NORX

v3.0

Designers & Submitters:

Jean-Philippe Aumasson

Philipp Jovanovic

Samuel Neves

https://norx.io

https://github.com/norx

contact@norx.io

September 15, 2016

Contents

1 Changelog

Changes from v2.0 to v3.0:

- Increased nonce size from $2w$ to $4w$. Thus, NORX64 (NORX32) has now a nonce size of 256 (128) bit.
- Adapted datagram layouts to handle larger nonces, see Tables B.1 to B.4.
- The key is additionally XORed to the capacity at the following places:
  - In initialisation after the state is transformed with F^l (see Fig. 3.6, initialise, line 6).
  - In finalisation between the two F^l permutations and after the last one (see Fig. 3.6, finalise, lines 3 and 5).
- The tag is extracted from the capacity instead of the rate part of the state (see Fig. 3.6, finalise, line 6).
- In the parallel versions the lane counter is reduced to a single w-bit word. Moreover, the counter is now XORed to every word of the rate $\overline{s}_{i,0}, \ldots, \overline{s}_{i,11}$ instead of $\overline{s}_{i,13}$ and $\overline{s}_{i,14}$ (see Fig. 3.6, branch, line 11).

Changes from v1.1 to v2.0:

- Complete re-write of the spec aiming at more clarity and consistency.
- Renaming of variables:

type	old	new	type	old	new
word size	W	w	header	H	A
round number	R	l	payload	P	M
parallelism degree	D	p	trailer	T	Z
tag size	$\lvert A \rvert$	t	tag	A	T

- New derivation scheme for initialisation constants, see §3.5.2.
- New arrangement of the elements in the initial state, see initialise line 3 in Fig. 3.6.
- Simplified integration of the parameters w, l, p, and t during initialisation, see initialise lines 4–7 in Fig. 3.6.
- Increasing rate by $2w$ and decreasing capacity by the same amount. New rate+capacity: NORX64: $768 + 256$, NORX32: $384 + 128$.

Changes from v1.0 to v1.1:

- Branching: Added a missing -1 in $0 \leq i \leq \lceil |P|/r \rceil - 1$ for the case $p = 0$.

- Branching: Added a note that the value $\lfloor i/2^w \rfloor$, which is XORed to $s_{i,14}$, is only non-zero for very large messages.
- Payload Processing: In the parallel processing modes $p = 0$ and $p > 1$ full plaintext blocks P_i are added now directly to lane L_i for processing without padding. Only the last plaintext block P_{n-1} is padded.
- Chapter 4: Added security bounds for the NORX mode of operations from [39].
- §5.1: Added a remark concerning extensibility of the design.
- §5.3: Added software performance measurements for the Apple A7 chip and visualisations for all platforms.

2 Introduction

The *Competition for Authenticated Encryption: Security, Applicability, and Robustness (CAESAR)* [3] invites cryptographers to submit authenticated encryption schemes supporting associated data (AEAD) [48], that offer advantages over AES-GCM [32, 45] and are suitable for widespread adoption.

NORX[1] is our candidate for CAESAR. It is a novel authenticated encryption scheme with associated data supporting an arbitrary parallelism degree, based on ARX primitives yet not using modular additions. NORX has a unique parallel architecture based on the monkeyDuplex construction [20, 23], where the parallelism degree and tag size can be tuned arbitrarily. An original domain separation scheme allows simple processing of header/payload/trailer data. NORX was optimized for efficiency in both software and hardware, with a SIMD-friendly core, almost byte-aligned rotations, no secret-dependent memory lookups, and only bitwise operations. The NORX core is inspired by the ARX primitive ChaCha [17], however it replaces integer addition with the approximation $a \oplus b \oplus (a \wedge b) \ll 1$[2]. This simplifies cryptanalysis and improves hardware efficiency. Furthermore, NORX specifies a dedicated datagram to facilitate interoperability and avoid users the trouble of defining custom encoding and signalling.

Notation. *Hexadecimal numbers* are denoted in `typewriter` style, for example $\texttt{ab} = 171$. A *word* is either a 32-bit or 64-bit string, which depends on the context. Unless stated otherwise we always use little-endian representation for integers, for example when converting data streams into word arrays. Table 2.1 summarises basic notation used throughout the document.

Table 2.1: Notation used throughout the document

Symbol	Meaning
ε	The empty bitstring of length 0.
0^n	The all-zero bitstring of length n.
$\lvert x \rvert$	Length of bitstring x in bits.
$\lvert x \rvert_n$	Length of bitstring x in n-bit blocks.
$x \parallel y$	Concatenation of bitstrings x and y.
$\mathsf{hw}(x)$	Hamming weight of bitstring x.
$\neg, \wedge, \vee, \oplus$	Bitwise negation, `AND`, `OR` and `XOR`.
$x \ll n, x \gg n$	Left-/Right-shift of bitstring x by n bits.
$x \lll n, x \ggg n$	Left-/Right-rotation of bitstring x by n bits.
$\longleftarrow$	Variable assignment.
$\mathsf{left}_l(x)$	Truncation of bitstring x to its l left-most bits.
$\mathsf{right}_r(x)$	Truncation of bitstring x to its r right-most bits.

[1] The name stems from "NO(T A)RX" and is pronounced like "norcks".
[2] Derived from the well-known identity $a + b = (a \oplus b) + (a \wedge b) \ll 1$ [14, 43].

Outline. Chapter 3 gives a complete specification of the NORX family of AEAD schemes. Chapter 4 lists the security goals for confidentiality and integrity of the plaintext and for integrity of associated data and public message numbers. Chapter 5 presents features of NORX, justifies our parameter choices, and reports on performance measurements of software implementations on 32- and 64-bit processors and presents preliminary results for an hardware evaluation of an ASIC implementation. Chapter 6 motivates design decisions and Chapter 7 presents preliminary results from the cryptanalysis of various aspects of NORX. Finally, we conclude with notes on the intellectual property, a consent of the CAESAR competition, acknowledgements, references and appendices.

3 Specification

This section gives a complete specification of NORX and its proposed instances.

3.1 Parameters and Interface

A NORX instance is parameterised by

- a *word size* of $w \in \{32, 64\}$ bits,
- a *round number* $1 \leq l \leq 63$,
- a *parallelism degree* $0 \leq p \leq 255$,
- a *tag size* of $t \leq 4w$ bits.

Encryption Mode

NORX encryption takes as input

- a *key* K of $k = 4w$ bits,
- a *nonce* N of $n = 4w$ bits,
- a *tuple* (A, M, Z) where
  - A is a *header*,
  - M is a *message*,
  - Z is a *trailer/footer*,

  and where any of A, M, Z can be the empty string (that is, of length 0).

NORX encryption produces as output

- a *ciphertext* (or *encrypted payload*) C of the same size as M,
- an *authentication tag* T of t bits.

In summary, NORX encryption $\mathcal{E}$ is specified as

$$\mathcal{E} : \{0,1\}^k \times \{0,1\}^n \times \{0,1\}^* \times \{0,1\}^* \times \{0,1\}^* \to \{0,1\}^* \times \{0,1\}^t$$

with

$$\mathcal{E}(K, N, A, M, Z) = (C, T)$$

where $|M| = |C|$.

Decryption Mode

NORX decryption takes as input

- a key K of $k = 4w$ bits,
- a nonce N of $n = 4w$ bits,
- a tuple (A, C, Z) where,
  - A is a header,
  - C is a ciphertext,
  - Z is a trailer,

  and where any of A, M, Z can be the empty string (that is, of length 0).
- an authentication tag T of t bits.

NORX decryption either returns a failure $\perp$, upon failed verification of the tag, or produces a plaintext M of the same size as C if the tag verification succeeds.

In summary, NORX decryption $\mathcal{D}$ is specified as

$$\mathcal{D} : \{0,1\}^k \times \{0,1\}^n \times \{0,1\}^* \times \{0,1\}^* \times \{0,1\}^* \times \{0,1\}^t \to \{0,1\}^* \cup \{\perp\}$$

with

$$\mathcal{D}(K, N, A, C, Z, T) = \begin{cases} M & \text{if } T = T' \\ \perp & \text{if } T \neq T' \end{cases}$$

where T denotes the received authentication tag, T' the one computed on the recipient's side and $|M| = |C|$.

3.2 Instances

A NORX instance is a choice of values for the four parameters w, l, p, and t. Table 3.1 proposes five NORX instances for different use cases: 128- or 256-bit security, four or six rounds, and a version with four-wise parallelism[1]. Table 3.1 also shows the corresponding nonce and key sizes n and k and the priority order of the recommended parameter sets from highest at the top to lowest at the bottom. The last column of Table 3.1 identifies the prioritized list of targeted use cases (most important on the left) as specified in `https://groups.google.com/forum/#!topic/crypto-competitions/DLv193SPSDc`.

We set the *default tag size* t for a given word size w to $t = 4w$, i.e. for $w = 32$ we get $t = 128$ and for $w = 64$ we get $t = 256$. A detailed discussion on the parameter combinations can be found in §5.2.

A NORX instance is denoted by NORXw-l-p-t, where w, l, p, and t are the parameters of the instance, see §3.1. If the default tag size is used, i.e. if $t = 4w$, the notation for an instance is shortened to NORXw-l-p. So for example, NORX64-6-1 has $(w, l, p, t) = (64, 6, 1, 256)$.

[1] For low-end systems we refer to NORX8 and NORX16 [12] which target security levels of 80- and 96-bit.

Table 3.1: NORX instances

Nr.	w	l	p	t	k	n	Use Cases
1.	64	4	1	256	256	256	2, 1
2.	32	4	1	128	128	128	1, 2
3.	64	6	1	256	256	256	2, 1
4.	32	6	1	128	128	128	1, 2
5.	64	4	4	256	256	256	2

3.3 Layout Overview

NORX uses the monkeyDuplex construction [20, 23] as a basis. Certain configurations of the mode allow to process the payload in parallel.

For the parallel mode the number i of parallel encryption lanes L_i is controlled by the parameter $0 \leq p \leq 255$. For the value $p = 1$, the layout of NORX corresponds to a standard (sequential) duplex construction, see Fig. 3.1. For $p > 1$, the number of lanes L_i is bounded by the latter value, e.g. for $p = 2$ see Fig. 3.2. If $p = 0$, the number of lanes L_i is bounded by the size of the payload. In that case, the layout of NORX is similar to that of the PPAE construction [27].

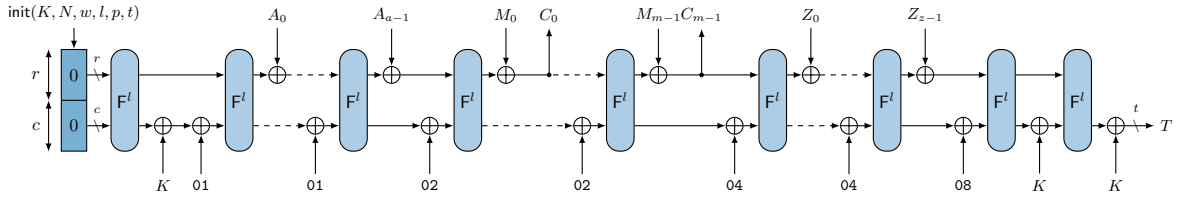

Figure 3.1: Layout of standard NORX ($p = 1$)

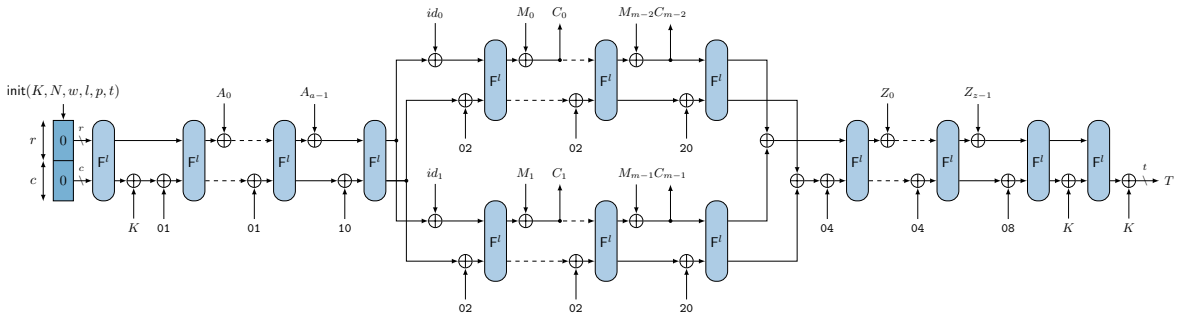

Figure 3.2: Layout of NORX with parallel encryption ($p = 2$)

The core algorithm F of NORX is a permutation of $b = r + c$ bits, where b is called the *width*, r the *rate* (or block length), and c the *capacity*. We call F a *round* and F^l denotes its l-fold iteration. The organisation of the internal state S of NORX is as follows:

w	b	r	c
32	512	384	128
64	1024	768	256

The state is viewed as a concatenation of 16 words, i.e. $S = s_0 \parallel \cdots \parallel s_{15}$, where $s_0, \ldots, s_{11}$ are called the *rate words* (where data blocks are injected) $s_{12}, \ldots, s_{15}$ are called the *capacity words* (which remain untouched). Conceptually, the 16 state words are arranged in a 4×4 matrix:

$$S = \left(\begin{array}{cccc} s_0 & s_1 & s_2 & s_3 \\ s_4 & s_5 & s_6 & s_7 \\ s_8 & s_9 & s_{10} & s_{11} \\ \hline s_{12} & s_{13} & s_{14} & s_{15} \end{array}\right)$$

3.4 The Permutation F^l

The complete pseudocode for the NORX core permutation F^l is given in Fig. 3.4. A single NORX round F processes the state S by first transforming its columns with

$$\mathsf{G}(s_0, s_4, s_8, s_{12}) \quad \mathsf{G}(s_1, s_5, s_9, s_{13}) \quad \mathsf{G}(s_2, s_6, s_{10}, s_{14}) \quad \mathsf{G}(s_3, s_7, s_{11}, s_{15})$$

and then transforming its diagonals with

$$\mathsf{G}(s_0, s_5, s_{10}, s_{15}) \quad \mathsf{G}(s_1, s_6, s_{11}, s_{12}) \quad \mathsf{G}(s_2, s_7, s_8, s_{13}) \quad \mathsf{G}(s_3, s_4, s_9, s_{14})$$

Those two operations are called *column step* and *diagonal step*, as in BLAKE2 [13], and will be denoted by col and diag, respectively. An illustration of these operations is shown in Fig. 3.3.

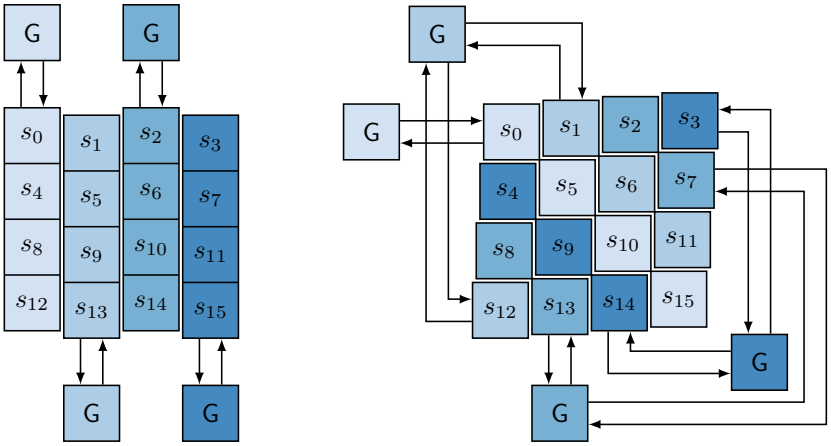

Figure 3.3: Column step and diagonal step of F

The G function uses *cyclic rotations* $\ggg$ and a *non-linear operation* H interchangeably to update its four input words a, b, c and d. The rotation offsets r_0, r_1, r_2, and r_3 for the cyclic rotations of 32- and 64-bit NORX are specified in Table 3.2.

Table 3.2: Rotation offsets for 32- and 64-bit NORX

w	r_0	r_1	r_2	r_3
32	8	11	16	31
64	8	19	40	63

Algorithm: $\mathsf{F}^l(S)$

1. **for** $i \in \{0, \ldots, l-1\}$ **do**
2. $\quad S \leftarrow \mathsf{diag}(\mathsf{col}(S))$
3. **end**
4. **return** S

Algorithm: $\mathsf{G}(a, b, c, d)$

1. $a \leftarrow \mathsf{H}(a, b)$
2. $d \leftarrow (a \oplus d) \ggg r_0$
3. $c \leftarrow \mathsf{H}(c, d)$
4. $b \leftarrow (b \oplus c) \ggg r_1$
5. $a \leftarrow \mathsf{H}(a, b)$
6. $d \leftarrow (a \oplus d) \ggg r_2$
7. $c \leftarrow \mathsf{H}(c, d)$
8. $b \leftarrow (b \oplus c) \ggg r_3$
9. **return** a, b, c, d

Algorithm: $\mathsf{col}(S)$

1. $(s_0, s_4, s_8, s_{12}) \leftarrow \mathsf{G}(s_0, s_4, s_8, s_{12})$
2. $(s_1, s_5, s_9, s_{13}) \leftarrow \mathsf{G}(s_1, s_5, s_9, s_{13})$
3. $(s_2, s_6, s_{10}, s_{14}) \leftarrow \mathsf{G}(s_2, s_6, s_{10}, s_{14})$
4. $(s_3, s_7, s_{11}, s_{15}) \leftarrow \mathsf{G}(s_3, s_7, s_{11}, s_{15})$
5. **return** S

Algorithm: $\mathsf{diag}(S)$

1. $(s_0, s_5, s_{10}, s_{15}) \leftarrow \mathsf{G}(s_0, s_5, s_{10}, s_{15})$
2. $(s_1, s_6, s_{11}, s_{12}) \leftarrow \mathsf{G}(s_1, s_6, s_{11}, s_{12})$
3. $(s_2, s_7, s_8, s_{13}) \leftarrow \mathsf{G}(s_2, s_7, s_8, s_{13})$
4. $(s_3, s_4, s_9, s_{14}) \leftarrow \mathsf{G}(s_3, s_4, s_9, s_{14})$
5. **return** S

Algorithm: $\mathsf{H}(x, y)$

1. **return** $(x \oplus y) \oplus ((x \wedge y) \ll 1)$

Figure 3.4: The NORX permutation F^l

3.5 The NORX Mode

The NORX mode is divided into a high-level and a low-level interface discussed in §§3.5.1 and 3.5.2, respectively. The *high-level interface* consists of only two functions: AEADEnc and AEADDec. These provide functionality for encryption and authentication of a message on the one hand and decryption and verification of an encrypted payload on the other. Both functions support processing of associated data. The *low-level interface* defines the concrete implementation of padding, domain separation, absorption or encryption of data block sequences, tag generation, etc.

3.5.1 High-level Structure

The two high-level interface functions AEADEnc and AEADDec are depicted in Fig. 3.5

3.5.2 Low-level Structure

The low-level functions of NORX are depicted in Fig. 3.6. Before going into the details of those methods, we first introduce the mechanisms for padding and domain separation which are required later on.

Algorithm: AEADEnc(K, N, A, M, Z)

1. $S \leftarrow \mathsf{initialise}(K, N)$
2. $S \leftarrow \mathsf{absorb}(S, A, 01)$
3. $\overline{S} \leftarrow \mathsf{branch}(S, |M|, 10)$
4. $\overline{S}, C \leftarrow \mathsf{encrypt}(\overline{S}, M, 02)$
5. $S \leftarrow \mathsf{merge}(\overline{S}, |M|, 20)$
6. $S \leftarrow \mathsf{absorb}(S, Z, 04)$
7. $S, T \leftarrow \mathsf{finalise}(S, 08)$
8. **return** C, T

Algorithm: AEADDec(K, N, A, C, Z, T)

1. $S \leftarrow \mathsf{initialise}(K, N)$
2. $S \leftarrow \mathsf{absorb}(S, A, 01)$
3. $\overline{S} \leftarrow \mathsf{branch}(S, |C|, 10)$
4. $\overline{S}, M \leftarrow \mathsf{decrypt}(\overline{S}, C, 02)$
5. $S \leftarrow \mathsf{merge}(\overline{S}, |C|, 20)$
6. $S \leftarrow \mathsf{absorb}(S, Z, 04)$
7. $S, T' \leftarrow \mathsf{finalise}(S, 08)$
8. **if** $T = T'$ **then return** M **else return** $\perp$ **end**

Figure 3.5: High-level interface functions of the standard NORX mode

Padding

NORX adopts the so-called *multi-rate padding* [23]. This padding rule is defined by the map

$$\mathsf{pad}_r : X \longmapsto X \parallel 10^u 1$$

where X is a bitstring and $u = (-|X| - 2) \bmod r$. If r and $|X|$ are divisible by 8 and X is viewed as a sequence of bytes, then the multi-rate padding can be written as

$$\mathsf{pad}_r : \begin{cases} X \longmapsto X \parallel 01 \parallel 00^u \parallel 80 & \text{if } u > 0 \\ X \longmapsto X \parallel 81 & \text{if } u = 0 \end{cases}$$

where $u = (-|X|_8 - 2) \bmod (r/8)$.

Domain Separation

NORX has a very simple and lightweight domain separation mechanism: different *domain separation constants* are XORed to the least significant byte of s_{15} before the state is transformed by F^l to distinguish different phases of the algorithm. Table 3.3 gives the specification of those constants and Figs. 3.1 and 3.2 illustrate their integration into the state of NORX. Figs. 3.5 and 3.6 show their concrete usage.

Table 3.3: Domain separation constants

header	payload	trailer	tag	branching	merging
01	02	04	08	10	20

Initialisation

The method initialise sets up the $16w$-bit internal state $S = (s_0, \ldots, s_{15})$ of NORX by processing a $4w$-bit key $K = k_0 \parallel k_1 \parallel k_2 \parallel k_3$, a $4w$-bit nonce $N = n_0 \parallel n_1 \parallel n_2 \parallel n_3$, the instance parameters w, l, p, and t and some initialisation constants. These constants are given in Table 3.4 and can be derived by

$$(u_0, \ldots, u_{15}) = \mathsf{F}^2(0, \ldots, 15)$$

Algorithm: initialise(K, N)

$k_0 \parallel k_1 \parallel k_2 \parallel k_3 \leftarrow K$, s.t. $|k_i| = w$
$n_0 \parallel n_1 \parallel n_2 \parallel n_3 \leftarrow N$, s.t. $|n_i| = w$
$S \leftarrow (n_0, n_1, n_2, n_3, k_0, k_1, k_2, k_3, u_8, u_9, u_{10}, u_{11}, u_{12}, u_{13}, u_{14}, u_{15})$
$(s_{12}, s_{13}, s_{14}, s_{15}) \leftarrow (s_{12}, s_{13}, s_{14}, s_{15}) \oplus (w, l, p, t)$
$S \leftarrow \mathsf{F}^l(S)$
$(s_{12}, s_{13}, s_{14}, s_{15}) \leftarrow (s_{12}, s_{13}, s_{14}, s_{15}) \oplus (k_0, k_1, k_2, k_3)$
return S

Algorithm: encrypt$(\overline{S}, M, v)$

$C \leftarrow \varepsilon$
$M_0 \parallel \cdots \parallel M_{m-1} \leftarrow M$, s.t. $|M_i| = r, 0 \leq |M_{m-1}| < r$
if $|M| > 0$ **then**
  for $i \in \{0, \ldots, m-2\}$ **do**
    $j \leftarrow i \bmod |\overline{S}|_b$
    $\overline{s}_{j,15} \leftarrow \overline{s}_{j,15} \oplus v$
    $\overline{S}_j \leftarrow \mathsf{F}^l(\overline{S}_j)$
    $C_i \leftarrow \mathsf{left}_r(\overline{S}_j) \oplus M_i$
    $\overline{S}_j \leftarrow C_i \parallel \mathsf{right}_c(\overline{S}_j)$
  end
  $j \leftarrow (m-1) \bmod |\overline{S}|_b$
  $\overline{s}_{j,15} \leftarrow \overline{s}_{j,15} \oplus v$
  $\overline{S}_j \leftarrow \mathsf{F}^l(\overline{S}_j)$
  $C_{m-1} \leftarrow \mathsf{left}_{|M_{m-1}|}(\overline{S}_j) \oplus M_{m-1}$
  $\overline{S}_j \leftarrow \overline{S}_j \oplus (\mathsf{pad}_r(M_{m-1}) \parallel 0^c)$
  $C \leftarrow C_0 \parallel \cdots \parallel C_{m-1}$
end
return $\overline{S}, C$

Algorithm: decrypt$(\overline{S}, C, v)$

$M \leftarrow \varepsilon$
$C_0 \parallel \cdots \parallel C_{m-1} \leftarrow C$ s.t. $|C_i| = r, 0 \leq |C_{m-1}| < r$
if $|C| > 0$ **then**
  for $i \in \{0, \ldots, m-2\}$ **do**
    $j \leftarrow i \bmod |\overline{S}|_b$
    $\overline{s}_{j,15} \leftarrow \overline{s}_{j,15} \oplus v$
    $\overline{S}_j \leftarrow \mathsf{F}^l(\overline{S}_j)$
    $M_i \leftarrow \mathsf{left}_r(\overline{S}_j) \oplus C_i$
    $\overline{S}_j \leftarrow C_i \parallel \mathsf{right}_c(\overline{S}_j)$
  end
  $j \leftarrow (m-1) \bmod |\overline{S}|_b$
  $\overline{s}_{j,15} \leftarrow \overline{s}_{j,15} \oplus v$
  $\overline{S}_j \leftarrow \mathsf{F}^l(\overline{S}_j)$
  $M_{m-1} \leftarrow \mathsf{left}_{|C_{m-1}|}(\overline{S}_j) \oplus C_{m-1}$
  $\overline{S}_j \leftarrow \overline{S}_j \oplus (\mathsf{pad}_r(M_{m-1}) \parallel 0^c)$
  $M \leftarrow M_0 \parallel \cdots \parallel M_{m-1}$
end
return $\overline{S}, M$

Algorithm: absorb(S, X, v)

$X_0 \parallel \cdots \parallel X_{m-1} \leftarrow X$, s.t. $|X_i| = r, 0 \leq |X_{m-1}| < r$
if $|X| > 0$ **then**
  for $i \in \{0, \ldots, m-2\}$ **do**
    $s_{15} \leftarrow s_{15} \oplus v$
    $S \leftarrow \mathsf{F}^l(S)$
    $S \leftarrow S \oplus (X_i \parallel 0^c)$
  end
  $s_{15} \leftarrow s_{15} \oplus v$
  $S \leftarrow \mathsf{F}^l(S)$
  $S \leftarrow S \oplus (\mathsf{pad}_r(X_{m-1}) \parallel 0^c)$
end
return S

Algorithm: branch(S, m, v)

$\overline{S} \leftarrow 0^b$
if $p \neq 1$ and $m > 0$ **then**
  $s \leftarrow p$
  if $p = 0$ **then**
    $s \leftarrow \lceil m/r \rceil$
  end
  $\overline{S} = (\overline{S}_0, \ldots, \overline{S}_{s-1}) \leftarrow (0^b, \ldots, 0^b)$
  $s_{15} \leftarrow s_{15} \oplus v$
  $S \leftarrow \mathsf{F}^l(S)$
  for $i \in \{0, \ldots, s-1\}$ **do**
    $\overline{S}_i \leftarrow S \oplus (i, i, i, i, i, i, i, i, i, i, i, i, 0, 0, 0, 0)$
  end
else
  $\overline{S} \leftarrow S$
end
return $\overline{S}$

Algorithm: merge$(\overline{S}, m, v)$

$S \leftarrow 0^b$
if $p \neq 1$ and $m > 0$ **then**
  for $i \in \{0, \ldots, |\overline{S}|_b - 1\}$ **do**
    $\overline{s}_{i,15} \leftarrow \overline{s}_{i,15} \oplus v$
    $\overline{S}_i \leftarrow \mathsf{F}^l(\overline{S}_i)$
    $S \leftarrow S \oplus \overline{S}_i$
  end
else
  $S \leftarrow \overline{S}$
end
return S

Algorithm: finalise(S, v)

$s_{15} \leftarrow s_{15} \oplus v$
$S \leftarrow \mathsf{F}^l(S)$
$(s_{12}, s_{13}, s_{14}, s_{15}) \leftarrow (s_{12}, s_{13}, s_{14}, s_{15}) \oplus (k_0, k_1, k_2, k_3)$
$S \leftarrow \mathsf{F}^l(S)$
$(s_{12}, s_{13}, s_{14}, s_{15}) \leftarrow (s_{12}, s_{13}, s_{14}, s_{15}) \oplus (k_0, k_1, k_2, k_3)$
$T \leftarrow \mathsf{right}_t(S)$
return S, T

Figure 3.6: Low-level interface functions of the NORX mode

Table 3.4: Initialisation constants

w	32	64	w	32	64
u_0	0454EDAB	E4D324772B91DF79	u_8	A3D8D930	B15E641748DE5E6B
u_1	AC6851CC	3AEC9ABAAEB02CCB	u_9	3FA8B72C	AA95E955E10F8410
u_2	B707322F	9DFBA13DB4289311	u_{10}	ED84EB49	28D1034441A9DD40
u_3	A0C7C90D	EF9EB4BF5A97F2C8	u_{11}	EDCA4787	7F31BBF964E93BF5
u_4	99AB09AC	3F466E92C1532034	u_{12}	335463EB	B5E9E22493DFFB96
u_5	A643466D	E6E986626CC405C1	u_{13}	F994220B	B980C852479FAFBD
u_6	21C22362	ACE40F3B549184E1	u_{14}	BE0BF5C9	DA24516BF55EAFD4
u_7	1230C950	D9CFD35762614477	u_{15}	D7C49104	86026AE8536F1501

which allows on-the-fly computation if necessary. Note, however, that only $u_8, \ldots, u_{15}$ are actually used in initialise.

Data Absorption

The method absorb takes an arbitrary long bitstring X as input and absorbs it in blocks of r bits into the internal state thereby ensuring authenticity of X. If the last block is smaller then r bits, it is extended to the block size through pad_r. For domain separation the constant v is used. Data absorption is skipped entirely in case the input has length 0, i.e. if X corresponds to the empty bitstring ε.

In NORX the function absorb is used for authenticating associated data in the form of header data A using domain separation constant $v = 01$ and/or trailer data Z using domain separation constant $v = 04$. Refer to the high-level interface in Fig. 3.5 to see where and how absorb is used concretely in NORX.

Branching

If the parallelism degree $p \neq 1$ then branch is used to prepare parallel payload processing. branch is skipped entirely if either $p = 1$ or $|M| = 0$. The state S is extended to a multi-state vector $\overline{S}$ having either p elements if $p > 1$ or $\lceil |M|_r \rceil$ elements if $p = 0$. Note that in order to ensure that each lane produces a unique bitstream for encryption, a w-bit *lane number* i is integrated into state copy $\overline{S}_i$ (XORed to the rate words $\overline{s}_{i,0}, \ldots, \overline{s}_{i,11}$) immediately after branching.
Due to the above lane number being a single w-bit word, the number of lanes is limited to 2^w. This implies that when $p = 0$, the maximum message size is $2^w r/8$ bytes, which is approximately 2^{36} (≈ 64 GiB) for NORX32 and 2^{70} (≈ 1024 EiB) for NORX64.

Data Encryption and Decryption

The method encrypt (decrypt) takes an arbitrary long bitstring M (C) as input and encrypts (decrypts) it thereby producing the encrypted (decrypted) payload C (M). Since M is also absorbed into the state S, its authenticity is ensured as well. As in absorb, data is processed in r-bit blocks and the last block is padded using pad_r. Note that in the latter case only a truncated data block of the same size as the unpadded input block is extracted such that $|M| = |C|$ holds. The constant $v = 02$ is used for domain separation.

The different cases for p are handled as follows. For $p = 1$ the NORX mode corresponds to a regular sequential sponge construction and no special steps have to be taken for data encryption or decryption. For $p > 1$ a fixed number of p parallel lanes is available for data processing. Data blocks are rotated in a round-robin fashion across the states by assigning the i-th data block to state $i \bmod p$. In the last case, if $p = 0$, each data block is processed on its own separate lane.

Merging

The merge function is only executed if $p \neq 1$ and $|M| > 0$. After parallel-processing all payload data blocks, the states $\overline{S}_i$ are merged back into a single state S. The domain separation constant for merge is $v = 20$.

Finalisation

The finalise function generates an authentication tag T by first injecting the domain separation constant $v = 08$ then transforming S *twice* with the permutation F^l interleaved by two key additions to the capacity, and finally extracting the t rightmost bits (=capacity) from S which are returned as the tag T.

Tag Verification

Note that tag verification is not listed explicitly among the low-level interface functions in Fig. 3.6 but rather in Fig. 3.5, see the last step of AEADDec.

Tag verification consists of comparing the *received tag* T to the *generated tag* T'. If $T = T'$, tag verification succeeds; otherwise tag verification fails, the decrypted payload is discarded and an error $\perp$ is returned.

Implementations of tag verification should satisfy the following requirements:

- Tag verification should not leak information on the (relative) values of the strings compared. In particular tag verification should be implemented in constant time, so that a comparison of identical strings take the same time as a comparison of distinct strings.
- The decrypted payload should not be returned to the user if tag verification fails. Ideally, extracted bytes should be securely erased from any temporary memory if tag verification fails.

4 Security Goals

We expect NORX with $l \geq 4$ rounds to provide the maximum security for any AEAD scheme with the same interface (input and output types and lengths). The following requirements should be satisfied in order to use NORX securely:

1. **Unique nonces.** Each key and nonce pair should not be used to process more than one message.
2. **Abort on verification failure.** If the tag verification fails, only an error is returned. In particular, the decrypted plaintext and the wrong authentication tag must not be given as an output and should be erased from memory in a safe way.

We do not make any claim regarding attackers using "related keys", "known keys", "chosen keys", etc. We also exclude from the claims below models where information is "leaked" on the internal state or key.

The security of NORX is limited by the key length (128 or 256 bits) and by the tag length (128 or 256 bits). Plaintext confidentiality should thus have the order of 128 or 256 bits of security. The same level of security should hold for integrity of the plaintext or of associated data (based on the fact that an attacker trying 2^n tags will succeed with probability 2^{n-256}, $n < 256$). In particular, recovery of a k-bit NORX key should require resources ("computations", energy, etc.) comparable to those required to recover the key of an ideal k-bit key cipher. Table 4.1 summarizes the security goals of NORX.

Table 4.1: Overview on the security levels (in bits)

security goal	NORX32	NORX64
plaintext confidentiality	128	256
plaintext integrity	128	256
associated data integrity	128	256
public message number integrity	128	256

Note that NORX restricts the number of messages processed with a given key: in [19] the *usage exponent* e is defined as the value such that the implementation imposes an upper limit of 2^e uses to a given key. In NORX we set it to $e_{64} = 128$ for 64-bit and $e_{32} = 64$ for 32-bit.

5 Features

NORX was designed for users, provides several features desirable for practical applications and offers a couple of advantages over AES-GCM [45]. First we list these characteristics in detail, then give a justification of our recommended parameter sets and finally present our performance results.

5.1 List of Characteristics

- **High security.** NORX supports 128- and 256-bit keys and authentication tags of arbitrary size, thanks to its duplex construction. The core permutation of NORX was designed and evaluated to be cryptographically strong. The minimal number of 8 rounds for initialisation / finalisation (i.e. 16 steps consisting of 8 column and 8 diagonal steps interleaved with each other) and of 4 rounds (i.e. 8 steps consisting of 4 column and 4 diagonal steps interleaved with each other) for the data processing part ensure a high security margin against cryptanalytic attacks. Large internal states of 512 and 1024 bits and the duplex construction offer protection against generic attacks.

- **Efficiency.** NORX was designed with 64-bit processors in mind, but is also compatible with smaller architectures like 8- to 32-bit platforms. Software implementations of NORX are able to take advantage of multi-core processors, due to the parallel duplex construction, and specialised instruction sets like AVX / AVX2 or NEON. Moreover, state sizes of 512 and 1024 bits make NORX very cache-friendly. Hardware implementations benefit from hardware-friendly operations, next to the arbitrary parallelism degree for payload processing, which results in highly competitive hardware performance of NORX.

- **Simplicity.** The core algorithm iterates a simple round function and can be implemented by translating our pseudocode into the programming language used: NORX requires no SBoxes, no Galois field operations, and no integer arithmetic; `AND`, `XOR`, and shifts are the only instructions required. This simplifies cryptanalysis and the task of implementing the cipher.

- **High key agility.** NORX requires no key expansion when setting up a new key, in contrast to many block-cipher based schemes, like AES-GCM. Switching the secret key is therefore very cheap. As an additional benefit, there are also no hidden costs of loading precomputed expanded keys from DRAM into L1 cache.

- **Adjustable tag sizes.** The NORX family uses a default tag size of $4w$ bits for our proposed instances. Thanks to the duplex construction, tag sizes can be easily adapted to the demands of any given application.

- **Simple integration.** NORX can be easily integrated into a protocol stack, as it supports flexible processing of arbitrary datagrams: any header and trailer are authenticated (and left in clear) and the payload is both encrypted and authenticated.

- **Interoperability.** Dedicated datagrams encode parameters of the cipher and encapsulate the protected data. This aims to increase interoperability across implementations.
- **Single pass.** Encryption and decryption of data is done in a single pass of the algorithm.
- **Online.** NORX supports encryption of data streams, i.e. the size of processed data needs not to be known in advance.
- **High data processing volume.** NORX allows to process very large data sizes from a single key-nonce pair. The usage exponent (see Chapter 4) theoretically limits the number of calls to the core permutation to values of 2^{64} (NORX32) and 2^{128} (NORX64). This translates to data sizes, which are orders of magnitude beyond everything relevant for current real-world applications. Especially, these values are a lot higher than the maximum of 2^{32} calls to the authenticated encryption function of AES-GCM, which could be easily reached already nowadays in practical applications.
- **Minimal overhead.** Payload encryption is non-expanding, i.e. the ciphertext has the same length as the plaintext. The authentication tag, has a length of 16 or 32 bytes depending on the concrete instance of NORX.
- **Robustness against timing attacks.** By avoiding data-dependent table look-ups, like SBoxes, and integer additions, the goal to harden soft- and hardware implementations of NORX against timing attacks should be comparably easy to achieve.
- **Moderate misuse resistance.** NORX retains its security on nonce reuse as long as it can be guaranteed that header data is unique[1]. For comparison, nonce reuse in AES-GCM is a major security issue, allowing an attacker to recover the secret key [38].
- **Autonomy.** NORX requires no external primitive.
- **Diversity.** The cipher does not depend on AES instructions, thereby adding to the diversity among cryptographic algorithms.
- **Extensibility.** Thanks to the duplex construction and a simple, yet powerful domain separation scheme, NORX can be easily extended to support additional features, like secret message numbers, sessions, or forward secrecy without losing its security guarantees.

5.2 Recommended Parameter Sets

We consider NORX32-4-1 and NORX64-4-1 as the standard instances for the respective word sizes of 32 and 64 bit. These configurations offer a good balance between performance and security. We recommend NORX32-4-1 for low resource applications on 8- to 32-bit platforms and NORX64-4-1 for software implementations on modern 64-bit CPUs or standard hardware implementations. Applications that require a higher security margin and where performance has less priority are advised to use the instances NORX32-6-1 and NORX64-6-1.

For use cases where very high data throughput is necessary, we recommend NORX64-4-4, which allows payload encryption on four parallel lanes, thus enabling very high data processing

[1] Nevertheless, the designers discourage this approach, and recommend that nonce freshness should be ensured by all means.

speeds. Finally, we advise hardware implementers not to realise multiple instances of NORX with different parameter combinations at the same time. This holds especially for different values of the parallelism degree p. An implementation should rather be optimised for one set of parameters to gain higher efficiency.

5.3 Performance

NORX was designed to perform well across both software and hardware. This section details our implementations and performance results.

5.3.1 Generalities

In this part we analyse some general performance-relevant properties of NORX, like number of operations in G and F^l, parallelism degree, and upper bounds for the speed of NORX on different platforms.

Number of Operations

Table 5.1 shows the number of operations required for the NORX core functions. We omit the overhead of initialisation, integration of parameters, domain separation constants, padding messages, and so on, as those costs are negligible compared to that of the core permutation F^l.

Table 5.1: Overview on the number of operations of the NORX functions

function	#XOR	#AND	#shifts	#rotations	total
G	12	4	4	4	24
F	96	32	32	32	192
F^4	384	128	128	128	768
F^6	576	192	192	192	1152
F^8	768	256	256	256	1536
F^{12}	1152	384	384	384	2304

Memory

NORX32 and NORX64 require at least 16 bytes to be stored in ROM to generate the initialisation constants[2]. To store all initialisation constants 32 and 64 bytes of ROM are necessary.

Processing a message in NORX requires enough RAM to store the internal state, i.e., 64 bytes in NORX32 and 128 bytes in NORX64. The data being processed need not be in memory for more than 1 byte at a time. In practice, however, it is preferable to process blocks of 48 (resp. 96) bytes at a time.

[2]Note that the 8 constants can be generated on-the-fly from $0, \ldots, 15$, see §3.5.2.

Parallelism

The core permutation F of NORX has a natural parallelism of 4 independent G applications. Additionally, NORX allows for greater parallelism levels using multiple lanes. Using the $p = 0$ mode, see Line 7, the internal parallelism level of NORX is effectively unbounded for long enough messages.

5.3.2 Software

NORX is easily implemented for 32-bit and 64-bit processors, as it works on 32- and 64-bit words and uses only word-based operations (`XOR`, `AND`, shifts and rotations). The specification can directly be translated to code and requires no specific technique such as look-up tables or bitslicing. The core of NORX essentially consists of repeated usage of the G function, which allows simple and compact implementations (e.g., by having only one copy of the G code).

Furthermore, constant-time implementations of NORX are straightforward to write, due to the absence of secret-dependent instructions or branchings.

Bit Interleaving

While NORX's lack of integer addition avoids dealing with carry chains, the implementer may still have to perform full-word rotations and shifts in words wider than the natural CPU word size. In 8-bit processors, some of this burden is alleviated by 2 out of 4 rotations being multiples of 8. However, this is only a half-measure.

Instead, the implementer can employ the *bit interleaving* technique presented in [24]. This technique consists of splitting an n-bit word w into $s = n/m$ m-bit words b_i, with $b_{ij} = w_{i+jn/m}$. A rotation by r in this representation can be performed by rotating each b_i by $\lfloor r/w \rfloor + 1$ if $i + r \bmod m < r$, $\lfloor r/w \rfloor$ otherwise, and moving b_i to $b_{i+r \bmod m}$. Rotations by 1 or $n - 1$ are particularly attractive, since they result in a single m-bit rotation. For example, consider implementing NORX64 on a 32-bit CPU. Each state word w will be split into the 2 words b_0 and b_1. To rotate by r:

- If $r \bmod 2 = 0$, rotate both b_0 and b_1 by $\lfloor r/2 \rfloor$;
- If $r \bmod 2 = 1$, rotate b_1 by $\lfloor r/2 \rfloor + 1$, b_0 by $\lfloor r/2 \rfloor$, and swap them.

Conversion between representations can be performed in logarithmic time using bit "zip" and "unzip" operations [8].

Avoiding Latency

One drawback of G is that it has little instruction parallelism. In architectures where one is limited by the latency of the G function, an implementer can trade a few extra instructions by

reduced latency:

$$\begin{aligned} t_0 &\longleftarrow a \oplus b \\ t_1 &\longleftarrow a \wedge b \\ t_1 &\longleftarrow t_1 \ll 1 \\ a &\longleftarrow t_0 \oplus t_1 \\ d &\longleftarrow d \oplus t_0 \\ d &\longleftarrow d \oplus t_1 \\ d &\longleftarrow d \ggg r_0 \end{aligned}$$

This tweak saves up to 1 cycle per instruction sequence, of which there are 4 per G, at the cost of 1 extra instruction (cf. Fig. 5.1). In a sufficiently parallel architecture, this can save at least $4 \times 2 \times l$ cycles, which translates to $6.4l/w$ cycles per byte saved overall. In our measurements, this translated to a performance improvement of NORX from 0.4 to 0.7 cycles per byte, depending on the target architecture, word size, and number of rounds.

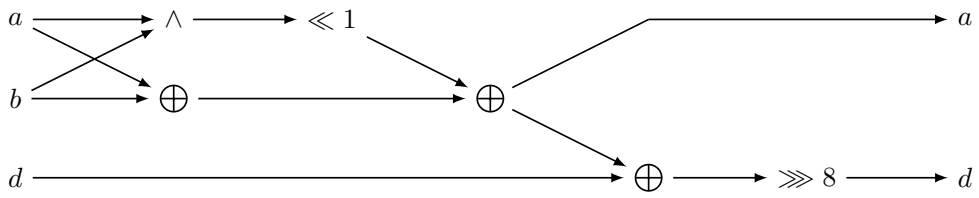

(a) Naïve implementation of the G instruction sequence

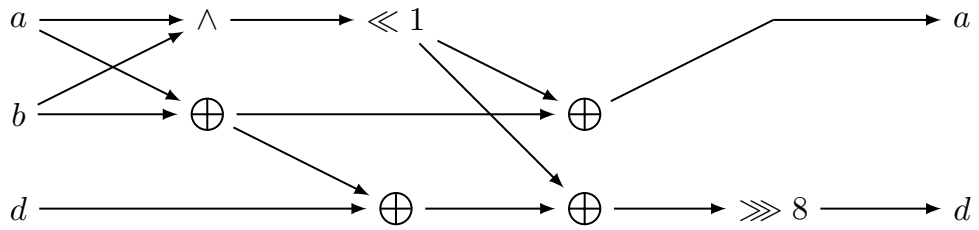

(b) Latency-oriented version of the G instruction sequence

Figure 5.1: Improving the latency of G.

Vectorization

NORX lends itself quite well to implementations taking advantage of SIMD extensions present in modern processors, such as AVX or NEON.

The typical vectorized implementation of NORX, when $p = 1$, works in full rows of the 4×4 state, and computes whole column and diagonal steps of F in parallel.

Results

We wrote portable C reference implementations for both NORX64 and NORX32, as well as optimized versions for CPUs supporting AVX and AVX2 and for NEON-enabled ARMs. Table 5.2

shows speed measurements on various platforms for messages with varying lengths. The listed CPU frequencies are nominal ones, i.e. without dynamic overclocking features like Turbo Boost, which improves the accuracy of measurements. Furthermore we listed only those platform-compiler combinations that achieved the highest speeds. Unless stated otherwise we used the compiler flags

```
-O3 -march=native -std=c89 -Wall -pedantic -Wno-long-long
```

The top speed of NORX (for $p = 1$), in terms of bytes per second, was achieved by an AVX2 implementation of NORX64-4-1 on a Haswell CPU, listed in Table 5.2. It achieves a throughput of about 1.75 GiBps (1.99 cycles per byte at 3.5 GHz). The overhead for short messages (≤ 64 bytes) is mainly due to the additional initialisation and finalisation rounds (see Fig. 3.1). However the cost per byte quickly decreases, and stabilizes for messages larger than about 1 KiB.

Note that the speed between reference and optimized implementations differs by a factor of less than 2, suggesting that straightforward and portable implementations will provide sufficient performance in most applications. Such consistent performance reduces development costs and improves interoperability.

5.3.3 Hardware

Hardware architectures of NORX are efficient and easy to design from the specification: vertical and parallel folding are naturally derived from the iterated and parallel structure of NORX. The cipher benefits from the hardware-friendliness of the function G, which requires only bitwise logical `AND`, `XOR`, and bit shifts, and the iterated usage of G inside the core permutation of NORX.

A hardware architecture was designed, supporting parameters $w \in \{32, 64\}$, $l \in \{2, \ldots, 16\}$ and $p = 1$. It was synthesized with the Synopsys Design Compiler for an ASIC using 180 nm UMC technology. The implementation was targeted at high data throughput. The requirements in area amounted to about 62 kGE. Simulations for NORX64-4-1 report a throughput of about 10 Gbps (1.2 GiBps), at a frequency of 125 MHz.

A more thorough evaluation of all hardware aspects of NORX is planned for the future. Due to the similarity of NORX to ChaCha and the fact that NORX has only little overhead compared to a blank stream cipher, we expect results similar to those of Chacha as presented in [37].

Table 5.2: Software performance of NORX in cycles per byte

data length [byte]		long	4096	1536	576	64	8
Samsung Exynos 4412 Prime (Cortex-A9) at 1.7 GHz							
NORX32-4-1	Ref	16.72	18.03	20.52	27.92	109.48	771.88
	NEON	9.27	10.20	11.95	16.46	72.30	521.00
NORX64-4-1	Ref	15.60	17.91	22.02	32.42	148.55	1177.12
	NEON	7.13	8.40	10.61	16.25	82.12	648.88
BeagleBone Black Rev B (Cortex-A8) at 1.0 GHz							
NORX32-4-1	Ref	16.66	17.90	20.28	26.49	102.34	708.00
	NEON	9.49	10.52	12.36	17.92	75.62	550.12
NORX64-4-1	Ref	17.24	19.81	24.34	35.73	164.86	1317.50
	NEON	7.00	8.35	10.67	16.44	85.66	680.00
Intel Core i7-2630QM at 2.0 GHz							
NORX64-6-1	Ref	6.33	7.02	8.24	13.96	70.62	607.50
	AVX	4.02	4.42	5.14	6.90	63.75	204.00
NORX64-4-1	Ref	4.83	5,35	6.30	8.66	50.00	400.62
	AVX	2.68	2.96	3.45	4.66	17.18	137.5
Intel Core i7-3667U at 2.0 GHz							
NORX64-6-1	Ref	8.15	9.01	10.49	14.15	53.20	425.62
	AVX	5.04	5.56	6.45	8.65	32.19	255.00
NORX64-4-1	Ref	5.58	6,17	7.22	9.82	38.05	303.75
	AVX	3.37	3.72	4.35	5.84	22.11	174.38
Intel Core i7-4770K at 3.5 GHz							
NORX64-6-1	Ref	5.37	5.94	6.92	9.40	36.44	292.00
	AVX2	2.98	3.29	3.84	5.17	19.00	153.00
NORX64-4-1	Ref	3.98	4.39	5.11	6.97	27.19	217.00
	AVX2	1.99	2.20	2.58	3.49	12.94	104.50

6 Design Rationale

In this chapter we motivate the design choices made in NORX. We pursue a top-down approach, starting with the general layout and going into the details of the cipher's components in the later sections.

6.1 The Parallel Duplex Construction

The layout of NORX is based on the monkeyDuplex construction [20, 23], but enhanced by the capability of parallel payload processing on multiple lanes (cf. Figs. 3.1 and 3.2). The *parallel duplex construction* is similar to the tree-hashing mode for sponge functions [22]. It allows NORX to take advantage of multi-core processors and enables high-throughput hardware implementations. Associated data can be authenticated as header and/or trailer data but only on a single lane. We felt that it is not worth the effort to enable processing of A and Z in parallel, as they are usually rather short. The number of encryption lanes is controlled by the parallelism degree $0 \leq p \leq 255$, which is a fixed instance parameter. Hence two instances with distinct p values cannot decrypt data from each other. Obviously the same holds for differing w and l values.

To ensure that the payload blocks on parallel lanes are encrypted with distinct key streams, we use the branching phase to include an id into each of the parallel lanes. For NORX the id is a simple counter. Once the parallel payload processing is finished, the states are re-combined in the merging phase and NORX advances to the processing of the trailer (if present) or generation of the authentication tag.

There does not exist a formal proof of security for the parallel duplex construction yet. Note that the most problematic step could be the merging phase for $p \neq 1$, due to the fact that (multi-)collisions could occur. However, we expect that the construction is safe in case of a nonce-respecting adversary. We will try to hand in the proof at a later point of time.

6.2 The G Function

The G function of NORX is inspired by the quarter-round function of the stream cipher ChaCha [17], which itself is an advancement of the quarter-round function of the eSTREAM finalist Salsa20 [1, 18]. Variants of ChaCha's quarter-round function can be found for example in the SHA-3 finalist BLAKE [2, 11] and its successor BLAKE2 [13].

Overview

One of the main goals for NORX was to design a core primitive, which does not rely on integer addition to introduce non-linearity. Instead it should use exclusively more hardware-friendly bitwise logic operations like `NOT`, `AND`, `OR`, or `XOR` and bit-shifts. Fig. 6.1 shows how the G function of NORX transforms an input (a, b, c, d) compared to the quarter-round function of

ChaCha . The rotation offsets for NORX are specified in Table 3.2. The offsets of ChaCha are $(s_0, s_1, s_2, s_3) = (16, 12, 8, 7)$ for 32-bit and $(s_0, s_1, s_2, s_3) = (32, 24, 16, 63)$ for 64-bit.[1]

$$
\begin{array}{lcl|lcl}
a & \longleftarrow & (a \oplus b) \oplus \big((a \wedge b) \ll 1\big) & a & \longleftarrow & a + b \\
d & \longleftarrow & (a \oplus d) \ggg r_0 & d & \longleftarrow & (a \oplus d) \ggg s_0 \\
c & \longleftarrow & (c \oplus d) \oplus \big((c \wedge d) \ll 1\big) & c & \longleftarrow & c + d \\
b & \longleftarrow & (b \oplus c) \ggg r_1 & b & \longleftarrow & (b \oplus c) \ggg s_1 \\
a & \longleftarrow & (a \oplus b) \oplus \big((a \wedge b) \ll 1\big) & a & \longleftarrow & a + b \\
d & \longleftarrow & (a \oplus d) \ggg r_2 & d & \longleftarrow & (a \oplus d) \ggg s_2 \\
c & \longleftarrow & (c \oplus d) \oplus \big((c \wedge d) \ll 1\big) & c & \longleftarrow & c + d \\
b & \longleftarrow & (b \oplus c) \ggg r_3 & b & \longleftarrow & (b \oplus c) \ggg s_3
\end{array}
$$

Figure 6.1: Comparison of NORX (left) and ChaCha (right) core functions

In NORX the integer additions is replaced by the following expression

$$x \longleftarrow (x \oplus y) \oplus \big((x \wedge y) \ll 1\big)$$

which uses bitwise logical `AND` to introduce non-linearity. It mimics integer addition of two bit strings x and y with a 1-bit carry propagation and thus provides, in addition to non-linearity, also a slight diffusion of bits. In conformity with the main design principle of NORX, we tried to make the non-linear operation as simple as possible in order to simplify cryptanalysis and to reduce the risk of overlooking potential security weaknesses. Moving to simple bitwise logical operations facilitates hardware implementations. One way to realise G as a circuit is depicted in Fig. 6.2.

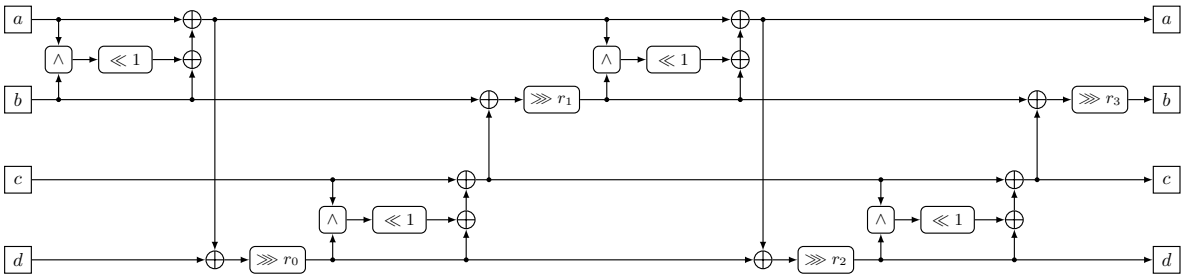

Figure 6.2: The G circuit

Bijectivity

The only expression in G which is not obviously invertible at a first glance, is the non-linear operation

$$z = (x \oplus y) \oplus ((x \wedge y) \ll 1)$$

[1]The original ChaCha stream cipher is defined for 32-bit words. For the 64-bit version we used the rotation offsets $(32, 24, 16, 63)$ from the BLAKE2 specification [13].

with n-bit words x, y and z. In order to proof bijectivity of the above expression we show how to invert it, under the assumption that one of its inputs is fixed. Therefore we write $x = \sum_{i=0}^{n-1} x_i \cdot 2^i$, $y = \sum_{i=0}^{n-1} y_i \cdot 2^i$ and $z = \sum_{i=0}^{n-1} z_i \cdot 2^i$ with x_i, y_i and $z_i \in \{0, 1\}$ and assume that y is fixed. Writing down the inverse non-linear operation at bit level is then straightforward:

$$\begin{aligned}
x_0 &= (z_0 \oplus y_0) \\
x_1 &= (z_1 \oplus y_1) \oplus (x_0 \wedge y_0) \\
&\vdots \\
x_i &= (z_i \oplus y_i) \oplus (x_{i-1} \wedge y_{i-1}) \\
&\vdots \\
x_{n-1} &= (z_{n-1} \oplus y_{n-1}) \oplus (x_{n-2} \wedge y_{n-2})
\end{aligned}$$

This proves that G is indeed a permutation. Further, it is a permutation when either of its input arguments is fixed, making it also a latin square.

Features

The only operations required to define G are bitwise `XOR`, `AND` and logical bit shifts, which has several advantages: All of the mentioned instructions can be implemented in constant time regardless of the word size. Especially for hardware implementations there are no carry-propagations to worry about, for example, as there would be for integer addition $\bmod 2^n$.
Moreover no table-lookup instructions, like SBoxes, are required, where the table index is data-dependent. Those operations, if not handled with extreme care, are often the reason for implementations leaking side-channel information, making the affected algorithm vulnerable, e.g., to timing-attacks [15]. By avoiding them, the task of hardening the cipher against side-channel attacks gets obviously much easier. No specialised implementations are required, e.g., bit-sliced SBoxes [5, 31], for table-lookups in constant time. Additionally, the waiving of more sophisticated instructions like integer addition, multiplication, Galois field arithmetic or other constructs based on linear algebra, has the effect that the algorithm is much easier to implement (both in soft- and hardware) and thus reduces the threat of introducing unwanted bugs.

6.3 The F Function

The layout of the round function F of NORX is the same as used in ChaCha [17].

Overview

Recall that F transforms a state $S = s_0 \parallel \cdots \parallel s_{15}$ in two phases. First a column step is applied

$$\mathsf{G}(s_0, s_4, s_8, s_{12}) \quad \mathsf{G}(s_1, s_5, s_9, s_{13}) \quad \mathsf{G}(s_2, s_6, s_{10}, s_{14}) \quad \mathsf{G}(s_3, s_7, s_{11}, s_{15})$$

followed by a diagonal step

$$\mathsf{G}(s_0, s_5, s_{10}, s_{15}) \quad \mathsf{G}(s_1, s_6, s_{11}, s_{12}) \quad \mathsf{G}(s_2, s_7, s_8, s_{13}) \quad \mathsf{G}(s_3, s_4, s_9, s_{14})$$

Bijectivity

As G is a permutation, F is obviously a permutation, too. This means that there exist no states S and S', with $S \neq S'$, which produce the same result, i.e. $\mathsf{F}^l(S) = \mathsf{F}^l(S')$, after any number of rounds l. This characteristic of F is important for the duplex construction [23, 20] in order to retain some desirable security properties.

Features

One great advantage of the ChaCha-related layout of F is, that the modification of a single bit in the input has the chance of affecting all 16 output words[2] after only one application of F. This features greatly enhances diffusion. Another benefit of the layout is the ability to execute the four applications of G in a step completely in parallel, which improves performance.

6.4 Number of Rounds

For a higher protection of the key and authentication tag, e.g. against differential cryptanalysis, we chose twice the number of rounds for initialisation and finalisation, compared to the data processing phases. This measure was already proposed in [20] and has only minor effects on the overall performance, but greatly increases the security of NORX. The minimal value of $l = 4$ is based on the following observations:

1. The best attacks on Salsa20 and ChaCha [10, 50, 52] break 8 and 7 rounds, respectively, which roughly corresponds to 4 and 3.5 rounds of the NORX core. However this is within a much stronger attack model than that provided by the duplex construction of NORX.
2. The preliminary cryptanalysis of NORX as presented in Chapter 7. The best differentials we were able to find, belong to a class of high-probability truncated differentials over 1.5 rounds and a class of impossible differentials over 3.5 rounds. Despite the fact that those differentials cannot be used to mount an attack on NORX, it might be possible to find similar differentials, using more advanced cryptanalytic techniques, which could be used for an attack.

The number of rounds may be adjusted according to the future cryptanalytic results on NORX.

6.5 Selection of Constants

6.5.1 Initialisation

The initialisation constants are listed in Table 3.4 and are derived through

$$(u_0, \ldots, u_{15}) = \mathsf{F}^2(0, \ldots, 15)$$

as already mentioned in §3.5.2. This approach allows an on-the-fly computation, if necessary, and is meant to provide transparency in order to show that the values belong to the "nothing-up-my-sleeves" category, i.e. that they were selected in such a way that there is no possibility

[2] In fact we found for NORX only one case where less than 16 words are affected. This can be achieved through the modification of three very specific bits in the input. See chapter Chapter 7 on cryptanalysis for more details.

to hide a backdoor. The main purpose of the initialisation constants is to provide some asymmetry during initialisation and to limit the freedom where differences can be injected by an attacker.

6.5.2 Domain Separation

The NORX algorithm is separated into different data processing phases. Each phase uses its own domain separation constant to mark the end of certain events like the absorbing of data blocks or merging and branching steps in case of an instance with parallelism degree $p \neq 1$. A domain separation constant is always added to the least significant byte of the capacity word s_{15}. The constants are given in Table 3.3. The separation of the processing phase is important for the security proofs of the indifferentiability of the duplex construction [21, 23]. In addition they help to break the self-similarity of the round function and thus increase the complexity of certain kind of attacks on NORX, for example, like slide attacks, see §7.4.2.

6.5.3 Rotation Offsets

The rotation offsets (r_0, r_1, r_2, r_3) used by NORX provide a good balance between security and efficiency. The values r_i, with $0 \leq i \leq 3$, were selected according to the following conditions:

1. At least two out of four offsets are multiples of 8.
2. The remaining offsets are odd and have the form $8n \pm 1$ or $8n \pm 3$, with a preference for the first shape.

The motivation behind those criteria has the following reasons: An offset which is a multiple of 8 preserves byte alignment and thus is much faster than an unaligned rotation on many non-64-bit architectures. Many 8-bit microcontrollers have only 1-bit shifts of bytes, so for example rotations by 5 bits are particularly expensive. Using aligned rotations, i.e. permutations of bytes, greatly increases the performance of the entire algorithm. Even 64-bit architectures benefit from such aligned rotations, for example when an instruction sequence of two shifts followed by `XOR` can be replaced by SSSE3's byte shuffling instruction `pshufb`. Odd offsets break up the byte structure and therefore increase diffusion.

In order to find good rotation offsets and assess their diffusion properties, we used an automated search combined with a diffusion test. Therefore let l denote a round number and let L and L_l be lists. For each offset tuple (r_0, r_1, r_2, r_3) with $r_i \in \{1, \ldots, w-1\}$ satisfying the above criteria, the following steps are repeated 10^6 times, after the offsets have been plugged into G:

1. Choose two b-bit sized states S and S' uniformly at random, such that $\mathsf{hw}(S \oplus S') = 1$.
2. Compute $X = \mathsf{F}^l(S) \oplus \mathsf{F}^l(S')$, where F denotes the round function of NORX.
3. Save $\mathsf{hw}(X)$ to L_l.

After the above loop is finished the test computes minimum, maximum, average and median values of the elements of L_l, saves the latter together with the offsets to L and resets L_l. Then it proceeds to the analysis of the next rotation tuple. This test is repeated until all candidate offsets have been processed.

Finally, we chose the offsets $(8, 19, 40, 63)$ for NORX64 and $(8, 11, 16, 31)$ for NORX32, which belonged to those having very high values for average and median Hamming weight for $l = 1$, achieve full diffusion after $l = 2$, and additionally offer good performance.

Table 6.1 lists the results of the test for 32- and 64-bit core functions with $l \leq 4$ and rotation offsets as specified above. The test results show that the diffusion speed of NORX's round function F is almost as high as ChaCha's and that full diffusion is reached after two rounds. Fig. 6.3 shows how single bit changes in the word s_0 propagate through the NORX state over the course of 5 steps ($= \mathrm{F}^{2.5}$). Unfortunately there seems to be no combination of rotation values with 3 offsets being a multiple of 8 and one being $w - 1$, like BLAKE2's $(32, 24, 16, 63)$, where F achieves a comparably strong diffusion as illustrated in Table 6.1. The reason for this can be traced back to the replacement of integer addition by the non-linear operation of NORX.

Table 6.1: Diffusion statistics for NORX and ChaCha round functions

	NORX32				ChaCha (32-bit)			
l	min	max	avg	med	min	max	avg	med
1	83	280	179.22	181	73	294	182.19	185
2	194	307	256.02	256	199	312	255.99	256
3	198	312	255.99	256	204	313	255.98	256
4	201	307	255.99	256	200	314	255.98	256
	NORX64				ChaCha (64-bit)			
l	min	max	avg	med	min	max	avg	med
1	95	429	230.13	222	73	506	248.84	246
2	440	589	511.98	512	430	591	512.01	512
3	434	589	512.00	512	439	589	511.97	512
4	428	589	511.98	512	435	585	512.00	512

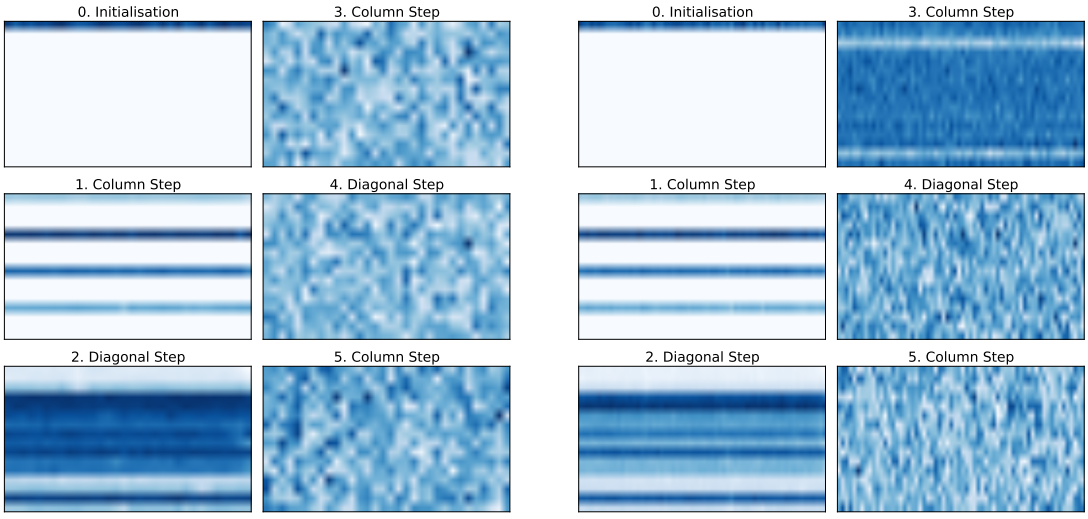

Figure 6.3: Visualisation of NORX diffusion

6.6 The Padding Rule

The sponge (or duplex) construction offers protection against generic attacks if the padding rule is sponge-compliant, i.e. if it is injective and ensures that the last block is different from the all-zero block. In [22] it has been proven that the multi-rate padding satisfies those properties. Moreover it is simple to describe, easy to implement and very efficient. Thus it was a natural choice to be used in NORX. Additionally, the multi-rate padding increases the complexity to mount certain kind of attacks on NORX, like slide attacks, see §7.4.2.

6.7 Absence of Backdoors

We, the designers of NORX, faithfully declare that we have not inserted any hidden weaknesses in this cipher.

7 Security Analysis

This chapter presents preliminary cryptanalysis of NORX.

7.1 Security Bounds for the Mode of Operation

NORX is, at its core, a keyed sponge. Let $\Pi = (\mathcal{E}, \mathcal{D})$ denote NORX, with encryption function $\mathcal{E}$, decryption function $\mathcal{D}$, and based on an ideal underlying permutation p. Then the following privacy and authenticity security bounds are satisfied

$$\mathbf{Adv}^{\mathsf{priv}}_{\Pi}(q_p, q_{\mathcal{E}}, \lambda_{\mathcal{E}}) \leq \frac{3(q_p + \sigma_{\mathcal{E}})^2}{2^{b+1}} + \left(\frac{8eq_p\sigma_{\mathcal{E}}}{2^b}\right)^{1/2} + \frac{rq_p}{2^c} + \frac{q_p + \sigma_{\mathcal{E}}}{2^k}$$

$$\begin{aligned}\mathbf{Adv}^{\mathsf{auth}}_{\Pi}(q_p, q_{\mathcal{E}}, \lambda_{\mathcal{E}}, q_{\mathcal{D}}, \lambda_{\mathcal{D}}) \leq{}& \frac{(q_p + \sigma_{\mathcal{E}} + \sigma_{\mathcal{D}})^2}{2^b} + \left(\frac{8eq_p\sigma_{\mathcal{E}}}{2^b}\right)^{1/2} + \frac{rq_p}{2^c} \\ &+ \frac{q_p + \sigma_{\mathcal{E}} + \sigma_{\mathcal{D}}}{2^k} + \frac{(q_p + \sigma_{\mathcal{E}} + \sigma_{\mathcal{D}})\sigma_{\mathcal{D}}}{2^c} + \frac{q_{\mathcal{D}}}{2^t}\end{aligned}$$

where r, c, b, k and t are rate, capacity, state, key and tag sizes, e is Euler's number, q_p are the number of permutation queries, $q_{\mathcal{E}}$ are the number of encryption queries of total length $\lambda_{\mathcal{E}}$ and $\sigma_{\mathcal{E}}$ is specified as follows:

$$\sigma_{\mathcal{E}} := \sum_{j=1}^{q_{\mathcal{E}}} \sigma_{\mathcal{E},j} \leq \begin{cases} 2\lambda_{\mathcal{E}} + 4q_{\mathcal{E}}, & \text{if } p = 0 \\ \lambda_{\mathcal{E}} + 3q_{\mathcal{E}}, & \text{if } p = 1 \\ \lambda_{\mathcal{E}} + (p+4)q_{\mathcal{E}}, & \text{if } p > 1 \end{cases}$$

The values $q_{\mathcal{D}}$, $\lambda_{\mathcal{D}}$ and $\sigma_{\mathcal{D}}$ for decryption $\mathcal{D}$ are specified analogously.

In summary, the NORX mode of operation achieves security levels of $\min\{2^{b/2}, 2^c, 2^k\}$ assuming an ideal underlying permutation p and, intuitively speaking, offers authenticity as long as it offers privacy and the $\frac{(q_p+\sigma_{\mathcal{E}}+\sigma_{\mathcal{D}})\sigma_{\mathcal{D}}}{2^c}$ term—quadratic on the number of forgery attempts—is negligible. For more information on the security proofs see [39, 6].

7.2 Differential Cryptanalysis

Differential attacks cover all attacks that exploit non-ideal propagation of differences in a cryptographic algorithm (or of its components). Differential cryptanalysis is one of the standard tools in the repertoire of every cryptanalyst and usually a lot of attacks on a cipher are at least partially differential. It is thus crucial to analyse the resistance of new designs to differential attacks.

First we introduce some of the required notations, then we analyse the propagation of differences through the G function, show how to construct high-probability truncated differentials of low weight for the core permutation F^l and finally study impossible differential cryptanalysis.

7.2.1 Notation

Definition 1. Let x and x' be n-bit strings. We call $\alpha = x \oplus x'$ the *difference* of x and x' with respect to bitwise XOR. Furthermore for tuples of n-bit strings $(x_0, \ldots, x_{m-1})$ and $(x'_0, \ldots, x'_{m-1})$ we call the component-wise difference

$$(\alpha_0, \ldots, \alpha_{m-1}) = (x_0, \ldots, x_{m-1}) \oplus (x'_0, \ldots, x'_{m-1}) = (x_0 \oplus x'_0, \ldots, x_{m-1} \oplus x'_{m-1})$$

a *tuple of differences*.

Definition 2. An n-bit difference α with $\mathrm{hw}(\alpha) = m$ and 1-entries at bit positions $0 \leq i_0 \leq \cdots \leq i_m \leq n-1$ is denoted by $\alpha[i_0, \ldots, i_m]$.

Definition 3. Let $f : \{0,1\}^{m \cdot n} \longrightarrow \{0,1\}^{k \cdot n}$, $f(a_0, \ldots, a_{m-1}) = (b_0, \ldots, b_{k-1})$ be a boolean function. Let $\alpha := (\alpha_0, \ldots, \alpha_{m-1}) = (x_0, \ldots, x_{m-1}) \oplus (x_0, \ldots, x_{m-1})$ and let $\beta := (\beta_0, \ldots, \beta_{k-1}) = f(x_0, \ldots, x_{m-1}) \oplus f(x'_0, \ldots, x'_{m-1})$ be tuples of differences. Then we call (α, β) a *differential* with respect to the function f and denote it by

$$\alpha \xrightarrow{f} \beta$$

If the context is clear we skip the f above the arrow and just write $\alpha \longrightarrow \beta$. Furthermore, we call α an *input difference* and β an *output difference* of f.

In our later analysis of NORX we usually consider functions f having $k = 1$ or $k = m$.

Definition 4. Let $f_0, \ldots, f_{l-1}$ be boolean functions defined by

$$f_i : \{0,1\}^{m \cdot n} \longrightarrow \{0,1\}^{m \cdot n},\ f_i(a_0, \ldots, a_{m-1}) = (b_0, \ldots, b_{m-1})$$

for $i \in \{0, \ldots, l-1\}$. Further let $\alpha^0 := (\alpha^0_0, \ldots, \alpha^0_{m-1}), \ldots, \alpha^l := (\alpha^l_0, \ldots, \alpha^l_{m-1})$ be tuples of differences such that

$$\alpha^i \xrightarrow{f_i} \alpha^{i+1}$$

Then we call $(\alpha^0, \ldots, \alpha^l)$ a *differential characteristic* with respect to the functions $f_0, \ldots, f_{l-1}$ and denote it by

$$\alpha^0 \xrightarrow{f_0} \ldots \xrightarrow{f_{i-1}} \alpha^i \xrightarrow{f_i} \ldots \xrightarrow{f_l} \alpha^l$$

The tuples α^j with $j \in \{1, \ldots, l-1\}$ are also called *internal differences*. In the case where $f := f_0 = \cdots = f_{l-1}$ we also say that $(\alpha^0, \ldots, \alpha^l)$ is a differential characteristic with respect to the *iterated function* f.

The notion of a differential characteristic can obviously be defined for arbitrary boolean functions f_i, but it is not required at this point. Thus, for reasons of simplicity, we decided to define it only for the special case, where the dimension of the domain equals the dimension of the codomain of f_i.

Definition 5. Every differential (α, β) of a function f has a probability $p \in [0,1]$ associated to it, which will be written as

$$\Pr(\alpha \xrightarrow{f} \beta) = p$$

To capture all those informations in a compact form, we denote a differential (α, β) of probability p with respect to a function f by:

$$\alpha \xrightarrow[p]{f} \beta$$

We use the commonly accepted assumption that the probability of a differential is equal to the sum of probabilities of all differential characteristics corresponding to this differential. Moreover it is commonly assumed that the probability of the best differential can accurately be estimated by the probability of the best differential characteristic.

7.2.2 Differential Properties of G

In this section we analyse how n-bit input differences α with $\mathsf{hw}(\alpha) = 1$ propagate through G and present the probabilities of the resulting output differences. Therefore, we decompose G into two functions G_1 and G_2 and initially analyse the behaviour of G_1.

Definition 6. Let $\mathsf{G}_1 : \{0,1\}^{4n} \longrightarrow \{0,1\}^{4n}$ be defined as

$$\begin{aligned} a &\longleftarrow (a \oplus b) \oplus ((a \wedge b) \ll 1) \\ d &\longleftarrow (a \oplus d) \ggg r_0 \\ c &\longleftarrow (c \oplus d) \oplus ((c \wedge d) \ll 1) \\ b &\longleftarrow (b \oplus c) \ggg r_1 \end{aligned}$$

The function G_2 is defined analogously to G_1 but with rotation offsets r_2 and r_3, instead of r_0 and r_1. Thus, we obviously have $\mathsf{G}(a,b,c,d) = \mathsf{G}_2(\mathsf{G}_1(a,b,c,d))$.

Let (x_0, x_1, x_2, x_3) and (x'_0, x'_1, x'_2, x'_3) be two tuples of n-bit strings having difference

$$(\alpha_0, \alpha_1, \alpha_2, \alpha_3) = (x_0, x_1, x_2, x_3) \oplus (x'_0, x'_1, x'_2, x'_3)$$

and let

$$(\beta_0, \beta_1, \beta_2, \beta_3) = \mathsf{G}_1(x_0, x_1, x_2, x_3) \oplus \mathsf{G}_1(x'_0, x'_1, x'_2, x'_3)$$

Further assume that $\mathsf{hw}(a_v) = 1$ for a fixed $v \in \{0, \ldots, 3\}$ where the 1-entry is a bit position i and $\mathsf{hw}(a_u) = 0$ for all $u \in \{0, \ldots, 3\} \setminus \{v\}$. Then we get differentials

$$(\alpha_0, \alpha_1, \alpha_2, \alpha_3) \xrightarrow{\mathsf{G}_1} (\beta_0, \beta_1, \beta_2, \beta_3)$$

and associated probabilities as presented in Table 7.1. Note that the output difference β_w, for $w \in \{0, \ldots, 3\}$ is the XOR sum of the 1-bit differences $\beta_w[j]$ in a given column. The resulting $\beta_w[j]$ do not hold for arbitrary $\alpha_v[i]$ with $i \in \{0, \ldots, n-1\}$. For example if $i = n-1$ the difference $\alpha_v[i]$ will be erased by the shift operation $\alpha_v[i] \ll 1$, thereby cancelling all output differences depending [1] on the latter.

The differentials in Table 7.1 only hold for input differences having exactly one active bit. Obviously, when allowing input differences with a larger number of active bits the situation gets immediately a lot more complex. This could lead to situations where active bits of different words interact and cancel each other out. For example an input difference $(\alpha_0[n-1], \alpha_1[n-1], 0, 0)$ leads to a cancellation of the probability 1 output difference $\alpha_0[n-1]$ in the output word a: The two active bits in the input words a and b neutralise each other during the update of the word a. We will see below how this property can be exploited to build differentials for G having high probability and low weight output differences.

[1] We refer to Fig. C.1 in the appendix for a visualisation of the relations between input and output differences of G_1.

Table 7.1: Output differences $\beta_w[j]$ and their probabilities after G_1 on an input difference $\alpha_v[i]$

	$\beta_0[j]$	$\beta_1[j]$	$\beta_2[j]$	$\beta_3[j]$	$\Pr(\alpha_v[i] \xrightarrow{G_1} \beta_w[j])$
$\alpha_0[i]$	$\alpha_0[i]$	0	0	0	1
	$\alpha_0[i] \ll 1$	0	0	0	2^{-1}
	0	$\alpha_0[i] \ggg (r_0 + r_1)$	0	0	1
	0	$((\alpha_0[i] \ggg r_0) \ll 1) \ggg r_1$	0	0	2^{-1}
	0	$(\alpha_0[i] \ll 1) \ggg (r_0 + r_1)$	0	0	2^{-1}
	0	$(((\alpha_0[i] \ll 1) \ggg r_0) \ll 1) \ggg r_1$	0	0	2^{-2}
	0	0	$\alpha_0[i] \ggg r_0$	0	1
	0	0	$(\alpha_0[i] \ggg r_0) \ll 1$	0	2^{-1}
	0	0	$(\alpha_0[i] \ll 1) \ggg r_0$	0	2^{-1}
	0	0	$((\alpha_0[i] \ll 1) \ggg r_0) \ll 1$	0	2^{-2}
	0	0	0	$\alpha_0[i] \ggg r_0$	1
	0	0	0	$(\alpha_0[i] \ll 1) \ggg r_0$	2^{-1}
$\alpha_1[i]$	$\alpha_1[i]$	0	0	0	1
	$\alpha_1[i] \ll 1$	0	0	0	2^{-1}
	0	$\alpha_1[i] \ggg r_1$	0	0	1
	0	$\alpha_1[i] \ggg (r_0 + r_1)$	0	0	1
	0	$((\alpha_1[i] \ggg r_0) \ll 1) \ggg r_1$	0	0	2^{-1}
	0	$(\alpha_1[i] \ll 1) \ggg (r_0 + r_1)$	0	0	2^{-1}
	0	$(((\alpha_1[i] \ll 1) \ggg r_0) \ll 1) \ggg r_1$	0	0	2^{-2}
	0	0	$\alpha_1[i] \ggg r_0$	0	1
	0	0	$(\alpha_1[i] \ggg r_0) \ll 1$	0	2^{-1}
	0	0	$(\alpha_1[i] \ll 1) \ggg r_0$	0	2^{-1}
	0	0	$((\alpha_1[i] \ll 1) \ggg r_0) \ll 1$	0	2^{-2}
	0	0	0	$\alpha_1[i] \ggg r_0$	1
	0	0	0	$(\alpha_1[i] \ll 1) \ggg r_0$	2^{-1}
$\alpha_2[i]$	0	$\alpha_2[i] \ggg r_1$	0	0	1
	0	$(\alpha_2[i] \ll 1) \ggg r_1$	0	0	2^{-1}
	0	0	$\alpha_2[i]$	0	1
	0	0	$\alpha_2[i] \ll 1$	0	2^{-1}
$\alpha_3[i]$	0	$\alpha_3[i] \ggg (r_0 + r_1)$	0	0	1
	0	$(\alpha_3[i] \ll 1) \ggg (r_0 + r_1)$	0	0	2^{-1}
	0	0	$\alpha_3[i] \ggg r_0$	0	1
	0	0	$(\alpha_3[i] \ll 1) \ggg r_0$	0	2^{-1}
	0	0	0	$\alpha_3[i] \ggg r_0$	1

To compute the output differences for G we can obviously proceed in the following way:

$$(\alpha_0, \alpha_1, \alpha_2, \alpha_3) \xrightarrow{G_1} (\beta_0, \beta_1, \beta_2, \beta_3) \xrightarrow{G_2} (\gamma_0, \gamma_1, \gamma_2, \gamma_3)$$

Listing all 1-bit output differences $\gamma_w[j]$ of G on an arbitrary input difference $\alpha_v[i]$ is quite a complex task. Thus we only give an estimation of the maximum number of active bits in the output difference $\gamma := (\gamma_0, \gamma_1, \gamma_2, \gamma_3)$ after one application of G. Table 7.2 lists the results, which were also confirmed experimentally.

Table 7.2: Maximum Hamming weight of an output difference γ_w after one application of G on an input difference $\alpha_v[i]$

	$a_0[i]$	$a_1[i]$	$a_2[i]$	$a_3[i]$
max. hw(γ_w)	102	115	34	39

7.2.3 Simple Differentials

In this section we show how to construct a class of high probability differentials for the round function F and a small number of iterations F^l. We will focus here on NORX64, but similar considerations should hold for NORX32.

We first consider a simple attack model where the initial state is assumed chosen uniformly at random and where one seeks differences in the initial state that give biased differences in the state obtained after a small number of iterations of F. High-probability truncated differentials wherein the output difference concerns only a small subset of bits (e.g., a single bit) are sufficient to distinguish a (reduced-round) permutation from a random one, and are easier to find for an adversary than differentials on all b bits of the state. To find such differentials we start from our previous analysis of G and extend it to F^l. First, we observe that it is easy to track differences during the first few steps, and in particular to find probability-1 (truncated) differential characteristics for a small number of iterations of F.

For example, by setting the active bit in the MSB of one of the input words a, b, c or d of G a lot of differences are erased due to the shift operation $\ll 1$, as already noted previously. Concretely, using two input words with the input difference $\alpha_0[63]$, i.e. the MSB being active in input a, six of the twelve output differences of G_1 (!) are erased by $\ll 1$ (cf. Table 7.1). As the shift is applied to the non-linear part of G a lot of non-probability-1 differences are deleted, while mainly probability-1 differences remain. Additionally, if distinct input words have active bits in the same positions it leads to further cancellations. Using this simple strategy we found three notable differentials for G of high probability and with low weight output differences:

$$\begin{aligned}
&(8000000000000000, 8000000000000000, 8000000000000000, 0000000000000000) \xrightarrow[1]{G} \\
&(0000000000000000, 0000000000000001, 8000000000000000, 0000000000000000) \\
&(0000000000000000, 8000000000000000, 8000000000000000, 8000000000000000) \xrightarrow[2^{-1}]{G} \\
&(8000000000000000, 0000000001000001, 8000000000800000, 0000000000800000) \\
&(0000000000000000, 8000000000000000, 8000000000000000, 8000000000000000) \xrightarrow[2^{-1}]{G} \\
&(8000000000000000, 0000000003000001, 8000000001800000, 0000000000800000)
\end{aligned}$$

Applying those differentials to F has the effect that the diffusion of the state is delayed by one step. Note that input differences with other combinations of active MSBs lead to similar output differences, but none with a lower or equal Hamming weight as the above. Using the first of the above differentials, we were able to easily derive a truncated differential over 3 steps (i.e. $F^{1.5}$), which has probability 1. This truncated differential can be used to construct an impossible differential over 3.5 rounds for the 64-bit version of F, which is shown in the next section.

We expect that advanced search techniques are able to find better differential distinguishers for a higher number of iterations of F, such that the sparse difference occurs at a later step than in the first. Nevertheless we expect that it is not possible to find differential distinguishers for as much rounds as specified for our instances, see Table 3.1, taking into account the reduced freedom an adversary has, when attacking the initialisation or round permutation.

7.2.4 Impossible Differentials

Cryptanalysis using impossible differentials was introduced in 1998 by Knudsen to attack the block cipher DEAL [42]. Later it was extended by Biham et al. in order to attack the block ciphers Skipjack [25] and IDEA [26]. The latter introduces the so called *miss-in-the-middle* technique. This approach combines two probability 1 differentials, one in forward and one in backward direction which exhibit a conflict when both directions are joined. This strategy leads to an impossible event, i.e. an incident having probability 0, and can be used to construct distinguishers or even mount key recovery attacks.

In our case we construct an impossible differential over 3.5 rounds of the 64-bit version of F, namely 3 steps in forward and 4 steps in backward direction, using the miss-in-the-middle approach from above. An illustration[2] of the used differentials and the resulting conflict is given in Fig. 7.1. A $*$ denotes a partially known and a ? an unknown entry. Our analysis

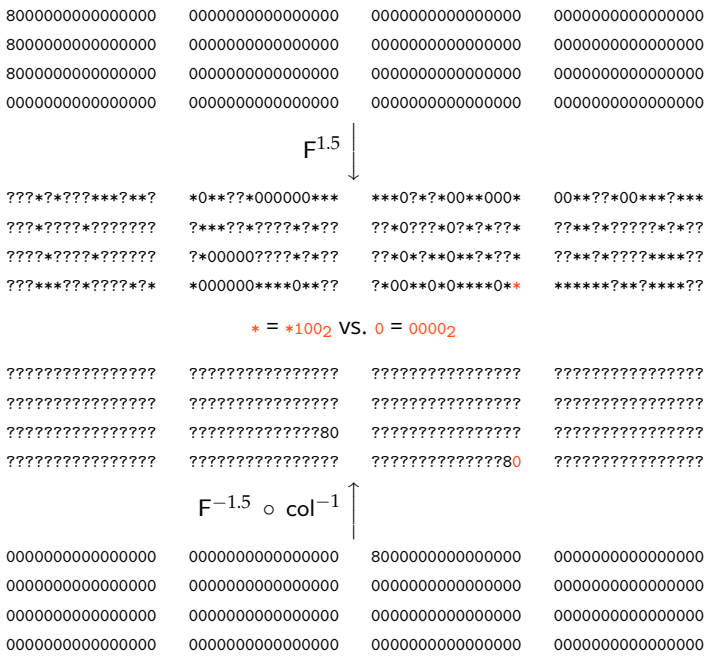

Figure 7.1: An impossible differential over 3.5 rounds of 64-bit F

shows that the conflict occurs in the 2nd bit of the 14th word. In forward direction this bit has always[3] value 1 whereas in backward direction it has always value 0. Note that there are many more impossible differentials of the above type starting from comparable input differences in

[2] We refer to Fig. C.2 in the appendix for the bit representation of the output differences.
[3] The impossible differential was validated empirically in about 2^{32} runs.

forward and backward direction. Nevertheless, using such a simple approach, we were not able to construct impossible differentials stretching over more than 3.5 rounds.

Those impossible differentials cannot be used to attack (round-reduced) NORX, due to the following reasons:

1. The state setup used during initialisation prevents an attacker from setting the required input difference in forward direction. It would be necessary to set differences in the first three consecutive MSBs of a column, which is impossible, as every column is initialised with at least two constant values (see initialisein Fig. 3.6). Thus, even in the *related-key attack model* it is not possible to exploit this class of impossible differentials.
2. Under the assumption that an attacker is *nonce-respecting* [49] and that F^l provides maximum security for $l \geq 4$, two states being set up with two different nonces lead to two distinct internal states after the initialisation phase. Therefore an attacker does not know how to set header blocks to construct the required input difference in forward direction. The same holds for the payload phase. In summary the impossible differential cannot be exploited at a later phase of the algorithm either.

7.3 Algebraic Cryptanalysis

Algebraic attacks on cryptographic algorithms discussed in the literature [7, 9, 30, 33] target ciphers whose internal state is mainly updated in a linear way and thus exploit a low algebraic degree of the attacked primitive. However, this is not the case for NORX, where the b inner state bits are updated in a strongly non-linear fashion. In the following we briefly discuss the non-linearity properties of NORX, demonstrating why it is unlikely that algebraic attacks can be successfully mounted against the cipher.

A convenient way of representing a Boolean function is through its *Algebraic Normal Form* (ANF). Given a Boolean function $f : \{0,1\}^n \longrightarrow \{0,1\}$, the ANF representing f is a multivariate polynomial, i.e. a sum of monomials in n input variables. Both a large number of monomials in the ANF and a good distribution of their degrees are important properties of non-linear building blocks in ciphers.

We constructed the ANF of G and measured the degree of every of the $4w$ polynomials and the distribution of the monomials. Table 7.3 reports the number of polynomials per degree for the 32- and 64-bit versions, as well as information on the distribution of monomials.

Table 7.3: Number of polynomials by degree, and number of monomials by polynomial

	#polynomials by degree						#monomials			
	3	4	5	6	7	8	min	max	avg	med
64-bit	2	6	122	2	8	116	12	489	253	49.5
32-bit	2	6	58	2	8	52	12	489	242	49.5

In both cases most polynomials have degree 5 or 8 and merely 2 have degree 3. Multiplying each of the above values by 4 gives the distribution of degrees for the ANF of the whole state

after one column or diagonal step. Due to memory constraints, we were unable to construct[4] the ANF for a single full round F, neither for the 64-bit nor for the 32-bit version. In summary, this shows that the state of NORX is updated in a strongly non-linear fashion. Due to the rapid growth of the degree and the huge state size of NORX we believe that it is unlikely that algebraic cryptanalysis can be used to successfully mount an attack on the AEAD scheme.

7.4 Other Attacks

In this section we briefly review other kinds of attacks that may be used against NORX.

7.4.1 Fixed Points

The G permutation and thus any iteration of the round function F have a trivial distinguisher: the fixed points $G(0) = 0$ and $F^l(0) = 0$. Nevertheless it seems hard to exploit this property, as hitting the all-zero state is as hard as hitting any other arbitrary state. Thus the ability to hit a predefined state implies the ability to recover the key, which is equivalent to completely breaking NORX. Therefore the zero-to-zero point is no significant threat to the security of NORX.

Furthermore, we used the constraint solver STP [34] to prove that there are no further fixed points. For NORX32 the solver was able to show that this is indeed the case, but for NORX64 the proof is a lot more complex. Even after over 1000 hours, STP was unable to finish its computation with a positive or negative result. We find it unlikely that there are any other fixed points in NORX64 besides the zero-to-zero point.

7.4.2 Slide Attacks

Slide attacks try to exploit the symmetry in a primitive that consists of the iteration of a number of identical rounds. They were introduced by Biryukov et al. [28, 29] to cryptanalyse block ciphers. Later they were also extended to stream ciphers [47] and hash functions [35]. To protect sponge constructions against slide attacks two simple countermeasures can be found in the literature:

1. In [35] it is proposed to add a non-zero constant to the state just before applying the permutation.
2. In [46] it is recommended to use a message padding, which ensures that the last processed data block is different from the all-zero message.

The duplex constructions is derived from sponge functions, hence the above countermeasures should hold for the former, too, and thus for NORX. Both defensive mechanisms are already integrated into NORX: the domain separation constants are added to the state just before the permutation F^l is applied and the multi-rate padding ensures that the last processed data block is different from the all-zero block. Hence, slide attacks should pose little to no threat to NORX.

[4] Using SAGE [51] on a workstation with 64 GiB RAM.

7.4.3 Rotational Cryptanalysis

Rotational cryptanalysis was introduced by Khovratovich and Nikolić in [40] to analyse ARX based primitives. The idea is to track the propagation of rotational relations through a cryptographic transformation. Once rotation-invariant behaviour is detected, it can be used to construct distinguishers, mount key recovery attacks and so on. Rotational cryptanalysis was successfully applied to several simplified cryptographic primitives including Skein [41] and Keccak [44].

NORX includes several defense mechanisms to increase the difficulty of finding exploitable rotation-invariant behaviour:

1. During state setup 8 out of 16 words are initialised with asymmetric constants, which impedes the occurrence of rotation-invariant behaviour and limits the freedom of an attacker. A similar approach is also used in Salsa20 [16].
2. The non-linear operation of NORX contains a non rotation-invariant bit-shift $\ll 1$.
3. NORX is based on the duplex construction, which prevents an attacker from modifying the complete internal state at a given time. He is only able to influence the rate bits, i.e. at most $r = 12w$ bits of the state, and has to "guess" the other $4w$ bits in order to mount an attack.

8 Intellectual Property

We, the designers of NORX, do hereby declare that

- NORX is free for everyone to use;
- We are not aware of any patent or patent application that may cover the practice of the NORX algorithm;
- We have not filed any patent application related to the NORX algorithm.

If any of this information changes, the submitter/submitters will promptly (and within at most one month) announce these changes on the `crypto-competitions` mailing list.

9 Consent

The submitter/submitters hereby consent to all decisions of the CAESAR selection committee regarding the selection or non-selection of this submission as a second-round candidate, a third-round candidate, a finalist, a member of the final portfolio, or any other designation provided by the committee. The submitter/submitters understand that the committee will not comment on the algorithms, except that for each selected algorithm the committee will simply cite the previously published analyses that led to the selection of the algorithm. The submitter/submitters understand that the selection of some algorithms is not a negative comment regarding other algorithms, and that an excellent algorithm might fail to be selected simply because not enough analysis was available at the time of the committee decision. The submitter/submitters acknowledge that the committee decisions reflect the collective expert judgments of the committee members and are not subject to appeal. The submitter/-submitters understand that if they disagree with published analyses then they are expected to promptly and publicly respond to those analyses, not to wait for subsequent committee decisions. The submitter/submitters understand that this statement is required as a condition of consideration of this submission by the CAESAR selection committee.

10 Acknowledgements

The authors thank Frank K. Gürkaynak, Mauro Salomon, Tibor Keresztfalvi and Christoph Keller for implementing NORX in hardware and for giving insightful feedback from their hardware evaluation.

Moreover, the authors would like to thank Alexander Peslyak (Solar Designer), for giving them access to one of his Haswell machines, so that they could test their AVX2 implementations of NORX.

Bibliography

[1] eSTREAM - the ECRYPT Stream Cipher Project, 2004–2008. http://www.ecrypt.eu.org/stream.

[2] SHA-3 Competition, 2007–2012. http://csrc.nist.gov/groups/ST/hash/sha-3/Round3/index.html.

[3] CAESAR — Competition for Authenticated Encryption: Security, Applicability, and Robustness, 2014. http://competitions.cr.yp.to/caesar.html.

[4] Official website of NORX, 2014. https://www.norx.io.

[5] Martin Albrecht, Nicolas T. Courtois, Daniel Hulme, and Guangyan Song. Bit-Slice Implementation of PRESENT in Pure Standard C, v1.5, 2011. Opensource code available at https://bitbucket.org/malb/research-snippets/src.

[6] Elena Andreeva, Joan Daemen, Bart Mennink, and Gilles Van Assche. Security of keyed sponge constructions using a modular proof approach. In Gregor Leander, editor, *Fast Software Encryption - 22nd International Workshop, FSE 2015, Istanbul, Turkey, March 8-11, 2015, Revised Selected Papers*, volume 9054 of *Lecture Notes in Computer Science*, pages 364–384. Springer, 2015.

[7] Frederik Armknecht. On the Existence of Low-Degree Equations for Algebraic Attacks. Cryptology ePrint Archive, Report 2004/185, 2004. http://eprint.iacr.org/2004/185.

[8] Jörg Arndt. *Matters Computational: Ideas, Algorithms, Source Code*. Springer-Verlag New York, Inc., New York, NY, USA, 1st edition, 2010. http://jjj.de/fxt/fxtpage.html#fxtbook.

[9] Jean-Philippe Aumasson, Itai Dinur, Luca Henzen, Willi Meier, and Adi Shamir. Efficient FPGA Implementations of High-Dimensional Cube Testers on the Stream Cipher Grain-128. Cryptology ePrint Archive, Report 2009/218.

[10] Jean-Philippe Aumasson, Simon Fischer, Shahram Khazaei, Willi Meier, and Christian Rechberger. New Features of Latin Dances: Analysis of Salsa, ChaCha and Rumba. In Kaisa Nyberg, editor, *FSE 2008*, volume 5086 of *LNCS*, pages 470–488. Springer, 2008.

[11] Jean-Philippe Aumasson, Luca Henzen, Willi Meier, and Raphael C.-W. Phan. SHA-3 Proposal BLAKE. In *NIST SHA-3 Proposal*, 2010. https://131002.net/blake.

[12] Jean-Philippe Aumasson, Philipp Jovanovic, and Samuel Neves. NORX8 and NORX16: Authenticated Encryption for Low-End Systems. Cryptology ePrint Archive, Report 2015/1154, 2015. http://eprint.iacr.org/2015/1154.

[13] Jean-Philippe Aumasson, Samuel Neves, Zooko Wilcox-O'Hearn, and Christian Winnerlein. BLAKE2: Simpler, Smaller, Fast as MD5. In Michael Jacobson, Michael Locasto, Payman Mohassel, and Reihaneh Safavi-Naini, editors, *ACNS 2013*, volume 7954 of *LNCS*, pages 119–135. Springer, 2013.

[14] Michael Beeler, R. William Gosper, and Richard Schroeppel. HAKMEM. Artificial Intelligence Memo 239, Massachusetts Institute of Technology, February 1972. http://dspace.mit.edu/handle/1721.1/6086.

[15] Daniel J. Bernstein. Cache-Timing Attacks on AES, 2005. http://cr.yp.to/antiforgery/cachetiming-20050414.pdf.

[16] Daniel J. Bernstein. Salsa20 Security, 2005. http://cr.yp.to/snuffle/security.pdf.

[17] Daniel J. Bernstein. ChaCha, a Variant of Salsa20. In *Workshop Record of SASC 2008: The State of the Art of Stream Ciphers*, 2008. http://cr.yp.to/chacha.html.

[18] Daniel J. Bernstein. The Salsa20 Family of Stream Ciphers. In Matthew Robshaw and Olivier Billet, editors, *New Stream Cipher Designs*, volume 4986 of *LNCS*, pages 84–97. Springer, 2008.

[19] Guido Bertoni, Joan Daemen, Michael Peeters, and Gilles Van Assche. On the Security of Keyed Sponge Constructions. Presented at SKEW 2011, 16–17 February 2011, Lyngby, Denmark, http://sponge.noekeon.org/SpongeKeyed.pdf.

[20] Guido Bertoni, Joan Daemen, Michael Peeters, and Gilles Van Assche. Permutation-based Encryption, Authentication and Authenticated Encryption. Presented at DIAC 2012, 05–06 July 2012, Stockholm, Sweden.

[21] Guido Bertoni, Joan Daemen, Michael Peeters, and Gilles Van Assche. On the Indifferentiability of the Sponge Construction. In Nigel Smart, editor, *EUROCRYPT 2008*, volume 4965 of *LNCS*, pages 181–197. Springer, Heidelberg, 2008.

[22] Guido Bertoni, Joan Daemen, Michael Peeters, and Gilles Van Assche. Cryptographic Sponge Functions, January 2011. http://sponge.noekeon.org/CSF-0.1.pdf.

[23] Guido Bertoni, Joan Daemen, Michael Peeters, and Gilles Van Assche. Duplexing the Sponge: Single-Pass Authenticated Encryption and Other Applications. In A. Miri and S. Vaudenay, editors, *SAC 2011*, volume 7118 of *LNCS*, pages 320–337. Springer, 2011.

[24] Guido Bertoni, Joan Daemen, Michael Peeters, Gilles Van Assche, and Ronny van Keer. KECCAK implementation overview, May 2012. http://keccak.noekeon.org.

[25] Eli Biham, Alex Biryukov, and Adi Shamir. Cryptanalysis of Skipjack Reduced to 31 Rounds Using Impossible Differentials. In Jacques Stern, editor, *EUROCRYPT 1999*, volume 1592 of *LNCS*, pages 12–23. Springer, Heidelberg, 1999.

[26] Eli Biham, Alex Biryukov, and Adi Shamir. Miss in the Middle Attacks on IDEA and Khufu. In Lars Knudsen, editor, *FSE 1999*, volume 1636 of *LNCS*, pages 124–138. Springer, Heidelberg, 1999.

[27] Alex Biryukov and Dmitry Khovratovich. PPAE: Parallelizable Permutation-based Authenticated Encryption. Presented at DIAC 2013, 11–13 August 2013, Chicago, USA, http://2013.diac.cr.yp.to/slides/khovratovich.pdf.

[28] Alex Biryukov and David Wagner. Slide Attacks. In Lars Knudsen, editor, *FSE 1999*, volume 1636 of *LNCS*, pages 245–259. Springer, Heidelberg, 1999.

[29] Alex Biryukov and David Wagner. Advanced Slide Attacks. In Bart Preneel, editor, *EUROCRYPT 2000*, volume 1807 of *LNCS*, pages 589–606. Springer, Heidelberg, 2000.

[30] Nicolas T. Courtois. Algebraic Attacks on Combiners with Memory and Several Outputs. In Choon sik Park and Seongtaek Chee, editors, *Information Security and Cryptology (ICISC)*, volume 3506 of *LNCS*, pages 3–20. Springer, Heidelberg, 2004. http://eprint.iacr.org/2003/125.

[31] Nicolas T. Courtois, Daniel Hulme, and Theodosis Mourouzis. Solving Circuit Optimisation Problems in Cryptography and Cryptanalysis. In *SHARCS*, 2012. http://eprint.iacr.org/2011/475.

[32] Joan Daemen and Vincent Rijmen. The Advanced Encryption Standard, 2001. http://csrc.nist.gov/publications/fips/fips197/fips-197.pdf.

[33] Itai Dinur and Adi Shamir. Cube Attacks on Tweakable Black Box Polynomials. In Antoine Joux, editor, *EUROCRYPT 2009*, volume 5479 of *LNCS*, pages 278–299. Springer, Heidelberg, 2009.

[34] Vijay Ganesh, Ryan Govostes, Khoo Yit Phang, Mate Soos, and Ed Schwartz. *STP — A Simple Theorem Prover*, 2006–2013. http://stp.github.io/stp.

[35] Michael Gorski, Stefan Lucks, and Thomas Peyrin. Slide Attacks on a Class of Hash Functions. In Josef Pieprzyk, editor, *ASIACRYPT 2008*, volume 5350 of *LNCS*, pages 143–160. Springer, Heidelberg, 2008.

[36] Robert Granger, Philipp Jovanovic, Bart Mennink, and Samuel Neves. Improved Masking for Tweakable Blockciphers with Applications to Authenticated Encryption. In Marc Fischlin and Jean-Sébastien Coron, editors, *Advances in Cryptology – EUROCRYPT 2016*, volume 9665 of *LNCS*, pages 263–293. Springer Berlin Heidelberg, 2016.

[37] Luca Henzen, Flavio Carbognani, Norbert Felber, and Wolfgang Fichtner. VLSI Hardware Evaluation of the Stream Ciphers Salsa20 and ChaCha, and the Compression Function Rumba. In *2nd International Conference on Signals, Circuits and Systems 2008*, pages 1–5. IEEE, 2008.

[38] Antoine Joux. Authentication Failures in NIST Version of GCM, 2006. http://csrc.nist.gov/groups/ST/toolkit/BCM/documents/Joux_comments.pdf.

[39] Philipp Jovanovic, Atul Luykx, and Bart Mennink. Beyond $2^{c/2}$ Security in Sponge-Based Authenticated Encryption Modes. Cryptology ePrint Archive, Report 2014/373, 2014. http://eprint.iacr.org/2014/373.

[40] Dmitry Khovratovich and Ivica Nikolić. Rotational Cryptanalysis of ARX. In Seokhie Hong and Tetsu Iwata, editors, *FSE 2010*, volume 6147 of *LNCS*, pages 333–346. Springer, 2010.

[41] Dmitry Khovratovich, Ivica Nikolić, and Christian Rechberger. Rotational Rebound Attacks on Reduced Skein. In Masayuki Abe, editor, *ASIACRYPT*, volume 6477 of *LNCS*, pages 1–19. Springer, Heidelberg, 2010.

[42] Lars R. Knudsen. DEAL — A 128-bit Block Cipher. In *NIST AES Proposal*, 1998.

[43] Donald E. Knuth. *The Art of Computer Programming, Volume 4A: Combinatorial Algorithms, Part 1*, volume 4A. Addison-Wesley, Upper Saddle River, New Jersey, 2011. http://www-cs-faculty.stanford.edu/~uno/taocp.html.

[44] Pawel Morawiecki, Josef Pieprzyk, and Marian Srebrny. Rotational Cryptanalysis of Round-Reduced KECCAK. Cryptology ePrint Archive, Report 2012/546, 2012. http://eprint.iacr.org/2012/546.

[45] National Institute of Standards and Technology. Recommendation for Block Cipher Modes of Operation: Galois/Counter Mode (GCM) and GMAC, 2007. http://csrc.nist.gov/publications/nistpubs/800-38D/SP-800-38D.pdf.

[46] Thomas Peyrin. Security Analysis of Extended Sponge Functions. Presented at the ECRYPT Workshop Hash Functions in Cryptology: Theory and Practice, Leiden, The Netherlands, June 4th 2008, http://www.lorentzcenter.nl/lc/web/2008/309/presentations/Peyrin.pdf.

[47] Deike Priemuth-Schmid and Alex Biryukov. Slid Pairs in Salsa20 and Trivium. In Dipanwita Roy Chowdhury, Vincent Rijmen, and Abhijit Das, editors, *Indocrypt*, volume 5365 of *LNCS*, pages 1–14. Springer, Heidelberg, 2008. http://eprint.iacr.org/2008/405.

[48] Phillip Rogaway. Authenticated-Encryption with Associated-Data. In *ACM Conference on Computer and Communications Security (CCS'02)*, pages 98–107. ACM press, 2002.

[49] Phillip Rogaway. Nonce-Based Symmetric Encryption. In Bimal Roy and Willi Meier, editors, *FSE 2004*, volume 3017 of *LNCS*, pages 348–358. Springer, Heidelberg, 2004.

[50] Zhenqing Shi, Bin Zhang, Dengguo Feng, and Wenling Wu. Improved Key Recovery Attacks on Reduced Round Salsa20 and ChaCha. In Taekyoung Kwon, Mun-Kyu Lee, and Daesung Kwon, editors, *ICISC 2012*, volume 7839 of *LNCS*, pages 337–351. Springer, 2012.

[51] W. A. Stein. *Sage Mathematics Software*. The Sage Development Team, 2005–2013. http://sagemath.org.

[52] Yukiyasu Tsunoo, Teruo Saito, Hiroyasu Kubo, Tomoyasu Suzaki, and Hiroki Nakashima. Differential Cryptanalysis of Salsa20/8. In *The State of the Art of Stream Ciphers (SASC)*, 2007.

A Test Vectors

A.1 Traces for F

To verify correctness of an F implementation we suggest to use the derivation of the NORX initialisation constants:

$$(u_0, \ldots, u_{15}) = \mathsf{F}^2(0, \ldots, 15)$$

The values of $u_0, \ldots, u_{15}$ are given in Table 3.4.

A.2 Full AEAD Computations

We assume that the following input data is given for NORX32 and NORX64, respectively:

NORX32		
type	data	#bytes
K :	00 01 ... 0E 0F	16
N :	20 21 ... 2E 2F	16
A :	00 01 ... 7E 7F	128
M :	00 01 ... 7E 7F	128
Z :	00 01 ... 7E 7F	128

NORX64		
type	data	#bytes
K :	00 01 ... 1E 1F	32
N :	20 21 ... 3E 3F	32
A :	00 01 ... 7E 7F	128
M :	00 01 ... 7E 7F	128
Z :	00 01 ... 7E 7F	128

The test vectors for the five proposed instances are given on the following pages. Intermediate values are snapshots of the state after initialisation (Fig. 3.5, line 1), after header processing (Fig. 3.5, line 2), after message encryption (Fig. 3.5, line 5), after trailer processing (Fig. 3.5, line 6), and after finalisation (Fig. 3.5, line 7).
For more tests we refer to the NORX software package available at [4]. The above test vectors were generated using `utils/debug.c`. See also `utils/check.c`, `utils/genkat.c`, and the various `kat.h` files for a more comprehensive test-framework.

NORX32-4-1

State after initialisation (Fig. 3.5, line 1)

```
7DD54975  C374FFC8  1DF66F83  08CEF7E9
CA5295E8  8E1E6324  538244DA  3091DC5D
5288E900  EDDAFB81  1A345AE0  933EC3AB
BED76EB5  8B64D948  A59BD31B  6BBBD034
```

State after header processing (Fig. 3.5, line 2)

```
2DFDA46B  956D99E2  DE62A45D  59A4AD56
F9A5411A  759C0658  45CF1EA3  A9515464
60CCA3C1  A29F076D  FAA12E42  EA22ED90
7D10BA9D  407E2C5B  97DC4FA4  80401262
```

State after message encryption (Fig. 3.5, line 5)

```
9769850C  41240274  A264E03A  B808815A
9285A6D3  8665C774  ED279CE2  9571FB11
F39624ED  3DCE8561  81879FF2  45B5E234
10D6694E  AFF8A691  9991AECE  BFFA4576
```

State after trailer processing (Fig. 3.5, line 6)

```
BBAB2C4A  42BF34A5  3AD53DFA  AF184F4D
66A33356  481AAE25  471E110F  9FBC7740
33A4CBDB  5CA77A41  ABCDF216  1A213FE2
353816EC  8EFF5ABE  3FB2298B  E4A9EC82
```

State after finalisation (Fig. 3.5, line 7)

```
97537D63  63AC168C  6CEF0F5B  EC0114E9
D6A022EC  FF4395E0  4F29B8B5  B8CC8998
D92C5C49  74BA3CEF  964EEDD3  23DF1024
BCE454D5  89B75B6B  EA597754  47CFFFCD
```

Ciphertext and authentication tag

```
C :  6C E9 4C B5 48 B2 0F ED 7B 68 C6 AC 60 AC 4C B5
     EB B1 F0 9A EC 5A 75 0E CF 50 EC 0E 64 93 8B F2
     40 17 A4 FF 06 84 F8 08 A6 7C 19 6C 31 A0 AF 12
     56 9B E5 F7 C5 6A D3 BC AC 88 DA 36 86 57 5F 93
     43 96 8D A2 20 77 EE CC E7 D6 63 17 49 08 A3 F7
     3C 9E 9A C1 49 B5 CE 6B E6 9C 9E 31 7C D7 E7 E8
     0C 85 69 97 74 02 24 41 3A E0 64 A2 5A 81 08 B8
     D3 A6 85 92 74 C7 65 86 E2 9C 27 ED 11 FB 71 95

T :  D5 54 E4 BC 6B 5B B7 89 54 77 59 EA CD FF CF 47
```

NORX32-6-1

State after initialisation (Fig. 3.5, line 1)

```
6AA9881B  E39461D4  72E17E31  D42E766A
773C7827  E93C085F  08201969  1E455C9D
016FDB2E  AB3DE913  69289A33  A5CE3028
C18142E4  71F18D99  CE8FD48D  1C141F99
```

State after header processing (Fig. 3.5, line 2)

```
2C6CF747  221AA528  E21D7954  70765726
48BED941  5DEC2685  628316CC  D117C238
A86F0D92  E606ECAF  60B283A8  5B654A7B
C9ACD7AC  89250D1C  F4DD1B13  84075320
```

State after message encryption (Fig. 3.5, line 5)

```
F357C5AC  11E7DCD2  480A43A5  A263168C
48816707  6B3AC79A  39FE6AFB  979FE76A
84CF2D94  357F0F56  883EDBA4  8D7BA338
85F97717  187F6B85  04ED8B2E  A5671D54
```

State after trailer processing (Fig. 3.5, line 6)

```
72FC85AC  5CE4F0D1  F67A024C  A2BA77E4
8EF884EF  7864FC01  24984D54  C00489BA
727789E5  EE8B24FE  546C3F3A  0D32D7B4
843B41CB  FC32B163  0E659E5A  ECAC8817
```

State after finalisation (Fig. 3.5, line 7)

```
8D3E098F  4D470E23  219829EE  96A87C5B
189B7D32  03413845  0D513C08  436ECF50
54ED6AE8  88BB869A  78A96727  13A396EF
8F1AB1B3  B1F1949A  E95318AC  4A26434C
```

Ciphertext and authentication tag

```
C :  20 9B 0B 2A FE 36 2A 83 3B B1 8A CF 03 E1 D0 C2
     7C 69 47 52 66 79 47 FC 73 8C 0E 40 E3 D5 97 C2
     2D 74 E9 06 E8 C4 73 AD F0 DB 63 61 D3 97 41 C4
     26 0F B3 D3 9F 84 22 A3 CF DF 93 0D 2D 17 75 EB
     3F 97 0E 52 95 23 07 C9 AA 07 3F C5 E1 19 BA DF
     B2 FF 00 9E 69 7C 8E 85 61 4F 44 78 C5 7B D2 B4
     AC C5 57 F3 D2 DC E7 11 A5 43 0A 48 8C 16 63 A2
     07 67 81 48 9A C7 3A 6B FB 6A FE 39 6A E7 9F 97

T :  B3 B1 1A 8F 9A 94 F1 B1 AC 18 53 E9 4C 43 26 4A
```

NORX64-4-1

State after initialisation (Fig. 3.5, line 1)

```
ED1C05E4E034B18B  A98C191C6015FA6D  288C3313ACF5E185  94E37DCA8C2B520F
841D5FBE319581DE  6BA9AE4E997C10DF  9ACC31C63498AAB8  BC4F4AA085B8FAD9
24A958D377B4FBBF  8DDB5DC488A3A710  7F776980AAA321EF  4D4C321A44EE66D9
C6439632673FBDC2  950244CDFEAEA45E  EB8B0AFF16BEDBE0  68A7A80B2838111F
```

State after header processing (Fig. 3.5, line 2)

```
07D9A7A131D4D6E0  5B60B0B0847E0416  57F3CB734EC314B3  F9CFDC4B605A6CCC
5E3F25A15BF57819  3501EA9EDDF5CC6C  69BAAF08D99F96C2  CF86E9721020F64E
3352D33F5677CBC4  331C29A0674FEF14  CB74AFFFA9BD69D9  5810E32F833F0370
44C3442263959E68  522FC8BBFE971C48  4EC92E818EA35AD3  BB223CBC51462414
```

State after message encryption (Fig. 3.5, line 5)

```
A4461CDB6586E74B  BDDF7652BF4F1AB0  DCF86684B8BFEB30  D870D0D016787A89
C5DC8F2CC92A2D60  404DE2D5457A5178  8A2475887B1ABF74  AD5BEFE2F99B111F
D258C60C34FC528A  69C0DA88A6C5CD25  3328D007C5C35CC3  3744B8E898EC83DD
70AB4D51F1570C40  5E3331A6663C18EE  BA01BA7CFDF2C4BF  36FA274968BF8B0A
```

State after trailer processing (Fig. 3.5, line 6)

```
B1B64376441A2AB0  2F5BE2578863D5EC  66F953E878E37E6B  EEE236C48DEDFFEE
6778F573276FDF5F  E3C3E60EDC6DB52D  B0AAEFFFF4764978  2A0F46F39ED63CF1
A9C34DCB7057873C  594CC2D6E926D398  D85F144A45107F10  EE584A7C1D80E6D3
7B763E9FCBB1F9D3  9A55D3CAC654F97A  9308DF76F6D7995E  6D9E59C21CC59E3B
```

State after finalisation (Fig. 3.5, line 7)

```
45D70450C188B282  44CB44A8ACC7D823  6CF99985A76DD706  F76D93B792F90C83
BCB8EC0B3370F727  011728D02D035E19  CC7972F3E89E595A  A75510060F10F800
D3314C7CDF7C4C99  52A16E0D4BD61F3C  4EA70ACD1A1F1D3A  B56927EF60BB58D4
7623A30533FAF2D1  3F3089C9D1613AE2  E4175BA55A93BDBF  8E4073C4334725E7
```

Ciphertext and authentication tag

```
C :  C0 81 6E 50 8A E4 A0 50 0B 93 38 7B BB AB C2 41
     AC 42 38 7E F5 E8 BF 0E C3 82 6C ED E1 66 A1 D5
     CA A3 E8 D6 2C D6 41 B3 FA F2 AA 2A DD E3 E5 ED
     0A 13 BD 8B 96 D5 F0 FB 7F E3 9C A7 80 95 31 75
     E2 45 BC 3E 53 4B 80 0E 96 46 77 1F 13 EA 40 85
     CB 3E 26 7F 10 6F 5F 17 A0 64 FF 23 4A 02 7C 64
     4B E7 86 65 DB 1C 46 A4 B0 1A 4F BF 52 76 DF BD
     30 EB BF B8 84 66 F8 DC 89 7A 78 16 D0 D0 70 D8

T :  D1 F2 FA 33 05 A3 23 76 E2 3A 61 D1 C9 89 30 3F
     BF BD 93 5A A5 5B 17 E4 E7 25 47 33 C4 73 40 8E
```

NORX64-6-1

State after initialisation (Fig. 3.5, line 1)

```
FBE74F8B8F87B637  D8CF91E950A5E7A3  BF8B7D6EE6F66D23  986681BC560A954D
E7EC41745D71BB3E  4EC9F3E9DDDD549A  04D80F76956ABE4E  40357E0A23B3A8DD
45CC935DD4C35559  6665AC3AE53EE95C  92652F352B2741AF  A793FACE1833962F
8E98DF92B5CC6486  811F60A1A563FCA3  C5208A2F80A021D7  E356DE42D6A7FFDF
```

State after header processing (Fig. 3.5, line 2)

```
86A6D4B372772D8E  442DD22404C14516  F545B5F6F138F2C8  BA8F26D2915DB59B
7E47657EFDBD4A19  227101E45166C034  57788639F923F1B7  F5D4510A6E2E43CE
95FE28CEDFA322B5  ACA45A91CDA378E0  AA3B58DAEB4C0C89  3CCE3333A7A6D105
D4E9DE665C5F55A5  360425F706D0732C  1F0AC3505DBDD8A4  0609356E2891D9FB
```

State after message encryption (Fig. 3.5, line 5)

```
6F50D4782B87F3C4  BED2B8E33C65DB40  B756F47FE9F9A2A7  232FE227928CDBF0
163AA5E6AA2B8F1F  BE9B93658DFEA66C  37AFC3D1C18E7B4E  1E405E25DC276BF0
0942EC8E4D5CF8A8  52642976ABF1E083  D92BABC7DFDCC66D  1A58A4388C3EC775
6C2C708571473F82  5D5415187A314FA7  E6438637D8DF9BBE  81AF89907BEEFBEF
```

State after trailer processing (Fig. 3.5, line 6)

```
803AADDC8C8A6C3C  240E45C8DDDF922F  46612516E3B59292  E8151BED6699777C
94EA2F311F5FD215  AE34A389A9E4D09A  2E1FC599862D05CE  6EF55D773D4E3841
37608855B6501422  790724A43D36A080  6D677BE24077921B  0B40C69E92D8D68E
0A31B5155E7C5E32  727A6A113C0DD6D6  7AAC1CFDD18EB31D  DF4BC15545BA3CB2
```

State after finalisation (Fig. 3.5, line 7)

```
A2F7E3AFA01DCFFB  14D342E71FE41B15  8B8AC118983D0B11  ACC01CEAF112614B
8ABC00E2F599919B  991AA67F0A29616B  98A27A56EF31B1AE  F294074232A3BA81
F1470D44E094DB73  F93ED9FD0F5395B4  02E9988E4A8272B2  CE07BC4DA6DD7D99
DBBE9152280DD1A0  38E20F7EC4BD7C7B  ABBC57F0C55BF55B  2C9BD283D0CC572C
```

Ciphertext and authentication tag

```
C :  50 CE 69 2C 19 CB 91 02 C6 12 96 6F 0F 62 6B 62
     96 DE 89 27 1C 98 29 10 AA C1 C3 55 52 2E 8F A7
     13 03 F8 D5 C9 DE 39 04 84 BA 91 A9 94 CF F9 1B
     F7 15 D6 CB 22 CC 00 F3 64 02 10 03 17 19 61 68
     72 39 DD 94 53 02 9B 87 85 9C 10 93 21 13 59 40
     BC 1B C8 1A 55 A9 51 C7 1B 29 42 FF DE BF 8D 13
     C4 F3 87 2B 78 D4 50 6F 40 DB 65 3C E3 B8 D2 BE
     A7 A2 F9 E9 7F F4 56 B7 F0 DB 8C 92 27 E2 2F 23

T :  A0 D1 0D 28 52 91 BE DB 7B 7C BD C4 7E 0F E2 38
     5B F5 5B C5 F0 57 BC AB 2C 57 CC D0 83 D2 9B 2C
```

NORX64-4-4

State after initialisation (Fig. 3.5, line 1)

```
09C1FAB275E26554  B88696E71C5F243B  6461DD69FAD686BD  C7E59562326FD541
9A6F849C8C700398  D98F499A54C2D279  89B94070BBFA7D93  B68DA702581F2FFB
4E411341A13B2DF4  E32F270FAFA72967  1E7026F60F7C89FF  D61E9C0926B80488
E5B95502271B6092  08B3C9B59806B885  D4DD9B5E6CB8B11F  E4EB9CCF3A9FDC04
```

State after header processing (Fig. 3.5, line 2)

```
C2E06BC5BFE0984A  7339373DA279C0EA  F70ABAF1783FAAC9  33CC3C2C5E8B7D0A
FB87B19894049DAE  87793479CB131B60  7B53C5803ACFF65F  44C8DF1D14396D3B
EC4024C6C6B74E89  BE2EECAFCB79A73D  00857DB83628CC10  B581AFB4D2F81FEA
612D5D628BAF35B4  1026D7DDDDE32B55  BAA88B60B5C7BB19  8090D73B27BAE22F
```

State after message encryption (Fig. 3.5, line 5)

```
B54F46C2037FE58D  0CB5F144A62F7D5A  AAC797F1084FACE7  465EF4939F3D9E9B
F8CBB14B3CDC1972  E94DB2745C7ABA62  2BAAD57DAF1BCEEF  4D130FB81726FBB6
8CE4926961F7C93C  BD25C2C9AB9F14C2  E1B9FC2C7658F531  D62F84C49F2376C8
EA3D7AEA6006DE4E  9936A13B9BEE9F38  0ACFF68DB00C231E  76A4911E8C2614EE
```

State after trailer processing (Fig. 3.5, line 6)

```
3EE2EB5074C21159  0182D18A98FBAC1C  EBDBAE586D91C714  879432FD0E488B39
2178909DD1991DF9  251ECA5BDC482B35  0C60578C3D60CFFD  DB5E7C8F9E000D45
FCDD16FE2F31D0A3  AA6060E3943E32E3  30BFAD62B34569BA  B9C711F41A85F4C8
ECD25D60681F60D4  FB14C2DCD8764106  77A2E80E1B3AC500  D776DEBF23C4B96D
```

State after finalisation (Fig. 3.5, line 7)

```
7915D8E31EF87114  1C9FDCFF1CCEF122  2F224AC3007FF3B1  9E2C80A3F55230CE
D2135CA3D669A42A  3D7EAA58FB4BF9F6  E5586C8CF476570A  E2A795E869BB2DCF
BE59054CEFE4E0E1  D6054EBB539B0FEA  BCFFE906339A2D7C  CDC60ED1A3147EA9
A70080497E3B6101  99FD8F3F35D5F567  1450DC1F7C057278  AC5CA7B8EB2782CF
```

Ciphertext and authentication tag

```
C :  B6 5A D4 9D 08 12 87 73 03 76 A0 38 F1 32 B2 0C
     33 E5 58 30 20 27 C0 D9 1C 03 0B 9C 7D DA 19 C7
     51 1A 4F 02 5A FD 40 FD A2 95 C9 22 29 FA EA 13
     A6 14 05 36 44 0B EB FC D3 62 72 5D 9E E9 0F 2C
     2A AC 10 6B 5F 49 86 9B 9F E2 2C D9 F1 84 84 FC
     70 C2 22 8C 1D A3 07 21 21 97 2C 2B D9 9A 29 2A
     15 51 52 B1 67 72 3F F7 CD A5 BB A3 DA 09 E3 69
     F2 7B FE 53 88 63 FF 56 18 40 01 28 8C C1 BE EC

T :  01 61 3B 7E 49 80 00 A7 67 F5 D5 35 3F 8F FD 99
     78 72 05 7C 1F DC 50 14 CF 82 27 EB B8 A7 5C AC
```

B Datagrams

Many issues with encryption interoperability are due to ad hoc ways to represent and transport cryptograms and the associated data. For example IVs are sometimes prepended to the ciphertext, sometimes appended, or sent separately. We thus specify datagrams that can be integrated in a protocol stack, encapsulating the ciphertext as a payload. Using a standardized encoding simplifies the transmission of NORX cryptograms across different APIs, and reduces the risk of insecure or suboptimal encodings. We specify two distinct types of datagrams, depending on whether the NORX parameters are fixed or need to be signaled in the datagram header.

B.1 Fixed Parameters

With *fixed parameters* shared by the parties (for example through the application using NORX), there is no need to include the parameters in the *header of the datagram*[1]. The datagram for fixed parameters thus only needs to contain N, A, C, Z, and T, as well as information to parse those elements.

We encode the byte length of A and Z on 16 bits, allowing for headers and trailers of up to 64 KiB, a large enough value for most real applications. The byte length of the encrypted payload is encoded on 32 bits for NORX32 and on 64 bits for NORX64, which translates to a maximum payload size of 4 GiB and 16 EiB, respectively[2]. Similarly to frame check sequences in data link protocols, the tag is added as a *trailer of the datagram* specified. The header, encrypted payload, and trailer of the underlying protocol are viewed as the *payload of the datagram*. The default tag length being a constant value of the NORX instance, it needs not be signalled.

Tables B.1 and B.2 show the fixed-parameters datagrams for NORX32 and NORX64. The length of the datagram header is 44 bytes for NORX64 and 24 bytes for NORX32.

Note that the CAESAR API (as per the final call, see [3]) receives the nonce and the associated data in two separate buffers, but the tag is included in the ciphertext buffer.

B.2 Variable Parameters

With *variable parameters*, the datagram needs to signal the values of w, l, and p. The header is thus extended to encode those values, as specified in Tables B.3 and B.4. To minimize bandwidth, w is encoded on one bit, supporting the two choices 32-bit ($w = 0$) and 64-bit ($w = 1$), l on 7 bits (with the MSB fixed at 0, i.e. supporting up to 63 rounds), and p on 8 bits (supporting parallelization degree up to 255). The datagram header is thus only 2 bytes longer than the header for fixed parameters.

[1] The header referred to is that of the datagram specified, not that of the data processed by the NORX instance.

[2] Note that NORX is capable of (safely) processing much larger data sizes, those are just the maximum values when our proposed datagrams are used.

Table B.1: NORX32 datagram for fixed parameters (offsets are in bytes)

Offset	0	1	2	3
0 4 8 12	Nonce N			
16	Header byte length $\|A\|$		Trailer byte length $\|Z\|$	
20	Encrypted payload byte length $\|C\|$			
24 ... ??	Header A			
?? ... ??	Encrypted payload C			
?? ... ??	Trailer Z			
?? ... ??	Tag T			

Table B.2: NORX64 datagram for fixed parameters (offsets are in bytes)

Offset	0	1	2	3
0 28	Nonce N			
32	Header byte length $\|A\|$		Trailer byte length $\|Z\|$	
36 40	Encrypted payload byte length $\|C\|$			
44 ... ??	Header A			
?? ... ??	Encrypted payload C			
?? ... ??	Trailer Z			
?? ... ??	Tag T			

Table B.3: NORX32 datagram for variable parameters (offsets are in bytes)

Offset	**0**	**1**	**2**	**3**
0 **4** **8** **12**	Nonce N			
16	Header byte length $\|A\|$		Trailer byte length $\|Z\|$	
20	Encrypted payload byte length $\|C\|$			
24	$w(1)\|\|l(7)$	p		
28 ... **??**	Header A			
?? ... **??**	Encrypted payload C			
?? ... **??**	Trailer Z			
?? ... **??**	Tag T			

Table B.4: NORX64 datagram for variable parameters (offsets are in bytes)

Offset	**0**	**1**	**2**	**3**
0 ... ... **28**	Nonce N			
32	Header byte length $\|A\|$		Trailer byte length $\|Z\|$	
36 **40**	Encrypted payload byte length $\|C\|$			
44	$w(1)\|\|l(7)$	p		
48 ... **??**	Header A			
?? ... **??**	Encrypted payload C			
?? ... **??**	Trailer Z			
?? ... **??**	Tag T			

C Addenda to Cryptanalysis

C.1 Diffusion Statistics for Inverse Round Functions

Table C.1 shows the diffusion statistics of the inverse round functions of NORX and ChaCha.

Table C.1: Diffusion statistics for inverse NORX and ChaCha round functions

	Inverse NORX32				Inverse ChaCha (32-bit)			
l	min	max	avg	med	min	max	avg	med
1	17	162	49.444	47	17	126	44.776	44
2	160	306	247.737	248	164	304	244.982	246
3	202	307	255.991	256	203	310	255.994	256
4	202	315	256.018	256	200	311	256.022	256
	Inverse NORX64				Inverse ChaCha (64-bit)			
l	min	max	avg	med	min	max	avg	med
1	17	203	51.346	49	17	142	46.129	45
2	262	568	433.742	435	194	543	382.667	383
3	440	593	511.995	512	440	591	511.964	512
4	435	585	512.011	512	433	596	511.991	512

C.2 Visualisation of Differentials for G_1

Fig. C.1 depicts the relations of the output differences of G_1 for input differences α_i with one active bit. The probability of an output difference in the tree can be computed by multiplying the values on the edges of the path leading from the root to the particular node.

C.3 Impossible Differential Cryptanalysis

Fig. C.2 shows the bit representations of the output differences of the impossible differential over 3.5 rounds of NORX64, which was presented in §7.2.4. The upper matrix illustrates the difference in forward direction and the lower matrix the one in backward direction. Each row corresponds to one of the 64-bit words of the state (denoted in little-endian), beginning with s_0 for the first row and ending with s_{15} for the last row. The conflict occurs in the 2nd bit of the 14th word.

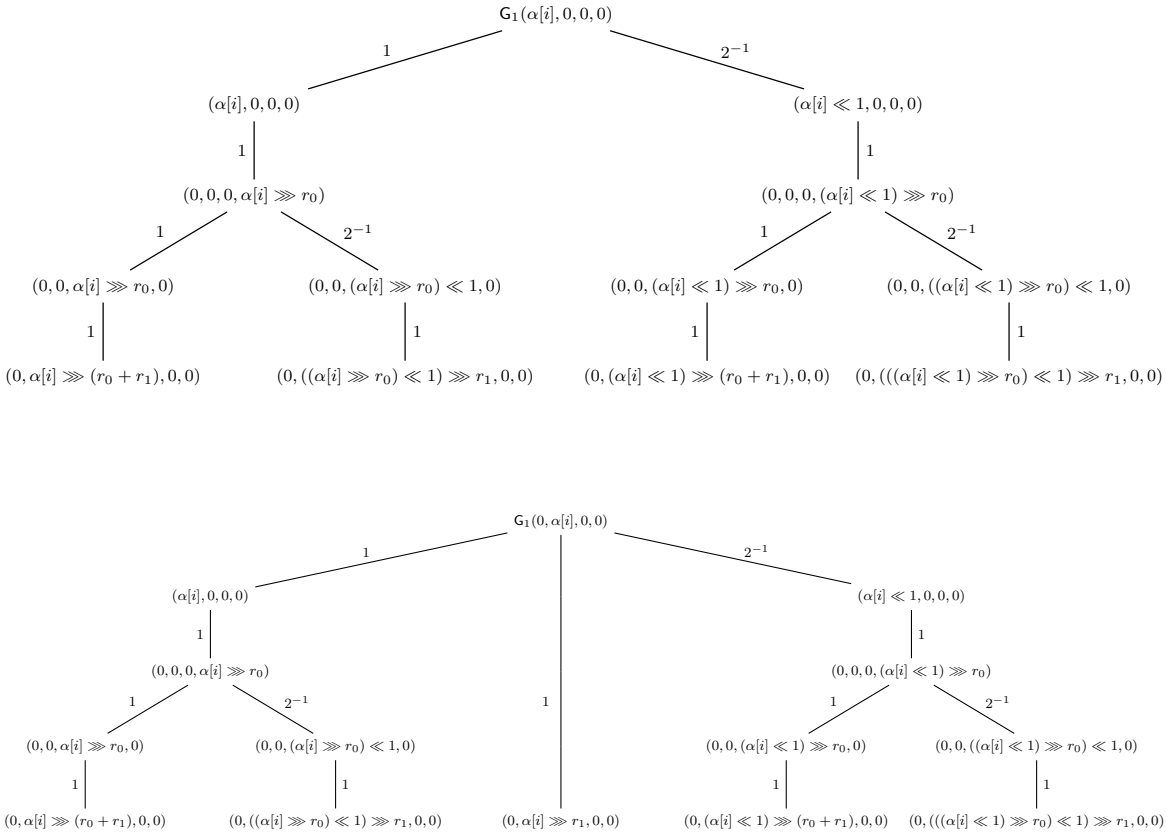

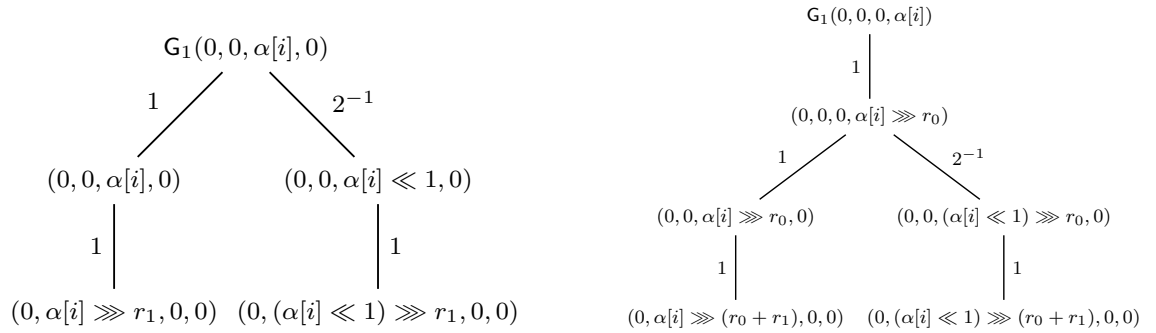

Figure C.1: Relations of the G_1 output differences

????	????	????	??1?	????	0???	????	????	????	??10	0???	??1?	????	1???	?00?	????
??10	0000	0???	?10?	????	????	??10	0000	0000	0000	0000	0000	0000	0???	??10	00??
??10	000?	??10	0000	????	?100	????	?100	0000	0000	00??	??10	0000	0000	0000	00??
0000	0000	00??	?10?	????	????	??10	0000	0000	?100	0???	?00?	????	?1??	??1?	???0
????	????	????	?10?	????	????	????	????	0???	????	????	????	????	????	????	????
????	??10	0???	100?	????	????	?10?	????	????	????	????	??1?	????	1???	????	????
????	????	?000	0000	????	????	????	??10	0000	????	?0??	????	?00?	????	????	10??
????	????	?1??	000?	????	1???	????	????	????	????	????	??0?	????	0???	????	????
????	????	????	??00	????	????	????	????	?1??	????	????	????	????	????	????	????
????	???1	0000	0000	0000	0000	0000	????	????	????	????	???1	????	?1??	????	????
????	????	??10	0000	0???	????	???1	0??1	0000	0???	??0?	????	??00	????	????	?100
????	????	??1?	?100	????	?0??	????	????	????	????	??1?	???0	0???	?1??	????	????
????	????	????	??10	0???	??0?	????	????	?00?	????	????	????	????	??0?	????	0???
??10	0000	0000	0000	0000	0000	0000	0???	??10	00??	???1	0000	0???	?10?	????	????
????	?100	0000	0000	00??	???1	0000	00?1	0000	00??	??00	000?	??00	0000	????	?100
??10	00??	??10	?100	0???	?00?	????	?1??	??1?	????	??10	???0	00??	?10?	????	????

vs.

????	????	????	????	????	????	????	????	????	????	????	????	????	????	????	????
????	????	????	????	????	????	????	????	????	????	????	????	????	????	????	????
????	????	????	????	????	????	????	????	????	????	????	????	????	????	????	????
????	????	????	????	????	????	????	????	????	????	????	????	????	????	????	????
????	????	????	????	????	????	????	????	????	????	????	????	????	????	????	????
????	????	????	????	????	????	????	????	????	????	????	????	????	????	????	????
????	????	????	????	????	????	????	????	????	????	????	????	????	????	????	????
????	????	????	????	????	????	????	????	????	????	????	????	????	????	????	????
????	????	????	????	????	????	????	????	????	????	????	????	????	????	????	????
????	????	????	????	????	????	????	????	????	????	????	????	????	????	1000	0000
????	????	????	????	????	????	????	????	????	????	????	????	????	????	????	????
????	????	????	????	????	????	????	????	????	????	????	????	????	????	????	????
????	????	????	????	????	????	????	????	????	????	????	????	????	????	????	????
????	????	????	????	????	????	????	????	????	????	????	????	????	????	????	????
????	????	????	????	????	????	????	????	????	????	????	????	????	????	1000	0000
????	????	????	????	????	????	????	????	????	????	????	????	????	????	????	????

Figure C.2: Bit representation of a 3.5-round impossible differential for 64-bit F

D Nonce Misuse-Resistant NORX

This section gives a preview on a nonce misuse-resistant version of NORX that is based on the misuse-resistant sponge (MRS) mode [36]. Instead of using MRS' full-state absorption in the data authentication part, NORX uses the usual rate-capacity layout though. This mode is basically a MAC-Then-Encrypt construction where the generated authentication tag over header, payload, and trailer is used as a nonce for the encryption function to encrypt the payload. Note that two passes over the payload data are required. Since this design is only a preview no parallel payload encryption is currently supported and therefore $p = 1$. The high-level interface is given in Fig. D.1 and the new initialise function is depicted in Fig. D.2. All other functions are unchanged and specified in Fig. 3.6.

Algorithm: AEADEnc(K, N, A, M, Z)

1. $S \leftarrow \mathsf{initialise}(K, N, \texttt{FE})$
2. $S \leftarrow \mathsf{absorb}(S, A, 01)$
3. $S \leftarrow \mathsf{absorb}(S, M, 02)$
4. $S \leftarrow \mathsf{absorb}(S, Z, 04)$
5. $S, T \leftarrow \mathsf{finalise}(S, 08)$
6. $S \leftarrow \mathsf{initialise}(K, T, \texttt{FF})$
7. $S, C \leftarrow \mathsf{encrypt}(S, M, 02)$
8. **return** C, T

Algorithm: AEADDec(K, N, A, C, Z, T)

1. $S \leftarrow \mathsf{initialise}(K, T, \texttt{FF})$
2. $S, M \leftarrow \mathsf{decrypt}(S, C, 02)$
3. $S \leftarrow \mathsf{initialise}(K, N, \texttt{FE})$
4. $S \leftarrow \mathsf{absorb}(S, A, 01)$
5. $S \leftarrow \mathsf{absorb}(S, M, 02)$
6. $S \leftarrow \mathsf{absorb}(S, Z, 04)$
7. $S, T' \leftarrow \mathsf{finalise}(S, 08)$
8. **if** $T = T'$ **then return** M **else return** $\perp$ **end**

Figure D.1: High-level interface functions of the nonce misuse-resistant NORX mode

Algorithm: initialise(K, N, v)

1. $k_0 \parallel k_1 \parallel k_2 \parallel k_3 \leftarrow K$, s.t. $|k_i| = w$
2. $n_0 \parallel n_1 \parallel n_2 \parallel n_3 \leftarrow N$, s.t. $|n_i| = w$
3. $S \leftarrow (n_0, n_1, n_2, n_3, k_0, k_1, k_2, k_3, u_8, u_9, u_{10}, u_{11}, u_{12}, u_{13}, u_{14}, u_{15})$
4. $(s_{12}, s_{13}, s_{14}, s_{15}) \leftarrow (s_{12}, s_{13}, s_{14}, s_{15}) \oplus (l, p, t, v)$
5. $S \leftarrow \mathsf{F}^l(S)$
6. $(s_{12}, s_{13}, s_{14}, s_{15}) \leftarrow (s_{12}, s_{13}, s_{14}, s_{15}) \oplus (k_0, k_1, k_2, k_3)$
7. **return** S

Figure D.2: Initialisation of the nonce-misuse-resistant NORX mode