAROOTWELVE’s security assurance directly benefits from nearly ten years of public scrutiny, including all cryptanalysis during and after the SHA-3 competition [13].

KANGAROOTWELVE gets its low computational workload per bit from using the KECCAK- $f[1600]$ permutation reduced to 12 rounds. Clearly, 12 rounds provide less safety margin than the full 24 rounds in SHA-3 and SHAKE functions. Still, the safety margin provided by 12 rounds is comfortable as, e.g., the best published collision attacks at time of writing break KECCAK only up to 6 rounds [15,16,36,37].

The other design choice that gives KANGAROOTWELVE great speed for long messages is the use of a tree hash mode. This mode is transparent for the user in the sense that the message length fully determines the tree topology. Basically, the mode calls an underlying sponge-based compression function for each 8192-byte chunk of message and finally hashes the concatenation of the resulting digests. We call this the *final node growing* approach. Clearly, the chunks can be hashed in parallel.

The main advantage of the final node growing approach is that implementers can decide on the degree of parallelism their programs support. A simple implementation could compute everything serially, while another would process two, four or more branches in parallel using multiple cores, or more simply, a SIMD instruction set such as the Intel® AVX2™. Future processors can even contain an increasing number of cores, or wider SIMD registers as exemplified by the recent AVX-512™ instruction set, and KANGAROOTWELVE will be readily able to exploit them. The fixed length of the chunks and the fact that the tree topology is fully determined by the message length improves interoperability: The hash result is independent of the amount of parallelism exploited in the implementation.

KANGAROOTWELVE is not the only KECCAK-based parallel hash mode. In late 2016, NIST published the SP 800-185 standard, including a parallelized hash mode called ParallelHash [31]. Compared to ParallelHash, KANGAROOTWELVE improves on the speed for short messages. ParallelHash compresses message chunks to digests in a first stage and compresses the concatenation of the digests in a second stage. This two-stage hashing introduces an overhead that is costly for short messages. In KANGAROOTWELVE we apply a technique called *kangaroo*

hopping: It merges the hashing of the first chunk of the message and that of the chaining values of the remaining chunks [9]. As a result, the two stages reduce to one if the input fits in one chunk with no overhead whatsoever.

Finally, KANGAROOTWELVE is a concrete application of the SAKURA encoding, which yields secure tree hash modes by construction [9].

After setting up some notation conventions in Sect. 2, we specify KANGAROOTWELVE in Sect. 3. Section 4 gives a rationale and Sect. 5 introduces a closely related variant called MARSUPILAMIFOURTEEN. In Sect. 6, we discuss implementation aspects and display benchmarks for recent processors.

2 Notation

A bit is an element of $\mathbb{Z}_2$. A string of bits is denoted using single quotes, e.g., ‘0’ or ‘111’. The concatenation of two strings a and b is denoted $a||b$. The truncation of a string s to its first n bits is denoted $\lfloor s \rfloor_n$. The n times repetition of a bit ‘s’ is denoted ‘sⁿ’, e.g. ‘110⁴’ = ‘110000’. The empty string is denoted as *.

A byte is a string of 8 bits. The byte $b_0, b_1, \dots, b_7$ can also be represented by the integer value $\sum_i 2^i b_i$ written in hexadecimal. E.g., the bit string 11110010 can be equivalently written as 0x4F as depicted in Fig. 1. The function $\text{enc}_8(x)$ encodes the integer x , with $0 \leq x \leq 255$, as a byte with value x .

The length in of a byte string s is denoted $\|s\|$. $(0x00)^n$ denotes the n times repetition of the byte 0x00.

Fig. 1. Example of byte representation

3 Specifications of KANGAROOTWELVE

KANGAROOTWELVE is an eXtensible Output Function (XOF). It takes as input a message M and an optional customization string C , both byte strings of variable length.

KANGAROOTWELVE produces unrelated outputs on different couples (M, C) . The customization string C is meant to provide *domain separation*, namely, for two different customization strings $C_1 \neq C_2$, KANGAROOTWELVE gives two independent functions of M . In practice, C is typically a short string, such as a name, an address or an identifier (e.g., URI, OID). KANGAROOTWELVE naturally maps to a XOF with a single input string M by setting the customization string input C to the empty string. This allows implementing it with a classical hash function API.

As a XOF, the output of KANGAROOTWELVE is unlimited, and the user can request as many output bits as desired. It can be used for traditional hashing simply by generating outputs of the desired digest size.

We provide a reference implementation and test vectors in [11].

3.1 The Inner Compression Function F

The core of KANGAROOTWELVE is the KECCAK- $p[1600, n_r = 12]$ permutation, i.e., a version of the permutation used in SHAKE and SHA-3 instances reduced to $n_r = 12$ rounds [30]. We build a sponge function F on top of this permutation with capacity set to $c = 256$ bits and therefore with rate $r = 1600 - c = 1344$. It makes use of multi-rate padding, indicated by `pad10*1`. Following [30], this is expressed formally as:

$$F = \text{SPONGE}[\text{KECCAK-}p[1600, n_r = 12], \text{pad10*1}, r = 1344].$$

On top of the sponge function F , KANGAROOTWELVE uses a SAKURA-compatible tree hash mode, which we describe shortly.

3.2 The Merged Input String S

First, we merge M and C to a single input string S in a reversible way by concatenating:

- the input message M ;
- the customization string C ;
- the length in bytes of C encoded using `length_encode( $\|C\|$ )` as in Algorithm 1.

Algorithm 1. The function `length_encode( $x$ )`

Input: an integer x in the range $0 \leq x \leq 256^{255} - 1$

Output: a byte string

Let l be the smallest integer in the range $0 \leq l \leq 255$ such that $x < 256^l$

Let $x = \sum_{i=0}^{l-1} x_i 256^i$ with $0 \leq x_i \leq 255$ for all i

return `enc8( $x_{l-1}$ ) || ... || enc8( $x_1$ ) || enc8( $x_0$ ) || enc8( $l$ )`

Examples:

`length_encode(0)` returns `0x00`

`length_encode(12)` returns `0x0C || 0x01`

`length_encode(65538)` returns `0x01 || 0x00 || 0x02 || 0x03`

Then, the input string S is cut into chunks of $B = 8192$ bytes, i.e.,

$$S = S_0 || S_1 || \dots || S_{n-1},$$

with $n = \left\lceil \frac{\|S\|}{B} \right\rceil$ and where all chunks except the last one must have exactly B bytes. Note that there is always one block as S consists of at least one byte.

3.3 The Tree Hash Mode

When $\|S\| > B$, we have $n > 1$ and KANGAROOTWELVE builds a tree with the following final node Node_* and inner nodes Node_i with $1 \leq i \leq n-1$:

$$\text{Node}_i = S_i \parallel '110'$$

$$\text{CV}_i = \lfloor F(\text{Node}_i) \rfloor_{256}$$

$$\begin{aligned} \text{Node}_* = S_0 \parallel '110^{62}' \parallel \text{CV}_1 \parallel \dots \parallel \text{CV}_{n-1} \parallel \text{length_encode}(n-1) \\ \parallel 0\text{xFF} \parallel 0\text{xFF} \parallel '01' \end{aligned}$$

$$\text{KANGAROOTWELVE}(M, C) = F(\text{Node}_*).$$

The chaining values CV_i have length $c = 256$ bits. This is illustrated in Fig. 2.

Fig. 2. Schematic of KANGAROOTWELVE for $\|S\| > B$, with arrows denoting calls to F .

When $\|S\| \leq B$, we have $n = 1$ and the tree reduces to its single final node Node_* and KANGAROOTWELVE becomes:

$$\text{Node}_* = S \parallel '11'$$

$$\text{KANGAROOTWELVE}(M, C) = F(\text{Node}_*).$$

3.4 Security Claim

We make a flat sponge claim [8] with 255 bits of claimed capacity in Claim 1. Informally, it means that KANGAROOTWELVE shall offer the same security strength as a random oracle whenever that offers a strength below 128 bits and a strength of 128 bits in all other cases. We discuss the implications of the claim more in depth in Sect. 4.1.

Claim 1 (Flat Sponge Claim [8]). *The success probability of any attack on KANGAROOTWELVE shall not be higher than the sum of that for a random oracle and*

$$1 - e^{-\frac{N^2}{2^{256}}},$$

with N the attack complexity in calls to KECCAK- p [1600, $n_r = 12$] or its inverse. We exclude from the claim weaknesses due to the mere fact that the function can be described compactly and can be efficiently executed, e.g., the so-called random oracle implementation impossibility [24], as well as properties that cannot be modeled as a single-stage game [33].

Note that $1 - e^{-\frac{N^2}{2^{256}}} < \frac{N^2}{2^{256}}$.

4 Rationale

In this section, we provide some more in-depth explanations on the design choices in KANGAROOTWELVE.

4.1 Implications of the Security Claim

The flat sponge claim covers all attacks up to a given *security strength* of 128 bits. Informally, saying that a cryptographic function has a security strength of s bits means that no attacks exist with complexity N and success probability p such that $N/p < 2^s$ [26].

The claim covers quasi all practically relevant security of KANGAROOTWELVE including that of traditional hashing: collision, preimage and second preimage resistance. To achieve 128-bit security strength, the output n must be chosen long enough so that there are no generic attacks (i.e., also applicable to a random oracle) that violate 128-bit security. So for 128-bit (second) preimage security the output should be at least 128 bits, and for 128-bit collision security the output should be at least 256 bits.

For many primitives the security strength that can be claimed degrades under multi-target attacks by $\log_2 M$ bits with M the number of targets. This is not the case for the flat sponge claim. As an example, let us take the case of a multi-target preimage attack versus a single-target preimage attack.

- In a (single-target) preimage attack, the adversary is given a n -bit challenge y and has to find an input x such that $\lfloor f(x) \rfloor_n = y$. A random oracle offers n bits of security strength: After N attempts, the total success probability is p with $p \approx N2^{-n}$. So we have that $N/p \approx 2^n$ for $N < 2^n$ and the security strength for a random oracle is n . For KANGAROOTWELVE we claim security strength $\min(n, 128)$ bits in this case.
- In an M -target preimage attack, the adversary is given M challenges, y_1 to y_M , and she succeeds if she find an input x such that $\lfloor f(x) \rfloor_n = y_i$ for any of the challenges. A random oracle with N attempts has success probability p with $p \approx MN2^{-n}$, and hence $N/p \approx 2^n/M$. So the security strength for the random oracle reduces to $n - \log M$ bits. For KANGAROOTWELVE we claim security strength $\min(n - \log M, 128)$ bits in this case.

Clearly, the reduction in security due to M targets is generic and independent of the security strength. It can be compensated for by increasing the output length n by $\log M$ bits.

4.2 Security of the Mode

The security of the mode, or the generic security, relies on both the sponge construction and on the tree hash mode. The latter is SAKURA-compatible so that it automatically satisfies the conditions of soundness and guarantees security against generic attacks, see [9, Theorem 1] and [10, Theorem 1]. In both cases, the bottleneck is the ability to generate collisions in the chaining values, or equivalently, collisions of the inner hash function.

The probability of inner collisions in the sponge construction is $N^2/2^{c+1}$, with N the number of blocks [7]. Regarding the collisions in the chaining values of the tree hash mode, the probability is at most $q^2/2^{c+1}$ [10, Theorem 1] with q the number of queries to F . Since each query to F implies at least one block to be processed by the sponge construction, we have $q \leq N$ and we can bound the sum of the two probabilities as $N^2/2^{c+1} + q^2/2^{c+1} \leq N^2/2^{(c-1)+1}$. This expression is equivalent as if c was one bit less than with a single source of collisions, and Claim 1 takes this into account by setting the claimed capacity to $c - 1 = 255$ bits.

We formalize the security of KANGAROOTWELVE's mode of operation in the following theorem. We can see the combination of the tree hash mode and the sponge construction as applied in KANGAROOTWELVE as a mode of operation of a permutation and call it $\mathcal{K}$.

Theorem 1. *The advantage of differentiating $\mathcal{K}$, where the underlying permutation is uniformly chosen among all the possible 1600-bit permutations, is upper bounded by*

$$\frac{2N^2 + N}{2^{c+1}},$$

with N the number of calls to the underlying permutation.

Proof. By the triangle inequality, the advantage in distinguishing $\mathcal{K}$ calling a random permutation from a random oracle is upper bounded by the sum of two advantages:

- that of distinguishing the tree hash mode calling as inner function a random function $\mathcal{F}$ from a random oracle;
- that of distinguishing the sponge construction calling a random permutation from a random function.

The former advantage is upper bounded by $q^2/2^{c+1}$, where q is the number of calls to $\mathcal{F}$. This follows from Theorem 1 of [10] for any sound tree hash mode, and from Theorem 1 of [9] that says that any SAKURA-compatible tree hash mode is sound. We show that the tree hash mode is indeed SAKURA-compatible in Sect. 4.3.

Following Theorem 2 of [7], the latter advantage is upper bounded by $(N^2 + N)/2^{c+1}$. Adding the two bounds and using $q \leq N$ proves our theorem. $\square$

4.3 SAKURA Compatibility

To show SAKURA-compatibility, we use the following terminology. The inputs to the underlying hash function are called *nodes*. Each node consists of one or more *hops*, and a hop is either a chunk of the message or a sequence of chaining values.

The encoding of the nodes follows [9, Sect. 3.1]:

- When $n = 1$, the tree reduces to a single node. This is the final node, and it contains a single message hop consisting of the input string S followed by the frame bits “message hop” ‘1’ and “final node” ‘1’.
- When $n > 1$, there are inner nodes and the final node.
 - Each inner node contains a message hop consisting of a chunk S_i followed by the frame bit “message hop” ‘1’; a simple padding bit ‘1’ and “inner node” ‘0’.
 - The final node contains two hops: a message hop followed by a chaining hop. The message hop is the first chunk of the input string S_0 followed by the frame bit “message hop” ‘1’ and a padding string ‘1’||‘0’⁶², to align the chaining hop to 64-bit boundaries. The chaining hop consists of the concatenation of the chaining values, the coded number of chaining values (`length_encode`($n - 1$)), the indication that there was no interleaving ($I = \infty$, coded with the bytes 0xFF||0xFF) and the frame bits “chaining hop” ‘0’ and “final node” ‘1’.

4.4 Choice of B

We fix the size of the message chunks to make KANGAROOTWELVE a function without parameters. This frees the user from the burden of this technical choice and facilitates interoperability.

In particular, we chose $B = 8192$. First, we chose a power of two as this can help fetching bulk data in time-critical applications. For instance, when hashing a large file, we expect the implementation to be faster and easier if the chunks contain a whole number of disk sectors.

As for the order of magnitude of B , we took into account following considerations. For each B -byte block there is a 32-byte chaining value in the final node, giving rise to a relative processing overhead of about $32/B$. Choosing $B = 2^{13}$, this is only $2^{-8} \approx 0.4\%$.

Another concern is the number of *unused bytes* in the last r -bit block of the input to F . We have $r = 1344$ bits or $R = r/8 = 168$ bytes. When cutting the chunk S_i into blocks of R bytes, it leaves $W = -(B + 1) \bmod R$ unused bytes in the last block. It turns out that W reaches a minimum for $B = 2^{7+6n}$ with $n \geq 0$ an integer. Its relative impact, $\frac{W}{B}$, decreases as B increases. For small values, e.g., $B \in \{128, 256, 512\}$, this is about 30%, while for $B = 8192$ it drops below 0.5%.

There is a tension between a larger B and the exploitable parallelism. Increasing B would further reduce these two overhead factors, but it would also delay the benefits of parallelism to longer messages.

Finally, the choice of B bounds the degree of parallelism that an implementation can fully exploit. An implementation can in principle compute the final node and leaves in parallel, but if more than $B/32$ leaves are processed at once, the final node grows faster than B bytes at a time. The chosen value of B allows a parallelism up to degree $B/32 = 256$.

4.5 Choice of the Number of Rounds

Opting for the $\text{KECCAK-}p[1600, n_r = 12]$ permutation is a drastic reduction in the number of rounds compared to the nominal KECCAK and to the SHA-3 standard. Still, there is ample evidence from third-party cryptanalysis that the switch to $\text{KECCAK-}p[1600, n_r = 12]$ leaves a safety margin similar to the one in the SHA-2 functions.

Currently, the best collision attack applicable to KANGAROOTWELVE or any SHA-3 instance works only when the permutation is reduced to 5 rounds [36]. The attack extends to 6 rounds if more degrees of freedom are available and requires a reduction of the capacity from 256 to 160 bits. Preimage attacks reach an even smaller number of rounds [21]. Hence our proposal has a safety margin of 7 out of 12 rounds w.r.t. collision and (second) preimage resistance.

Structural distinguishers is the term used for properties of a specific function that are very unlikely to be present in a random function. Zero-sum distinguishers were applied to the $\text{KECCAK-}p[1600, n_r]$ family of permutations in a number of publications [3, 14, 21]. They allow producing a set of input and of output values that both sum to zero, and this in about half the time it would be needed on a random permutation with only black-box access. These structural distinguishers are of nice theoretical interest, but they do not pose a threat as they do not extend to distinguishers on sponge functions that use $\text{KECCAK-}p[1600, n_r]$, see, e.g., [35].

The structural distinguisher on the KECCAK sponge function that does reach the highest number of rounds is the keystream prediction by Dinur et al. [17]. It works when the permutation is reduced to 9 rounds, with a time and data complexity of 2^{256} , and allows to predict one block of output. This is above the security claim of KANGAROOTWELVE, but the same authors propose a variant that works on 8 rounds with a time and data complexity of 2^{128} , leaving a safety margin of 4 rounds or 33% for KANGAROOTWELVE against this rather academic attack. Examples of structural distinguishers for the KECCAK sponge function with practical complexity and reaching the highest number of rounds are reported by Huang et al. and work up to 7 rounds [22].

In comparison, SHA-256 has a collision attack on 31 (out of 64) steps and its compression function on 52 steps [23, 25]. SHA-512's compression function admits collision attacks with practical complexities for more than half of its steps [18].

5 MARSUPILAMIFOURTEEN

While KANGAROOTWELVE claims a strong notion of 128-bit security strength, and we believe any security beyond it is purely of theoretical interest, some users may wish to use a XOF or hash function with higher security strength. In particular, when defining a cipher suite or protocol aiming for 256-bit security strength, all cryptographic functions shall have at least 256-bit security. Coming forward to such requests, in this section we present a variant of KANGAROOTWELVE with 511-bit claimed capacity.

Addressing a claimed capacity of 511 bits requires an increase of both the capacity in F and length of chaining values in the tree hash mode to at least 512 bits. Taking exactly $c = 512$ bits is sufficient for resisting generic attacks below the claim. As for KECCAK- p -specific attacks, the increase of the claimed capacity to 511 bits increases the available budget of attackers and hence reduces the safety margin. In many types of attack, adding a round in the primitive (permutation or block cipher) increases the attack complexity by a large factor. Or the other way round, if one wishes to keep the same safety margin, an increase of the attack complexity must be compensated by adding rounds.

We did the exercise and the result is MARSUPILAMIFOURTEEN. It has the same specifications as KANGAROOTWELVE, with the following exceptions:

- The capacity and chaining values are 64-byte long instead of 32 bytes. This reduces the sponge rate in F to 136 bytes.
- The number of rounds of KECCAK- $p[1600, n_r]$ is 14 instead of 12.
- The claimed capacity in the flat sponge claim is 511 bits instead of 255.

The computational workload per bit is roughly 45% higher than that of KANGAROOTWELVE.

Naturally, even thicker safety margins are achieved with the standard FIPS 202 instances or ParallelHash [30, 31].

6 Implementation

We implemented KANGAROOTWELVE in C and made it available in the KECCAK code package (KCP) [12]. We now review different aspects of this implementation and its performance.

6.1 Byte Representation

KANGAROOTWELVE assumes that its inputs M and C are byte strings. SAKURA encoding works at the bit level and adds padding and suffixes so that the input to the function F is a string of bits whose length is in general not a multiple of 8.

It is common practice in implementations of KECCAK to represent the last few bits of a string as a *delimited suffix* [12]. The delimited suffix is a byte that contains the last $|X| \bmod 8$ bits of a string X , with $|X|$ the length of X in bits,

followed by the delimiter bit ‘1’, and ending with the necessary number of bits ‘0’ to reach a length of 8 bits. When absorbing the last block in the sponge function F , the delimiter bit coincides with the first bit ‘1’ of the pad10*1 padding rule. An implementation can therefore process the first $\lfloor |X|/8 \rfloor$ bytes of the string S and, in the last block, simply add the delimited suffix and the second bit of the pad10*1 padding rule at the last position of the outer part of the state (i.e., at position $r - 1$, with r the rate).

Following the convention in Sect. 2, the delimited suffix of a string with last bits $(s_0, \dots, s_{n-1})$ can be represented by the value $2^n + \sum_{i=0}^{n-1} s_i 2^i$ in hexadecimal. For KANGAROOTWELVE, this concretely means that the final node with $\|S\| \leq B$ has suffix ‘11’ and delimited suffix 0x07. With $\|S\| > B$ the intermediate nodes with trailing bits ‘110’ use 0x0B (as depicted in Fig. 3), and the final node ending with ‘01’ will have 0x06 as delimited suffix.

Fig. 3. Example of delimited suffix

On a similar note, the 64-bit string ‘110⁶²’ in the final node is represented by the bytes 0x03||(0x00)⁷, still following the convention in Sect. 2.

This approach is taken by the Internet Research Task Force RFC draft describing KANGAROOTWELVE [38] and in the reference source code in [11].

6.2 Structuring the Implementation

The implementation has an interface that accepts the input message M in pieces of arbitrary sizes. This is useful if a file, larger than the memory size, must be processed. The customization string C can be given at the end.

We have integrated the KANGAROOTWELVE code in KCP as illustrated on Fig. 4. In particular, we instantiate the sponge construction on top of KECCAK- $p[1600, n_r = 12]$ to implement the function F , at least to compute the final node. The function F on the leaves is computed as much in parallel as possible, i.e., if at least $8B$ input bytes are given by the caller, it uses a function that computes 8 times KECCAK- $p[1600, n_r = 12]$ in parallel; if it is not available and if at least $4B$ bytes are given, it computes $4 \times$ KECCAK- $p[1600, n_r = 12]$ in parallel; and so on. If no parallel implementation exists for the given platform, or if not enough bytes are given by the caller, it falls back on a serial implementation like for the final node.

The KCP foresees that the serial and parallel implementations of the KECCAK- p permutation can be optimized for a given platform. In contrast, the code for the tree hash mode and the sponge construction is generic C, without optimizations for specific platforms, and it accesses the optimized permutation-level functions through an interface called SnP (for a single permutation) or $PlSnP$ (for permutations computed in parallel) [12].

Fig. 4. The structure of the code implementing KANGAROOTWELVE in the KCP.

To input large messages M , the state to maintain between two calls internally uses two queues: one for the final node and one for the current leaf. To save memory, the input bytes are absorbed directly into the state of F as they arrive. Hence, the state reduces to two times the state of F . Of course, if a message is known to be smaller than or equal to B bytes, one could further save one queue.

6.3 256-bit SIMD

Current mainstream PC processors, in the Intel[®] Haswell and Skylake families, support a 256-bit SIMD instruction set called AVX2[™]. We can exploit it to compute $4 \times \text{KECCAK-}p[1600, n_r = 12]$ efficiently, even on a single core.

On an Intel[®] Core[™] i5-6500 (Skylake), we measured that one evaluation of $\text{KECCAK-}p[1600, n_r = 12]$ takes about 450 cycles, while 2 in parallel about 730 cycles and $4 \times \text{KECCAK-}p[1600, n_r = 12]$ about 770 cycles. This does not include the time needed to add the input bytes to the state. Yet, this clearly points out that the time per byte decreases with the degree of parallelism.

Figure 5 displays the number of cycles for input messages up to 150,000 bytes. Microscopically, the computation time steps up for every additional $R = 168$ bytes, but this is not visible on the figure. The time needed to hash messages of length smaller than 168 bytes thus represents the smallest granularity and is reported in Table 1. Note that if many very short messages have to be processed, they can be batched so as to use a parallel implementation. This case is also reported in Table 1.

Macroscopically, when $\|S\| < B$, the time is a straight line with a slope of about 2.89 cycles/byte, i.e., the speed for F implemented serially. At $\|S\| = B = 8192$, there is a slight bump (a) as the tree gets a leaf, which causes an extra evaluation of $\text{KECCAK-}p[1600, n_r = 12]$. When $\|S\| = 3B = 24,576$, two leaves can be computed in parallel and the number of cycles drops. When $\|S\| = 5B = 40,960$, four leaves can be computed in parallel and we see another drop. From then on, the same pattern repeats and one can easily identify the slopes of serial, $\times 2$ and $\times 4$ parallel implementations of $\text{KECCAK-}p[1600, n_r = 12]$.

In our implementation, the final node is always processed with a serial implementation. In principle, a more advanced implementation could process the final node in parallel with the leaves. In more details, it would process the first chunk S_0 in parallel with the first few leaves, and it would buffer about B bytes of chaining values and so as to process them in parallel with leaves. However, at

Table 1. The overall speed for very short messages ($\|S\| < 168$) in cycles, very short messages when batched in cycles/message, for short messages ($\|S\| \leq 8192$) and for long ($\|S\| \gg 8192$) messages in cycles/byte. The figures assume a single core in each case.

Intel®Core™	Very short m.	Batched v.s.m.	Short m.	Long m.
Intel®Core™ i5-4570 (Haswell)	618 c	242 c/m	3.68 c/b	1.44 c/b
Intel®Core™ i5-6500 (Skylake)	486 c	205 c/m	2.89 c/b	1.22 c/b
Intel®Core™ i7-7800X (SkylakeX)	395 c	92 c/m	2.35 c/b	0.55 c/b

this point, we preferred code simplicity over speed optimization. Similarly, one could in principle remove the peaks of Fig. 5 and make it monotonous. It could be achieved by using, e.g., the fast $4 \times \text{KECCAK-}p[1600, n_r = 12]$ implementation even if there are less than $4B$ bytes available, with some dummy input bytes.

Fig. 5. The number of cycles of KANGAROOTWELVE on an Intel®Core™ i5-6500 (Skylake) as a function of the input message size.

Figure 6 shows the implementation cost in cycles per bytes. To determine the speed in cycles per byte for long messages in our implementation, we need to take into account both the time to process $4B$ input bytes in 4 leaves (or a multiple thereof) and 4 chaining values in the final node. Regarding the latter, 21 chaining values fit in exactly 4 blocks of $R = 168$ bytes. Hence, we measure

Fig. 6. The number of cycles per byte of KANGAROOTWELVE on an Intel®Core™ i5-6500 (Skylake) as a function of the input message size.

the time taken to process an extra $84B = \text{lcm}(4B, 21B)$ bytes. These results are reported in Table 1, together with measurements on short messages.

6.4 512-bit SIMD

Recently, Intel® started shipping processors with the AVX-512™ instruction set. It supports 512-bit SIMD instructions, enabling efficient implementations of $8 \times \text{KECCAK-}p[1600, n_r = 12]$. In addition to a higher degree of parallelism, some new features of AVX-512™ benefit to the implementation of KANGAROOTWELVE, of ParallelHash and of KECCAK in general.

- *Rotation instructions.* With the exception of AMD’s XOP™, earlier SIMD instruction sets did not include a rotation instruction. This means that the cyclic shifts in θ and ρ had to be implemented with a sequence of three instructions (shift left, shift right, XOR). With a rotation instruction, cyclic shifts are thus reduced from three to one instruction.
- *Three-input binary functions.* AVX-512™ offers an instruction that produces an arbitrary bitwise function of three binary inputs. In θ , computing the parity takes four XORs, which can be reduced to two applications of this new instruction. Similarly, the non-linear function χ can benefit from it to directly compute $a_x + (a_{x+1} + 1)a_{x+2}$.

Table 2. Speed comparison. All figures are in cycles per byte for long messages, unless otherwise specified.

Function	SkylakeX	Skylake	Haswell
KANGAROOTWELVE	0.55	1.22	1.44
KANGAROOTWELVE ($\leq 8\text{KiB}$)	2.35	2.89	3.68
ParallelHash128	0.96	2.31	2.73
Blake2bp	1.39	1.34	1.37
Blake2sp	1.22	1.29	1.39
SHAKE128	4.28	5.56	7.09
MD5	4.33	4.54	4.93
SHA-1	3.05	3.07	4.15
SHA-256	6.65	6.91	9.27
SHA-512	4.44	4.64	6.54
Blake2b	2.98	3.04	3.08
Blake2s	4.26	4.85	5.34
Blake-256	5.95	6.76	7.52
Blake-512	4.48	5.19	5.68
Grøstl-256	7.24	8.13	9.35
Grøstl-512	9.95	11.31	13.51
JH	13.04	15.14	15.09
Skein	4.48	5.18	5.34

- *32 registers.* Compared to AVX2TM, the new processors increase the number of registers from 16 to 32. As KECCAK- p has 25 lanes, this significantly decreases the need to move data between memory and registers.

We report in Table 1 the speed of our current implementation on a machine equipped with a processor in the Intel[®] SkylakeX family, supporting this instruction set [12].

6.5 Comparison with Other Functions

To put the speed of KANGAROOTWELVE in perspective, we compare it to typical hash functions, including the traditional standards MD5, SHA-1 and SHA-2 [28, 29, 34], the SHA-3 finalists [2, 19, 20, 39], the popular Blake2 functions [4] and some SHA-3 instances [30, 31]. For consistency, wherever possible we performed benchmarks on three machines in our possession. Moreover, we cross-checked with the publicly available eBASH results [5] and in case of discrepancy, we selected the fastest. For the traditional hash functions, the fastest implementation often came from OpenSSL [32]. For Blake2, we included some specific AVX2TM code by Samuel Neves [27]. Note that the comparison on SkylakeX

must be taken with care, as not all implementations available at the time of this benchmarking are fully optimized for the AVX-512TM instruction set.

Table 2 shows the results. We first list hash functions that explicitly exploit SIMD instructions with a built-in tree hash mode, such as ParallelHash and Blake2{b,s}p, and compare them to KANGAROOTWELVE for long messages (or when it is used for hashing multiple messages in parallel).

It is interesting to compare the other hash functions to KANGAROOTWELVE when it is restricted to serial processing (as for short messages), to see its speed gain already before the parallelism kicks in. Of course, such a restriction does not exist when hashing a large file, and in practice the comparison should also be made with KANGAROOTWELVE for long messages.

7 Conclusion

KANGAROOTWELVE can be seen as a new member of the KECCAK family. It inherits all the properties of the family such as suitability in hardware and resistance against side-channel attacks, but grew up with a strong focus on software performance and interoperability. We tuned the mode and the primitive to offer a tremendous computational speedup in many applications while keeping a comfortable security margin. The latter is confirmed by the cryptanalysis results on KECCAK accumulated over the last ten years, which are directly applicable to the new sibling. Also, all existing KECCAK implementations can be reused with minimal effort thanks to the layered approach in the design. For instance, KANGAROOTWELVE benefits immediately from the new SHA-3 hardware support recently introduced in the ARMv8.2 instruction set [1].

The speedup benefits to both low-end and high-end processors. For the low end, one immediately benefits from the reduction in the number of rounds, and care was taken not to add overhead in the case of short messages.

At the high end, we observed that KANGAROOTWELVE gets significant performance improvements in recent processors, which go beyond the mere gain due to parallelism. Part of these improvements come from the choice of low-latency Boolean operations in the primitive that superscalar architectures can implement efficiently, as demonstrated in the latest Intel[®]'s SkylakeX processors with the introduction of three-input binary functions.

On such a processor, KANGAROOTWELVE processes long messages at 0.55 cycles/byte. At this speed, it would require only one of its cores to process, in real-time, the output of 10 high-speed solid-state drives (SSD), i.e., accumulated bandwidth of more than 7 GB/s per core (assuming a clock frequency of 4 GHz). This simply illustrates that with KANGAROOTWELVE the speed of hashing is no longer a bottleneck in software applications.

Acknowledgements. Our implementation for the serial processing is based on the AVX2TM code written by Andy Polyakov for OpenSSL. We would also like to thank the anonymous reviewers for their constructive comments.

References

1. ARM corporation: ARM architecture reference manual ARMv8, for ARMv8-A architecture profile, document ARM DDI 0487C.a (ID121917). <http://www.arm.com/>
2. Aumasson, J.-P., Henzen, L., Meier, W., Phan, R. C.-W., SHA-3 proposal BLAKE. Submission to NIST (2008)
3. Aumasson, J.-P., Meier, W.: Zero-sum distinguishers for reduced Keccak-f and for the core functions of Luffa and Hamsi (2009). <http://131002.net/data/papers/AM09.pdf>
4. Aumasson, J.-P., Neves, S., Wilcox-O’Hearn, Z., Winnerlein, C.: BLAKE2: simpler, smaller, fast as MD5. In: Jacobson, M., Locasto, M., Mohassel, P., Safavi-Naini, R. (eds.) ACNS 2013. LNCS, vol. 7954, pp. 119–135. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-38980-1_8
5. Bernstein, D.J., Lange, T., (eds.) eBACS: ECRYPT benchmarking of cryptographic systems. <http://bench.cr.yp.to>
6. Bertoni, G., Daemen, J., Peeters, M., Van Assche, G.: KECCAK specifications. NIST SHA-3 Submission, October 2008
7. Bertoni, G., Daemen, J., Peeters, M., Van Assche, G.: On the indistinguishability of the sponge construction. In: Smart, N. (ed.) EUROCRYPT 2008. LNCS, vol. 4965, pp. 181–197. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-78967-3_11
8. Bertoni, G., Daemen, J., Peeters, M., Van Assche, G.: Cryptographic sponge functions, January 2011. <https://keccak.team/files/SpongeFunctions.pdf>
9. Bertoni, G., Daemen, J., Peeters, M., Van Assche, G.: Sakura: a flexible coding for tree hashing. In: Boureau, I., Owesarski, P., Vaudenay, S. (eds.) ACNS 2014. LNCS, vol. 8479, pp. 217–234. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-07536-5_14
10. Bertoni, G., Daemen, J., Peeters, M., Van Assche, G.: Sufficient conditions for sound tree and sequential hashing modes. *Int. J. Inf. Secur.* **13**, 335–353 (2014). <https://doi.org/10.1007/s10207-013-0220-y>
11. Bertoni, G., Daemen, J., Peeters, M., Van Assche, G., Van Keer, R.: KangarooTwelve: fast hashing based on KECCAK- p . Cryptology ePrint Archive, Report 2016/770 (2016). <http://eprint.iacr.org/2016/770>
12. Bertoni, G., Daemen, J., Peeters, M., Van Assche, G., Van Keer, R.: KECCAK code package, June 2016. <https://github.com/gvanas/KeccakCodePackage>
13. Bertoni, G., Daemen, J., Peeters, M., Van Assche, G., Van Keer, R.: KECCAK third-party cryptanalysis (2017). https://keccak.team/third_party.html
14. Boura, C., Canteaut, A., De Cannière, C.: Higher-order differential properties of KECCAK and Luffa. In: Joux, A. (ed.) FSE 2011. LNCS, vol. 6733, pp. 252–269. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-21702-9_15
15. Dinur, I., Dunkelman, O., Shamir, A.: Collision attacks on up to 5 rounds of SHA-3 using generalized internal differentials. In: Moriai, S. (ed.) FSE 2013. LNCS, vol. 8424, pp. 219–240. Springer, Heidelberg (2014). https://doi.org/10.1007/978-3-662-43933-3_12
16. Dinur, I., Dunkelman, O., Shamir, A.: Improved practical attacks on round-reduced Keccak. *J. Cryptol.* **27**(2), 183–209 (2014)
17. Dinur, I., Morawiecki, P., Pieprzyk, J., Srebrny, M., Straus, M.: Cube attacks and cube-attack-like cryptanalysis on the round-reduced Keccak sponge function. In: Oswald, E., Fischlin, M. (eds.) EUROCRYPT 2015. LNCS, vol. 9056, pp. 733–761. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-46800-5_28

18. Dobraunig, C., Eichlseder, M., Mendel, F.: Analysis of SHA-512/224 and SHA-512/256. In: Iwata, T., Cheon, J.H. (eds.) ASIACRYPT 2015. LNCS, vol. 9453, pp. 612–630. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-48800-3_25
19. Ferguson, N., Lucks, S., Schneier, B., Whiting, D., Bellare, M., Kohno, T., Callas, J., Walker, J.: The skein hash function family. Submission to NIST (Round 2) (2009)
20. Gauravaram, P., Knudsen, L.R., Matusiewicz, K., Mendel, F., Rechberger, C., Schlffer, M., Thomsen, S.S.: Grstl - a SHA-3 candidate. Submission to NIST (Round 3) (2011)
21. Guo, J., Liu, M., Song, L.: Linear structures: applications to cryptanalysis of round-reduced KECCAK. In: Cheon, J.H., Takagi, T. (eds.) ASIACRYPT 2016. LNCS, vol. 10031, pp. 249–274. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-53887-6_9
22. Huang, S., Wang, X., Xu, G., Wang, M., Zhao, J.: Conditional cube attack on reduced-round Keccak sponge function. In: Coron, J.-S., Nielsen, J.B. (eds.) EUROCRYPT 2017. LNCS, vol. 10211, pp. 259–288. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-56614-6_9
23. Li, J., Isobe, T., Shibutani, K.: Converting meet-in-the-middle preimage attack into pseudo collision attack: application to SHA-2. In: Canteaut, A. (ed.) FSE 2012. LNCS, vol. 7549, pp. 264–286. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-34047-5_16
24. Maurer, U., Renner, R., Holenstein, C.: Indifferentiability, impossibility results on reductions, and applications to the random oracle methodology. In: Naor, M. (ed.) TCC 2004. LNCS, vol. 2951, pp. 21–39. Springer, Heidelberg (2004). https://doi.org/10.1007/978-3-540-24638-1_2
25. Mendel, F., Nad, T., Schlffer, M.: Improving local collisions: new attacks on reduced SHA-256. In: Johansson, T., Nguyen, P.Q. (eds.) EUROCRYPT 2013. LNCS, vol. 7881, pp. 262–278. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-38348-9_16
26. Micciancio, D., Walter, M.: On the bit security of cryptographic primitives. Eurocrypt (2018, to appear)
27. Neves, S.: BLAKE2 AVX2 implementations. <https://github.com/sneves/blake2-avx2>
28. NIST: Federal information processing standard 180–1, secure hash standard, April 1995
29. NIST: Federal information processing standard 180–2, secure hash standard, August 2002
30. NIST: Federal information processing standard 202, SHA-3 standard: Permutation-based hash and extendable-output functions, August 2015. <http://dx.doi.org/10.6028/NIST.FIPS.202>
31. NIST: NIST special publication 800–185, SHA-3 derived functions: cSHAKE, KMAC, TupleHash and ParallelHash, December 2016. <https://doi.org/10.6028/NIST.SP.800-185>
32. OpenSSL community: OpenSSL - cryptography and SSL/TLS toolkit. <https://github.com/openssl/openssl>
33. Ristenpart, T., Shacham, H., Shrimpton, T.: Careful with composition: limitations of the indifferentiability framework. In: Paterson, K.G. (ed.) EUROCRYPT 2011. LNCS, vol. 6632, pp. 487–506. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-20465-4_27

34. Rivest, R.: The MD5 message-digest algorithm. Internet Request for Comments, RFC 1321, April 1992
35. Saha, D., Kuila, S., Chowdhury, D.R.: Symsum: symmetric-sum distinguishers against round reduced SHA3. *IACR Trans. Symmetric Cryptol.* **2017**(1), 240–258 (2017)
36. Song, L., Liao, G., Guo, J.: Non-full sbox linearization: applications to collision attacks on round-reduced KECCAK. In: Katz, J., Shacham, H. (eds.) *CRYPTO 2017*. LNCS, vol. 10402, pp. 428–451. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-63715-0_15
37. Song, L., Liao, G., Guo, J.: Solution to the 6-round collision challenge (2017). https://keccak.team/crunchy_contest.html
38. Viguier, B.: KangarooTwelve. Internet Research Task Force draft, March 2018. <https://datatracker.ietf.org/doc/draft-viguier-kangarootwelve/>
39. Wu, H.: The hash function JH. Submission to NIST (Round 3) (2011)