TinyJAMBU: A Family of Lightweight Authenticated Encryption Algorithms (Version 2)

Designers and Submitters: Hongjun Wu and Tao Huang

Division of Mathematical Sciences
Nanyang Technological University
wuhongjun@gmail.com

17 May 2021

Contents

2

Chapter 1

Introduction

JAMBU is a lightweight authenticated encryption mode submitted to the CAE-SAR competition [6]. JAMBU is the smallest block cipher authenticated encryption mode in the CAESAR competition, and it was selected to the Third Round of the competition. JAMBU has been presented at the NIST Lightweight Cryptography Workshop 2015 [22].

In this report, we propose the TinyJAMBU mode which is a small variant of the JAMBU mode. TinyJAMBU mode is based on a keyed permutation. The state size of TinyJAMBU is only two thirds of that of JAMBU, the message block size of TinyJAMBU is half of that of JAMBU mode. When nonce is reused, TinyJAMBU provides better authentication security than JAMBU mode. The authentication security of TinyJAMBU mode is better than the Duplex mode [2] when nonce is reused (for the same permutation size and message block size).

In this report, we propose a lightweight 128-bit keyed permutation with no key schedule. The permutation is based on a 128-bit nonlinear feedback shift register. This lightweight permutation is used in the TinyJAMBU mode. The keyed permutation supports three possible key sizes: 128 bits, 192 bits, 256 bits.

We implemented TinyJAMBU-128 with the NIST LWC hardware API in VHDL and synthesized it on ASIC using 180nm UMC technology standard cell library. For the implementation with 8-bit input port and 8 rounds of state update function, only 4352 Gate Equivalent (GE) are used (the CryptoCore without the API is 2708 GE). The throughput is 206 Mbps for long associated data and 128 Mbps for long plaintext. The hardware area can be reduced to 3224 GE if the secret key is an embedded fixed key.

Chapter 2

TinyJAMBU Authenticated Encryption Mode

The TinyJAMBU mode is a small variant of the JAMBU mode which is a third-round candidate of the CAESAR competition. In the TinyJAMBU mode, a 128-bit keyed permutation is used, the state size is 128 its, and the message block size is 32 bits. When nonce is reused, the TinyJAMBU mode provides better authentication security than the JAMBU mode.

When nonce is reused, the TinyJAMBU mode provides better authentication security than the Duplex mode (for the same permutation size and the same message block size). The reason is that the attacker can easily set part of the state to arbitrary value when nonce is reused in the Duplex mode, while it is difficult to do that in the TinyJAMBU mode.

The TinyJAMBU mode is shown in Fig. 2.1. If the last block of the associated data (or plaintext) is not a full block, the length of the partial block (the number of bytes) is xored to the state.
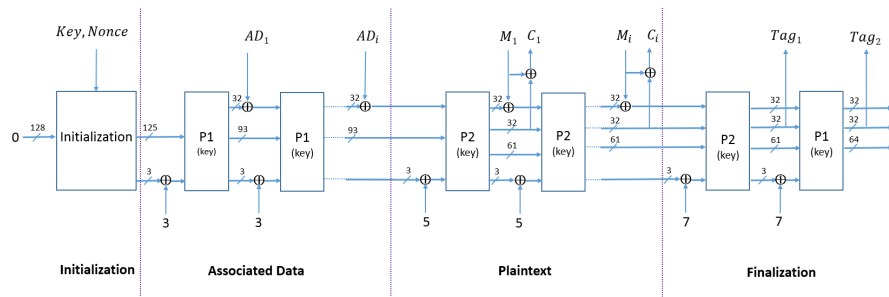
Figure 2.1: The TinyJAMBU mode for 128-bit state and keyed-permutations

Chapter 3

Specification

3.1   Recommended parameter sets

TinyJAMBU supports three key sizes: 128 bits, 192 bits and 256 bits.

- Primary member: TinyJAMBU-128
  128-bit key, 96-bit nonce, 64-bit tag, 128-bit state

- TinyJAMBU-192
  192-bit key, 96-bit nonce, 64-bit tag, 128-bit state

- TinyJAMBU-256
  256-bit key, 96-bit nonce, 64-bit tag, 128-bit state

3.2   Operations, Variables and Functions

The operations, variables and functions used in TinyJAMBU are defined below.

3.2.1   Operations

The following operations are used in the description of TinyJAMBU:

$\oplus$	:	bit-wise exclusive OR
$\&$	:	bit-wise AND
$\sim$	:	bit-wise NOT
$\|$	:	concatenation
$\lfloor a \rfloor$	:	floor operator, gives the integer part of a

3.2.2 Variables and Constants

The following variables and constants are used in TinyJAMBU:

$a_{\{i\cdots j\}}$	:	the word consists of $a_i\|\|a_{i+1}\|\|\cdots\|\| a_j$, where a_i is the ith bit of a.
AD	:	associated data, a sequence of bytes.
ad_i	:	one bit of associated data.
$adlen$	:	the length of associated data in bits.
C	:	ciphertext, a sequence of bytes
c_i	:	the ith ciphertext bit.
$FrameBits$	:	Three-bit FrameBits.
		FrameBits = 1 for nonce
		FrameBits = 3 for associated data
		FrameBits = 5 for plaintext and ciphertext
		FrameBits = 7 for finalization
$FrameBits_i$	:	The ith bit of FrameBits.
K	:	the key.
k_i	:	the ith bit of K.
$klen$	:	the key length in bits.
M	:	the plaintext, a sequence of bytes
m_i	:	the ith bit of the plaintext.
$mlen$	:	the length of the plaintext in bits.
NONCE	:	the 96-bit nonce.
nonce_i	:	the ith bit of the 96-bit nonce.
P_n	:	the 128-bit permutation with n rounds
S	:	the 128-bit state of the permutation.
s_i	:	the ith bit of the state of the permutation.
T	:	the 64-bit authentication tag.
t_i	:	the ith bit of the authentication tag.

3.2.3   The Keyed Permutation P_n

In TinyJAMBU, a 128-bit keyed permutation is used. The permutation P_n consists of n rounds. In the ith round of the permutation, a 128-bit nonlinear feedback shift register is used to update the state as follows (shown in Fig. 3.1):

> StateUpdate(S, K, i):
>   feedback $= s_0 \oplus s_{47} \oplus (\sim (s_{70} \& s_{85})) \oplus s_{91} \oplus k_{i \bmod klen}$
>   for j from 0 to 126: $s_j = s_{j+1}$
> $s_{127} =$ feedback
> end

For example, P_{640} means that the state of the permutation is updated using the function StateUpdate() for 640 times. 32 rounds of the permutation can be computed in parallel on 32-bit CPU.

Figure 3.1: The 128-bit Nonlinear Feedback Shift Register in TinyJAMBU

3.3   TinyJAMBU-128

TinyJAMBU-128 uses a 128-bit key and a 96-bit nonce. The associated data length and the plaintext length are less than 2^{50} bytes. The authentication tag is 64-bit. The TinyJAMBU authenticated encryption mode is used in TinyJAMBU-128.

3.3.1   The initialization

In the keyed permutation of TinyJAMBU-128, the 128-bit key of TinyJAMBU-128 is used, and the $klen$ is set to 128.

The initialization of TinyJAMBU-128 consists of two stages: key setup and nonce setup.

Key Setup. The key setup is to randomize the state using the keyed permutation P_{1024}.

1. Set the 128-bit state S as 0.

2. Update the state using P_{1024}.

Nonce Setup. The nonce setup consists of three steps. In each step, the Framebits of nonce (the value is 1) are XORed with the state, then we update the state using the keyed permutation P_{640}, then 32 bits of the nonce are XORed with the state.

> for i from 0 to 2:
> $s_{\{36\cdots38\}} = s_{\{36\cdots38\}} \oplus FrameBits_{\{0\cdots2\}}$
>   Update the state using P_{640}
> $s_{\{96\cdots127\}} = s_{\{96\cdots127\}} \oplus nonce_{\{32i\cdots32i+31\}}$
> end for

3.3.2 Processing the associated data

After the initialization, we process the associated data AD. In each step, the Framebits of associated data (the value is 3) are XORed with the state, then we update the state using the keyed permutation P_{640}, then 32 bits of the associated data are XORed with the state.

Processing the full blocks of associated data:

> for i from 0 to $\lfloor adlen/32 \rfloor$:
> $s_{\{36\cdots38\}} = s_{\{36\cdots38\}} \oplus FrameBits_{\{0\cdots2\}}$
>   Update the state using P_{640}
> $s_{\{96\cdots127\}} = s_{\{96\cdots127\}} \oplus ad_{\{32i\cdots32i+31\}}$
> end for

Processing the partial block of associated data. If the last block is not a full block (it is called a partial block), the last block is XORed to the state, and the number of bytes of associated data in the partial block is XORed to the state.

> if $(adlen \bmod 32) > 0$:
> $s_{\{36\cdots38\}} = s_{\{36\cdots38\}} \oplus FrameBits_{\{0\cdots2\}}$
>   Update the state using P_{640}
> $lenp = adlen \bmod 32$ /* number of bits in the partial block */
> $startp = adlen - lenp$ /* starting position of the partial block */
> $s_{\{96\cdots96+lenp-1\}} = s_{\{96\cdots96+lenp-1\}} \oplus ad_{\{startp\cdots adlen-1\}}$
>
>   /* the number of bytes in the partial block is XORed to the state*/
> $s_{\{32\cdots33\}} = s_{\{32\cdots33\}} \oplus (lenp/8)$
> end if

3.3.3 The encryption

After processing the associated data, we encrypt the plaintext M. In each step, the Framebits of plaintext (the value is 5) are XORed with the state, then we update the state using the keyed permutation P_{1024}, then 32 bits of the plaintext are XORed with the state, and we obtain 32 bits of ciphertext by XORing the plaintext with another part of the state.

Processing the full blocks of plaintext:

> for i from 0 to $\lfloor mlen/32 \rfloor$:
>> $s_{\{36\cdots38\}} = s_{\{36\cdots38\}} \oplus FrameBits_{\{0\cdots2\}}$
>> Update the state using P_{1024}
>> $s_{\{96\cdots127\}} = s_{\{96\cdots127\}} \oplus m_{\{32i\cdots32i+31\}}$
>> $c_{\{32i\cdots32i+31\}} = s_{\{64\cdots95\}} \oplus m_{\{32i\cdots32i+31\}}$
>
> end for

Processing the partial block of plaintext. If the last block is not a full block (it is a partial block), the last block is XORed to the state, and the number of bytes in the partial block is XORed to the state.

> if $(mlen \bmod 32) > 0$:
>> $s_{\{36\cdots38\}} = s_{\{36\cdots38\}} \oplus FrameBits_{\{0\cdots2\}}$
>> Update the state using P_{1024}
>> $lenp = mlen \bmod 32$ /* number of bits in partial block */
>> $startp = mlen - lenp$ /* starting position of partial block */
>> $s_{\{96\cdots96+lenp-1\}} = s_{\{96\cdots96+lenp-1\}} \oplus m_{\{startp\cdots mlen-1\}}$
>> $c_{\{startp\cdots mlen-1\}} = s_{\{64\cdots64+lenp-1\}} \oplus m_{\{startp\cdots mlen-1\}}$
>> /* the length (bytes) of the last partial block is XORed to the state*/
>> $s_{\{32\cdots33\}} = s_{\{32\cdots33\}} \oplus (lenp/8)$
>
> end if

3.3.4 The finalization

After encrypting the plaintext, we generate the 64-bit authentication tag T as follows. The Framebits of finalization (the value is 7) are XORed with the state.

> $s_{\{36\cdots38\}} = s_{\{36\cdots38\}} \oplus FrameBits_{\{0\cdots2\}}$
> Update the state using P_{1024}
> $t_{\{0\cdots31\}} = s_{\{64\cdots95\}}$

> $s_{\{36\cdots38\}} = s_{\{36\cdots38\}} \oplus FrameBits_{\{0\cdots2\}}$
> Update the state using P_{640}
> $t_{\{32\cdots63\}} = s_{\{64\cdots95\}}$

3.3.5 The decryption

In a decryption process, the initialization and processing the associate data are the same as the encryption process. After processing the associated data, we decrypt the ciphertext C. In each step, the Framebits of plaintext (the value is 5) are XORed with the state, then we update the state using the keyed permutation P_{1024}. We obtain 32 bits of plaintext by XORing the ciphertext with 32 state bits $s_{\{64\cdots95\}}$, then the plaintext is XORed with the state bits $s_{\{96\cdots127\}}$.

Processing the full blocks of ciphertext:

for i from 0 to $\lfloor mlen/32 \rfloor$:
$\quad\quad s_{\{36\cdots38\}} = s_{\{36\cdots38\}} \oplus FrameBits_{\{0\cdots2\}}$
$\quad\quad$ Update the state using P_{1024}
$\quad\quad m_{\{32i\cdots32i+31\}} = s_{\{64\cdots95\}} \oplus c_{\{32i\cdots32i+31\}}$
$\quad\quad s_{\{96\cdots127\}} = s_{\{96\cdots127\}} \oplus m_{\{32i\cdots32i+31\}}$
end for

Processing the partial block of ciphertext. If the last block is not a full block (it is a partial block), the number of bytes in the partial block is XORed to the state.

if $(mlen \bmod 32) > 0$:
$\quad\quad s_{\{36\cdots38\}} = s_{\{36\cdots38\}} \oplus FrameBits_{\{0\cdots2\}}$
$\quad\quad$ Update the state using P_{1024}
$\quad\quad lenp = mlen \bmod 32$ /* number of bits in partial block */
$\quad\quad startp = mlen - lenp$ /* starting position of partial block */
$\quad\quad m_{\{startp\cdots mlen-1\}} = s_{\{64\cdots64+lenp-1\}} \oplus c_{\{startp\cdots mlen-1\}}$
$\quad\quad s_{\{96\cdots96+lenp-1\}} = s_{\{96\cdots96+lenp-1\}} \oplus m_{\{startp\cdots mlen-1\}}$
$\quad\quad$ /* the length (bytes) of the last partial block is XORed to the state*/
$\quad\quad s_{\{32\cdots33\}} = s_{\{32\cdots33\}} \oplus (lenp/8)$
end if

3.3.6 The verification

After decrypting the plaintext, we generate a 64-bit authentication tag T', then compare T' with the received tag T. The Framebits of finalization are of value 7.

$s_{\{36\cdots38\}} = s_{\{36\cdots38\}} \oplus FrameBits_{\{0\cdots2\}}$
Update the state using P_{1024}
$t'_{\{0\cdots31\}} = s_{\{64\cdots95\}}$
$s_{\{36\cdots38\}} = s_{\{36\cdots38\}} \oplus FrameBits_{\{0\cdots2\}}$
Update the state using P_{640}
$t'_{\{32\cdots63\}} = s_{\{64\cdots95\}}$

$T' = t'_{\{0\cdots63\}}$. Accept the message if $T' = T$; otherwise, reject.

3.4 TinyJAMBU-192

TinyJAMBU-192 uses a 192-bit key and a 96-bit nonce. The associated data length and the plaintext length are less than 2^{50} bytes. The authentication tag is 64-bit.

The design of TinyJAMBU-192 is similar to that of TinyJAMBU-128. The differences between TinyJAMBU-192 and TinyJAMBU-128 are given below.

Keyed Permutation. In the keyed permutation of TinyJAMBU-192, a 192-bit key is used, and the *klen* is set to 192.

The keyed permutation P_{1024} in TinyJAMBU-128 is replaced with P_{1152} in TinyJAMBU-256. P_{640} in TinyJAMBU-128 is still P_{640} in TinyJAMBU-192.

3.5 TinyJAMBU-256

TinyJAMBU-256 uses a 256-bit key and a 96-bit nonce. The associated data length and the plaintext length are less than 2^{50} bytes. The authentication tag is 64-bit.

The design of TinyJAMBU-256 is similar to that of TinyJAMBU-128. The differences between TinyJAMBU-256 and TinyJAMBU-128 are given below.

Keyed Permutation. In the keyed permutation of TinyJAMBU-256, a 256-bit key is used, and the *klen* is set to 256.

The keyed permutation P_{1024} in TinyJAMBU-128 is replaced with P_{1280} in TinyJAMBU-256. P_{640} in TinyJAMBU-128 is still P_{640} in TinyJAMBU-256.

Chapter 4

The Tweak in the Design

4.1 The Tweak

The permutation P_{384} in the original TinyJAMBU is used to process the nonce and associated data. In the new TinyJAMBU (version 2), the permutation P_{640} is used to replace P_{384} with more rounds. There is no change to the processing of plaintext blocks.

There is no tweak of TinyJAMBU is identical to the "Planned Tweak Proposals" specified in the "TinyJAMBU Update", which was submitted to NIST on September 18, 2020.

4.2 Reason for the Tweak

The reason for the tweak is to provide larger security margin for the protection of nonce and associated data against differential forgery attack.

In [19], Saha et al. analysed the security margin of the nonce and associated data of TinyJAMBU against the differential forgery attack. In the TinyJAMBU design, our analysis shows that the differential forgery attack against nonce and associated data succeeds with probability at most 2^{-73}. In [19], it is shown that some NAND gates in the differential trail are not independent, so the forgery attack against nonce and associated data succeeds with probability $2^{-70.64}$. The MILP code being used in the analysis is available online [20].

In TinyJAMBU v2, we increased the round number of the permutation being used to process nonce and associated data. It is shown in [19] that around 338 rounds are needed to resist the differential forgery attack on nonce and associated data. The permutation P_{640} in TinyJAMBU v2 has 302 more rounds than 338 rounds, so the security margin is large.

In our security analysis, we take into account the related NAND gates, as analyzed in [19]. The differential analysis of TinyJAMBU v2 is much simpler than that of the original TinyJAMBU, because P_{640} itself is a strong permutation for protecting 32-bit data blocks against differential forgery attack.

Chapter 5

Security Goals

In TinyJAMBU, each pair of key and nonce is used to protect only one message. If verification fails, the new tag and the decrypted plaintext should not be given as output.

In TinyJAMBU, the associated data plays the same role as nonce. When the same nonce but different associated data are used for a key, it is equivalent to the use of unique nonce for the key.

The nonce misuse happens when the same nonce and the same associated data are reused for a key. When nonce is misused, TinyJAMBU provides strong protection of the secret key, and provides strong authentication security, but provides weak protection of the plaintext.

5.1 Security goals for unique nonce

The security goals of TinyJAMBU for unique nonce are given in Table 5.1. We assume that each key is used to process at most 2^{50} byes of messages (associated data, plaintext/ciphertext), and each message is at least 8 bytes. Note that the authentication security in Table 5.1 includes the integrity security of plaintext, associated data and nonce.

Table 5.1: Security Goals of TinyJAMBU with Unique Nonce

	Encryption	Authentication
TinyJAMBU-128	112-bit	64-bit
TinyJAMBU-192	168-bit	64-bit
TinyJAMBU-256	224-bit	64-bit

5.2   Security goals for misused nonce

When nonce is misused in TinyJAMBU (the same nonce and the same associated data are reused for a key), the secret key of TinyJAMBU remains strong, and the authentication of TinyJAMBU remains strong (suppose that a secret key is used to process at most 2^{50} bytes of data, and each message consists of at least 8 bytes of data).

When nonce is reused, an attacker is able to decrypt the ciphertext since the encryption of TinyJAMBU is somehow similar to the Cipher Feedback mode.

The security goals of TinyJAMBU for reused nonce are given in Table 5.2.

Table 5.2: Security Goals of TinyJAMBU with Repeated Nonce (an adversary has the maximum forgery advantage when a key was used to process adaptively chosen 2^{50} bytes of data, and each message is at least 8 bytes)

	Secret Key	Authentication	Max. Forgery Adv.
TinyJAMBU-128	112-bit	64-bit	2^{-15}
TinyJAMBU-192	168-bit	64-bit	2^{-15}
TinyJAMBU-256	224-bit	64-bit	2^{-15}

5.3   Security goals for unprotected decryption

For unprotected decryption, we mean that the attacker is able to retrieve the decrypted plaintext when the verification fails. It happens when the attacker has physical access to the decryption device, and the decryption device does not fully protect the decrypted message when the verification fails (for example, the device is ultra lightweight and strong protection of decrypted plaintext is not implemented).

Suppose that the attacker can process at most 2^{50} bytes of data (the attacker may have limited time accessing the device, and the lightweight device may not be very fast), and each message consists of at least 8 bytes of data. The security goals of TinyJAMBU for unprotected decryption are given in Table 5.3. The security goals are identical to the security goals for misused nonce. The security analysis of unprotected decryption is also identical to the analysis of misused nonce (for the unprotected decryption, we assume that the decrypted plaintext is known to the attacker).

Table 5.3: Security Goals of TinyJAMBU for Unprotected Decryption

	Secret Key	Authentication	Max. Forgery Adv.
TinyJAMBU-128	112-bit	64-bit	2^{-15}
TinyJAMBU-192	168-bit	64-bit	2^{-15}
TinyJAMBU-256	224-bit	64-bit	2^{-15}

Chapter 6

Security of the TinyJAMBU Mode

In this section, we analyse the security of TinyJAMBU mode. To simplify the analysis, we consider $P_K : \{0,1\}^b \times \{0,1\}^k \to \{0,1\}^b$ as an underlying ideal keyed-permutation used in the mode, and use it to replace $P1$ and $P2$.

Our security proof is inspired by the security proof for the sponge function [11] and the SAEB mode [17].

6.1 Security Model

Let $\mathcal{E}$ be the authenticated encryption function, which takes key, nonce, associated data and plaintext as input and outputs ciphertext and authentication tag. The authenticated encryption process can be denoted as $(C, T) \longleftarrow \mathcal{E}_K(N, AD, M)$.

Similarly, the authenticated decryption function is $\mathcal{D}$, which takes key, nonce, associated data, ciphertext and authentication tag as input, and outputs plaintext when verification is successful or $\perp$ otherwise. We denote the decryption and verification process as $P/\perp \longleftarrow \mathcal{D}_K(N, AD, C, T)$.

We consider an adversary $\mathcal{A}$ that has access to $(\mathcal{E}_K, \mathcal{D}_K)$ and an oracle $(\$, \perp)$ where $\$$ takes (N, AD, M) as input and returns a random bit string with the same length as C; $\perp$ is an oracle that returns the reject symbol $\perp$ for any query. We assume that the adversary does not make any trivial authenticated encryption queries or trivial authenticated encryption queries. In addition, we assume that key is unknown to the adversary. Hence, the adversary does not have access to the underlying permutation. Let $q_{\mathcal{E}}$ and $q_{\mathcal{D}}$ denote the maximum number of queries that the adversary make for authenticated encryption and authenticated decryption respectively. Let $\sigma_{\mathcal{E}}$ and $\sigma_{\mathcal{D}}$ denote the total number of blocks involved in authenticated encryption and authenticated decryption respectively. When consider j-th authenticated encryption query, $\sigma_{\mathcal{E},j}$ is used to denote the total number of keyed permutation invocations. Similarly, $\sigma_{\mathcal{D},j}$ is

the total number of keyed permutation invocations in the j-th decryption and verification query.

6.2 Privacy

In the following, we assume that the secret key is unknown to the adversary. We define the security of TinyJAMBU on privacy against nonce-respecting adversary $\mathcal{A}$ in a chosen plaintext attack setting as follows:

$$\mathbf{Adv}^{\text{priv}}_{\text{TinyJAMBU}} = |\Pr[\mathcal{A}^{\text{TinyJAMBU}_{\mathcal{E}_K}} = 1] - \Pr[\mathcal{A}^{\$} = 1]|,$$

where $\mathcal{A}^{\text{TinyJAMBU}_{\mathcal{E}_K}}$ refers to the case that the adversary interacts with Tiny-JAMBU authenticated encryption and $\mathcal{A}^{\$}$ refers to the case that the adversary interacts with a random oracle.

The following theorem gives the security bound on privacy for TinyJAMBU.

Thoerem 1. *Let $\mathcal{A}$ be a nonce-respecting adversary makes $q_{\mathcal{E}}$ encryption queries of at most $\sigma_{\mathcal{E}}$ blocks. Then*

$$\mathbf{Adv}^{priv}_{TinyJAMBU} \leq \frac{\sigma_{\mathcal{E}}^2}{2^{n+1}}.$$

Proof. We use $s := ([s]_v \parallel [s]_r \parallel [s]_c)$ to denote the internal state of the authenticated encryption scheme, where $[s]_v$ is the v-bit state segment that the message blocks inject into; $[s]_r$ is the r-bit state segment that is used as keystream, which is similar to the *rate* segment in the Duplex sponge mode; $[s]_c$ is the c-bit state segment that only frame bits are xored with, which is similar to the *capacity* segment in the Duplex sponge mode. $s_{j,k}$ is used to denote the internal state at the k-th block of j-th query, where $1 \leq k \leq \sigma_{\mathcal{E},j}$, $1 \leq j \leq q_{\mathcal{E}}$.

We start out proof by defining the following event:

Coll: $(s_{j,k-1} \neq s_{j',k'-1})$ and $(s_{j,k} = s_{j',k'})$ for $j, j' \in 1, ..., q_{\mathcal{E}}$, $k \in 1, ..., \sigma_{\mathcal{E},j}$ and $k' \in 1, ..., \sigma_{\mathcal{E},j'}$.

Note that for nonce-respecting adversaries, Coll occurs if and only if an internal state collision occurs.

In regard to the difference between TinyJAMBU$_{\mathcal{E}_K}$ and $\$$, we have the following Lemma:

Lemma 1. *The output of TinyJAMBU$_{\mathcal{E}_K}$ and $\$$ are identically distributed until the event Coll occurs.*

Proof. From the definition of $\$$, the output (C, T) is a uniformly random bit string. Consider $\mathcal{E}_K$, without event Coll, the input of each keyed-permutation invocation is a new value. Since P_K is an ideal keyed-permutation, the output of P_K is also uniformly random, which is identical to the output of $\$$. $\square$

From Lemma 1, the advantage of adversary is bounded by the probability that the event Coll occurs. Thus, we have

$$\mathbf{Adv}_{\mathrm{TinyJAMBU}}^{\mathrm{priv}} \leq \Pr[\mathtt{Coll}]$$

Given $\sigma_{\mathcal{E}}$ keyed permutation invocations, the probability of event Coll is upper bounded by $\binom{\sigma_{\mathcal{E}}}{2} \cdot 2^{-n}$.

Thus, we have

$$\mathbf{Adv}_{\mathrm{TinyJAMBU}}^{\mathrm{priv}} \leq \frac{\sigma_{\mathcal{E}}^2}{2^{n+1}}.$$

$\square$

In the TinyJAMBU authenticated cipher, we have $b = 128$, $r = v = 32$, $c = 64$ and $\sigma \leq 2^{48}$. Hence, after 2^{48} queries, the security bound for TinyJAMBU is $\mathbf{Adv}_{\mathrm{TinyJAMBU}}^{\mathrm{priv}} \leq 2^{96}/2^{129} \approx 2^{-33}$. It means that when nonce is not reused, the adversary has the advantage of 2^{-33} to distinguishing ciphertext from random after 2^{50} bytes of plaintext are processed.

6.3 Authenticity

For the authenticity of TinyJAMBU mode, we will prove the security for a nonce-reuse adversary $\mathcal{A}$. In the following, we assume that the secret key is unknown to the adversary.

We define the authenticity of TinyJAMBU as follows:

$$\mathbf{Adv}_{\mathrm{TinyJAMBU}}^{auth} = \Pr[\mathcal{A}^{\mathrm{TinyJAMBU}_{\mathcal{E}_K}, \mathrm{TinyJAMBU}_{\mathcal{D}_K}} forges.]$$

Thoerem 2. *Let $\mathcal{A}$ be a nonce-reuse adversary makes $q_{\mathcal{E}}$ authenticated encryption queries of at most $\sigma_{\mathcal{E}}$ blocks and $q_{\mathcal{D}}$ authenticated decryption queries of at most $\sigma_{\mathcal{D}}$ blocks. Then*

$$\mathbf{Adv}_{TinyJAMBU}^{auth} \leq \left(\frac{e\sigma_{\mathcal{E}}}{\rho 2^r}\right)^{\rho} \cdot \frac{2^r}{\sqrt{\rho}} + \frac{(\sigma_{\mathcal{E}} + \sigma_{\mathcal{D}})(\rho - 2)}{2^{c+v/2+1}} + \frac{q_{\mathcal{D}}}{2^t},$$

where e is the Euler's number and ρ is a positive integer constant.

Proof. We use similar notations as in the proof of Theorem 1 but with a few adjustments to include the decryption. We use $s^{\mathcal{E}}$ and $s^{\mathcal{D}}$ to denote the internal state in encryption and decryption queries respectively. So $s_{j,k}^{\mathcal{E}}$ refers to the k-th keyed-permutation invocation in the j-th encryption query.

Then we define the following events:

1. $\mathtt{Coll}_{\mathcal{E}}$: $(s_{j,k-1}^{\mathcal{E}} \neq s_{j',k'-1}^{\mathcal{E}})$ and $(s_{j,k}^{\mathcal{E}} = s_{j',k'}^{\mathcal{E}})$ for $j, j' \in 1, ..., q_{\mathcal{E}}$, $k \in 1, ..., \sigma_{\mathcal{E},j}$ and $k' \in 1, ..., \sigma_{\mathcal{E},j'}$;

2. $\mathtt{Coll}_{\mathcal{D}}$: $(s_{j,k-1}^{\mathcal{D}} \neq s_{j',k'-1}^{\mathcal{E}})$ and $(s_{j,k}^{\mathcal{D}} = s_{j',k'}^{\mathcal{E}})$ for $j \in 1, ..., q_{\mathcal{D}}$, $j' \in 1, ..., q_{\mathcal{E}}$, $k \in 1, ..., \sigma_{\mathcal{D},j}$ and $k' \in 1, ..., \sigma_{\mathcal{E},j'}$;

3. `MultiColl`: there exists ρ multi-collisions on the rate segment of state for a positive integer ρ.

The event $\texttt{Coll}_{\mathcal{E}}$ is exactly the same as `Coll` in the proof of Theorem 1. The event $\texttt{Coll}_{\mathcal{D}}$ is an authenticated decryption query has a non-trivial state collision with some authenticated encryption query. Here 'trivial state collision' refers to a state collision due to common prefix blocks.

A successful forgery can either from an internal state collision or a random guess without state collision. Here we only need to consider the internal state collisions in the two collision events $\texttt{Coll}_{\mathcal{E}}$ and $\texttt{Coll}_{\mathcal{D}}$. Note that the collisions in two authenticated decryption queries will not lead to a forgery. The probability that a successful random guess without state collision is $1/2^t$. Hence, we upper bound the probability of forgeries as follows:

$$\Pr[\mathcal{A}^{\mathrm{TinyJAMBU}_{\mathcal{E}_K}, \mathrm{TinyJAMBU}_{\mathcal{D}_K}} \mathit{forges}] \leq \Pr[\texttt{Coll}_{\mathcal{E}}] + \Pr[\texttt{Coll}_{\mathcal{D}}] + \frac{q_{\mathcal{D}}}{2^t}.$$

Now we compute the probability of the state collision events under the consideration of the number of multi-collisions ($\#MC$) on the rate segment of the state. As nonce reuse in our setting, it is possible to query multiple messages with the same prefix blocks. However, we consider this case as the same class of multi-collision. Only when the previous blocks are different, two blocks can be considered as different multi-collisions. We set a positive integer ρ. Let $f(t)$ be the probability that a state collision when t multi-collisions on the rate segment present. When the number of multi-collisions is less than ρ, the probability of state collision is given by:

$$\Pr[\texttt{Coll}_{<\rho}] = \sum_{t=0}^{\rho-1} \Pr[\#MC = t] \cdot f(t).$$

When the number of multi-collisions is at least ρ and of course at most σ, the probability of state collision is given by:

$$\Pr[\texttt{Coll}_{>=\rho}] = \sum_{t=\rho}^{\sigma} \Pr[\#MC = t] \cdot f(t) \leq \Pr[\#MC >= \rho].$$

With the definition of the event `MultiColl`, we have the following inequality:

$$\begin{aligned}\Pr[\texttt{Coll}_{\mathcal{E}}] + \Pr[\texttt{Coll}_{\mathcal{D}}] \leq \quad & \Pr[\texttt{Coll}_{\mathcal{E}} | \neg(\texttt{MultiColl})] + \\ & \Pr[\texttt{Coll}_{\mathcal{D}} | \neg(\texttt{MultiColl})] + \\ & \Pr[\texttt{MultiColl}].\end{aligned}$$

Using the result from Equation (4) in [11], we can bound the probability of `MultiColl` as:

$$\Pr[\texttt{MultiColl}] \leq (\frac{e\sigma_{\mathcal{E}}}{\rho 2^r})^{\rho} \cdot \frac{2^r}{\sqrt{\rho}}. \tag{6.1}$$

Next, we analyse the probability for event $\mathtt{Coll}_{\mathcal{E}}$ when $\mathtt{MultiColl}$ is false. Consider a state $s_{\alpha,i}^{\mathcal{E}}$, which is the i-th block in the α-th query. Suppose that $\mathtt{Coll}_{\mathcal{E}}$ occurs between $s_{\alpha,i}^{\mathcal{E}}$ and a previous block $s_{\beta,j}^{\mathcal{E}}$, we compute the probability as follows.

Case 1. $s_{\alpha,i-1}^{\mathcal{E}}$ is a new value which is not equal to any of the previous states. In this case, $P_K(s_{\alpha,i-1}^{\mathcal{E}})$ is randomly chosen from $\{0,1\}^b$. Hence, the probability that $\mathtt{Coll}_{\mathcal{E}}$ occurs between $s_{\alpha,i}^{\mathcal{E}}$ and a previous block $s_{\beta,j}^{\mathcal{E}}$ is bounded by $(\sum_{i=1}^{\alpha-1} \sigma_{\mathcal{E},i} + i - 1) \cdot \frac{1}{2^n}$.

Case 2. $s_{\alpha,i-1}^{\mathcal{E}}$ is a state that has been queried before. In this case, the output of $P_K(s_{\alpha,i-1}^{\mathcal{E}})$ is a fixed value. Since $s_{\alpha,i}^{\mathcal{E}} = P_K(s_{\alpha,i-1}^{\mathcal{E}}) \oplus (M \parallel 0^r \parallel (frameBits \parallel 0^{c-4})$, $[s_{\alpha,i}^{\mathcal{E}}]_r$ is a known value. The event $\mathtt{Coll}_{\mathcal{E}}$ implies that $[s_{\alpha,i}^{\mathcal{E}}]_r = [s_{\beta,j}^{\mathcal{E}}]_r$ and $s_{\alpha,i-1}^{\mathcal{E}} \neq s_{\beta,j-1}^{\mathcal{E}}$. Suppose that there are θ classes of multi-collisions on $[s_{\alpha,i}^{\mathcal{E}}]_r$, which are $\Gamma_1, \Gamma_2, ..., \Gamma_\theta$ such that the previous states are mutually distinct. For each Γ_i, there are β_i different messages been queried before. Therefore, the probability that the message segment $[s_{\alpha,i}^{\mathcal{E}}]_v$ collide with a previous blocks queried in class Γ_i is given by $\beta_i/2^v$. The rate segments $[s_{\alpha,i}^{\mathcal{E}}]_r$ must collide from the definition of multi-collision. The capacity segment has collision probability $1/2^c$ since the previous blocks are different. In summary, the probability that $s_{\alpha,i}^{\mathcal{E}}$ collides with a previous state is $\Sigma_{i=1}^{\theta} \beta_i/2^{c+v}$.

Next, we compute the state collision probability for all authenticated encryption queries. We divide the state s into disjoint sets $\Theta_1 \bigcup \Theta_2 \bigcup ... \bigcup \Theta_t$ according to the values of the rate segment $[s]_r$. Clearly, any state collision can only occur within the same set.

Then we compute the collision probability within a set, say Θ_x. Suppose that there are θ multi-collision classes, $\Gamma_1, \Gamma_2, ..., \Gamma_\theta$, in Θ_x. Each class Γ_i is corresponding to β_i different messages with the same input state. Then, the probability that a state collision between any two classes Γ_i and Γ_j is given by $\min\{\frac{\beta_i \beta_j}{2^v}, 1\} \cdot \frac{1}{2^c}$. Thus, the overall probability that there is a state collision in set Θ_x is

$$\Pr[Coll\mathcal{E} \text{ in } \Theta_x | \neg(\mathtt{MultiColl})] = \sum_{i \neq j, 1 \leq i,j \leq \theta} \min\{\frac{\beta_i \beta_j}{2^v}, 1\} \cdot \frac{1}{2^c} \quad (6.2)$$

When $\beta_i = 2^{v/2}$ for all $1 \leq i \leq \theta$, the above Equation 6.2 has the maximum probability $\binom{\theta}{2} \cdot \frac{1}{2^c}$ with the minimum number of block queries $\gamma = \sum_{i=1}^{\theta} \beta_i = 2^{v/2}\theta$. We explain the details in Appendix A.

Since the event $\mathtt{MultiColl}$ is false, the maximum possible number of multi-collisions on the rate segment is $\rho - 1$. Thus, $\theta \leq \rho - 1$. We have

$$\Pr[Coll\mathcal{E} \text{ in } \Theta_x | \neg(\mathtt{MultiColl})] \leq \binom{\rho-1}{2} \frac{1}{2^c} = \frac{(\rho-1)(\rho-2)}{2^{c+1}}.$$

Given the total number of encryption query blocks $\sigma_{\mathcal{E}}$, the maximum number of disjoint sets that can reach maximum probability is $\frac{\sigma_{\mathcal{E}}}{\gamma} = \frac{\sigma_{\mathcal{E}}}{(\rho-1)2^{v/2}}$.

Finally, we are able to compute

$$\begin{aligned} \Pr[Coll\mathcal{E}|\neg(\texttt{MultiColl})] &\leq \frac{\sigma_{\mathcal{E}}}{(\rho-1)2^{v/2}} \cdot \frac{(\rho-1)(\rho-2)}{2^{c+1}} \\ &= \frac{\sigma_{\mathcal{E}}(\rho-2)}{2^{c+v/2+1}}. \end{aligned}$$

Generally, we have $\sigma_{\mathcal{E}} \cdot 2^{-r} \leq \rho$ from the Pigeonhole principle. So the probability of *Case 2* is at least the probability of *Case 1*. Hence, the sum of probabilities in *Case 2* gives the bound for $\Pr[\texttt{Coll}]_{\mathcal{E}}$:

$$\Pr[\texttt{Coll}_{\mathcal{E}}|\neg(\texttt{MultiColl})] \leq \frac{\sigma_{\mathcal{E}}(\rho-2)}{2^{c+v/2+1}}.$$

We bound the probability for event $\texttt{Coll}_{\mathcal{D}}$ when $\texttt{MultiColl}$ is false. The probability that a $\texttt{Coll}_{\mathcal{D}}$ occurs between $s_{\alpha,i}^{\mathcal{D}}$ and some block from authenticated encryption $s_{\beta,j}^{\mathcal{E}}$ can be computed in a similar way as the computation for $\texttt{Coll}_{\mathcal{E}}$, resulting the following bound:

$$\Pr[\texttt{Coll}_{\mathcal{D}}|\neg(\texttt{MultiColl})] \leq \frac{\sigma_{\mathcal{D}}(\rho-2)}{2^{c+v/2+1}}.$$

In summary, we have

$$\mathbf{Adv}_{\text{TinyJAMBU}}^{auth} \leq (\frac{e\sigma_{\mathcal{E}}}{\rho 2^r})^{\rho} \cdot \frac{2^r}{\sqrt{\rho}} + \frac{(\sigma_{\mathcal{E}}+\sigma_{\mathcal{D}})(\rho-2)}{2^{c+v/2+1}} + \frac{q_{\mathcal{D}}}{2^t} \qquad (6.3)$$

$\square$

In the TinyJAMBU authenticated cipher, $b = 128$, $r = v = 32$, $c = 64$ and $\sigma \leq 2^{48}$. For authenticity security, we choose $\rho = 2^{17.445}$, after substituting the values of r and σ, the security bound for authenticity is $\mathbf{Adv}_{\text{TinyJAMBU}}^{auth} \leq 2^{-387} + 2^{17.445} \cdot (\sigma_{\mathcal{E}} + \sigma_{\mathcal{D}})/2^{81} + (q_{\mathcal{E}} + q_{\mathcal{D}})/2^{64}$. When 2^{48} blocks are queried, the adversary has the advantage of less than $2^{-15.5}$ to forge a message successfully.

Note that when nonce is reused, if Duplex mode is used for the same parameters of block size and rate size, when 2^{48} blocks are queried, the adversary has the advantage around 2^{-1} to forge a message successfully.

Chapter 7

Security Analysis

7.1 Properties of the TinyJAMBU Mode

In an authenticated encryption mode, the probability of state collision plays the role for protecting the confidentiality of plaintext and for resisting the forgery attack on the message. The security proof of the TinyJAMBU mode was given in Chapter 5. In the following, we provide a simple analysis on the state collision of the TinyJAMBU mode.

7.1.1 State collision for unique nonce

For unique nonce, a new state is generated for each message. It is impossible to apply the adaptively chosen message attack on the TinyJAMBU mode when the nonce is unique since each message is processed using a new initial state.

Only the state size affects state collision when the nonce is unique. When we use the 128-bit state in the TinyJAMBU mode, the probability is $\binom{2^{48}}{2} \times 2^{-128} = 2^{-33}$ when a key is used to process 2^{50} message bytes (there are about 2^{48} message blocks if assume that each message is at least 8 bytes long).

7.1.2 State collision for repeated nonce

For repeated nonce, the adaptively chosen messages can be used to improve the probability of state collision.

In TinyJAMBU, after processing a plaintext block, 32 bits of the ith state are known to the attacker as keystream (these 32 state bits are denoted as U^i_{32}); 32 bits of the state are unknown to the attacker, but the attacker is able to modify these 32 bits through message injection (these 32 state bits are denoted as V^i_{32}); 64 bits of the state are unknown to the attacker (these 64 state bits are denoted as W^i_{64}).

When 2^{48} message blocks are processed, the multicollisions at U^i_{32} can be observed. On average, each 32-bit value of U^i_{32} appears 2^{16} times.

We consider 2^{16} states with the same value at U^i_{32}, and inject 2^{16} message blocks to V^i_{32} for each state to introduce collision at V^i_{32} (so $2^{16} \times 2^{16} = 2^{32}$ message blocks are used). The probability that each V^i_{32} with 2^{16} injected message blocks has a collision with V^j_{32} with 2^{16} injected message blocks is about 1. So there are about $\binom{2^{16}}{2} = 2^{31}$ state pairs which are the same at $U^i_{32} \parallel V^i_{32}$. The chance that there is a whole state collision is about $2^{31} \times 2^{-64} = 2^{-33}$ after the message injections. We need to inject 2^{32} message blocks here to carry out this attack.

Since the attacker is able to adaptively choose up to 2^{48} message blocks, the attacker is able to perform the above attack for 2^{16} times. Thus the success rate of the whole state collision is $2^{-33} \times 2^{16} = 2^{-17}$.

Experiment 1. We implemented the above attack using 48-bit state size instead of the 128-bit state size. The U and V are 12 bits, W is 24 bits. The attacker can use 2^{18} chosen message blocks. We repeated the attack for 2^{15} times, the success rate for a state collision is $2^{-7.02}$. (The experimental result matches very well with the theoretical result, which is 2^{-7} for the 48-bit state version.)

7.2 Properties of the Keyed Permutation P_n

In this section, we analyse the differential, linear and algebraic properties of the keyed permutation P_n. We analyse the differential and linear properties using the Mixed Integer Linear Programming (MILP) [16]. The Gurobi optimizer [10] is used to find the best differential and linear probability bounds.

7.2.1 Differential properties of the keyed permutation P_n

To evaluate the security against differential forgery attack and differential key recovery attack, we analyse the differential properties [3, 5] of the TinyJAMBU permutation P_n. The following type of differences is analysed, and the largest differential probabilities are summarized in Table 7.1.

- Type 1. Input differences at $S_{96\cdots127}$, output differences at $S_{0\cdots127}$

Remarks. Note that in the previous TinyJAMBU submissions, we considered four types of differential properties. In this submission, P_{640} is used to replace P_{384} so as to protect the nonce and associated data with larger security margin. P_{640} is a strong permutation for 32-bit input and 32-bit output, so analysing Type 1 difference is sufficient for showing that TinyJAMU has large security margin against differential forgery attack and differential key recovery attack. The security analysis of TinyJAMBU is significantly simplified with the use of P_{640}.

Table 7.1: Type 1 Differential Properties of P_n

Round	Probability	Method
128	2^{-6}	MILP
256	2^{-20}	MILP
384	2^{-41}	MILP
512	2^{-64}	MILP
640	2^{-88}	MILP

Experiment 2. Gurobi is applied to find the optimal differential probabilities in Table 7.1. Gurobi also gives the optimal differential paths. In the experiment, we used differential pairs (with a random key for each pair) to verify the probabilities of the best differential paths of P_{128}, P_{256}, P_{384}. Some experiment results are given below.

- P_{128}
  ```
  input difference  = 0x80000000000000000000000000000000
  output difference = 0x80004000000102000000001000000000
  ```
  Gurobi differential probability: 2^{-6}
  Experiment differential probability: $2^{-6.00}$ with 2^{20} pairs

- P_{256}
  ```
  input difference  = 0x80000000000000000000000000000000
  output difference = 0x8004001020000240010000204080000
  ```
  Gurobi differential probability: 2^{-20}
  Experiment differential probability: $2^{-19.55}$ with 2^{30} pairs

- P_{384}
  ```
  input difference  = 0x20000000000000000000000000000000
  output difference = 0x20410000090040000008102480020001
  ```
  Gurobi differential probability: 2^{-41}
  Experiment differential probability: $2^{-40.79}$ with 2^{47} pairs

The above experiment results show that the Gurobi differential probabilities are close to the experimental differential probabilities.

Experiment 3. The differential probabilities of P_{512} and P_{640} in Table 7.1 are too small, so it is impossible to verify the probabilities of the differential paths directly in experiment. We use Gurobi to find the multiple differential paths for fixed input and output differences, then sum up the probabilities of the differential paths. The results for some fixed input differences and output differences are given below.

- P_{512}

  ```
  input difference  = 0x80000000000000000000000000000000
  output difference = 0x01028110000110080100009040080000
  ```
  Gurobi best differential probability: 2^{-64}
  Differential probability with multiple paths: $2^{-63.7}$

- P_{640}

  ```
  input difference  = 0x80000000000000000000000000000000
  output difference = 0x90004800040102400100081200000000
  ```
  Gurobi best differential probability: 2^{-88}
  Differential probability with multiple paths: $2^{-84.6}$

The above results show that the differential probabilities of P_{512} and P_{640} are affected by mutiple differential paths. The optimal differential probability of P_{640} is increased by about 10 times.

Considering the results of Experiment 2 and 3, the actual differential probabilities need to be adjusted slightly. The differential probabilities in Table 7.1 are adjusted as follows.

Table 7.2: The adjusted Type 1 differential Properties of P_n

Round	Probability
128	2^{-6}
256	2^{-19}
384	2^{-40}
512	2^{-62}
640	2^{-83}

7.2.2 Linear properties of the keyed permutation P_n

In this section, we analyse the linear properties [13, 14] of the TinyJAMBU permutation P_n. The linear bias being analysed is for arbitrary input bits and output bits at $S_{64...95}$. The best linear bias of the permutations are summarized in Table 7.3. The linear bias given in Table 7.3 are sufficient for demonstrating that TinyJAMBU has large security margin against the linear cryptanalysis since the linear attack involves 1024-round permutation. (Note that the linear bias that can be exploited in the attack is much smaller than the linear bias in Table 7.3, because the attacker can only read 32 input bits, write 32 input bits, and read 32 ouput bits of the permutation.)

Table 7.3: Linear bias of P_n

Round	Bias
256	2^{-11}
384	2^{-23}
512	2^{-30}

7.2.3 Algebraic properties of the keyed permutation P_n

We consider the algebraic property for the input bits at $S_{96\dots127}$. Our experiment shows that after 512 rounds, every output bit at $S_{64\dots95}$ is affected by the 32-bit input cube tester [9] at $S_{96\dots127}$.

7.3 Differential Forgery Attacks

It was analysed in Chapter 5 and Sect. 7.1.2 that the TinyJAMBU mode provides strong authentication security even when nonce is repeated. The attacker has the advantage of 2^{-17} to forge a message with 2^{48} adaptively chosen message blocks and repeated nonces. Note that an attacker always has the advantage of 2^{-16} to successfully forge a message using 2^{48} random trials since the tag size is 64 bits. Thus an adversary has advantage of $2^{-17} + 2^{-16} < 2^{-15}$ to forge a message successfully when less than 2^{50} bytes of data are authenticated using the same key when the nonce is repeated (under the assumption that each message is at least 8 bytes).

In the above analysis, it is assumed that the underlying permutation is perfect. In this section, we analyse the differential forgery attacks on the concrete permutations. To generate a collision, we inject difference into nonce, associated data or plaintext/ciphertext, then eliminate the difference in the state later.

7.3.1 Differential forgery attacks on nonce and associated data

In TinyJAMBU, each 32-bit nonce block and associated data block is processed using P_{640} (with different Framebits for nonce and associated data). The associated data also plays the role of nonce in TinyJAMBU. In the following, we only consider the forgery attacks on associated data (the attack on the nonce is identical).

When there is a difference at the 32-bit associated data block ad_i, there is input difference at $s_{96\dots127}$ of the permutation P_n. This difference must pass through at least one permutation P_{640}. According to Table 7.2, the largest differential probability for this type of input difference of P_{640} is 2^{-83}. It means that the differential forgery attack on associated data succeeds with probability at most 2^{-83}. (The success rate of the differential forgery attack on associated

data is expected to be much smaller than 2^{-83} since a forgery attack requires the output difference appears at $s_{96\dots127}$, while there is no constrain on the output difference in Table 7.2).

7.3.2 Differential forgery attacks on plaintext/ciphertext

When an adversary introduces a difference to a plaintext block M_i or ciphertext block C_i, there is input difference at $s_{96\dots127}$ of the permutation P_{1024}. This difference must pass through at least 1024 rounds of the permutation. According to Table 7.2, the largest differential probability for this type of input difference of P_{640} is 2^{-83}. At least 1024 rounds are used to process a plaintext block, so the differential probability of a 1024-round permutation is much smaller than than 2^{-83}. It indicates that the differential forgery attack on plaintext/ciphertext succeeds with probability much smaller than 2^{-83}.

7.4 Key Recovery Attacks

In the section, we analyse the security of the key against differential, linear and algebraic attacks.

7.4.1 Differential key recovery attack

There are two possible approaches to launch differential attacks against the key of TinyJAMBU. One approach is to introduce difference into nonce and/or associated data, then eliminate the difference in the state using the difference in the subsequent nonce and/or associated data blocks. Another approach is to introduce difference into plaintext or ciphertext when the nonce is reused, then observe the difference in the keystream.

Differential key recovery attack using nonce/associated data. The differential key recovery attack using nonce and associated data relies on eliminating the difference in the state. According to Section 7.3.1, the differential forgery attack on nonce and associated data succeeds with probability at most 2^{-83} (the actual success rate is expected to be much smaller than 2^{-83}). Since each key is used to protect at most 2^{50} bytes of data, there are at most 2^{47} messages if each message is at least 8 bytes. The attacker does not have enough messages to launch this differential attack successfully.

Differential key recovery attack using plaintext/ciphertext. When nonce is repeated for the same key, a difference can be injected into the state at $s_{96\dots127}$ through plaintext/ciphertext block, then the output difference can be observed in the next keystream block. According to Table 7.2, the differential probability is at most 2^{-83} for 640 rounds of the permutation. In TinyJAMBU, at least 1024 rounds are used to encrypt a 32-bit message block, so the differential probability for 1024 rounds is much smaller than 2^{-83}. There are at most

2^{48} 32-bit plaintext blocks being processed, so the attacker does not have enough message blocks to carry out the differential key recovery attack successfully.

7.4.2 Linear key recovery attack

When nonce is misused, an attacker can try to find the linear relation between the input bits and the output bits $s_{64\ldots95}$ of P_n.

According to Table 7.3, the linear bias is at most 2^{-30} for 512 rounds. This linear bias is much smaller than 2^{-16} (the message block size is 32 bits). At least 1024 rounds are used to process a plaintext block, so it is impossible to recover the key of TinyJAMBU using the linear cryptanalysis.

7.4.3 Algebraic attacks

We experimentally tested the number of rounds for each output bit being affected by 32-bit cube for the input. The experimental results show that after 512 rounds, every output bits is affected by the 32-bit cube tester. Hence, we believe that the 1024-round permutation provides large security margin against the algebraic cryptanalysis since the message block size is 32 bits.

7.5 Slide attack

The slide attack [4] is an effective tool to analyse the cipher with self-similarity round functions. Although TinyJAMBU permutation has the sliding property, the frame bits being added to the state will prevent the slide attack since the position of the frame bits is fixed. For example, for two related keys with slide property, the slide property between the two states gets eliminated with the introduction of Framebits of nonce.

7.6 Third-party security analysis

In [19], Saha et al. analyzed the security margin of the nonce and associated data of TinyJAMBU against the differential forgery attack. In the original Tiny-JAMBU design, our analysis shows that the differential forgery attack against nonce and associated data succeeds with probability probability at most 2^{-73}. In [19], it is shown that some NAND gates in the differential trail are not independent, so the forgery attack against nonce and associated data succeeds with probability $2^{-70.64}$. In TinyJAMBU v2, we increased the round number of the permutation being used to process nonce and associated data. TinyJAMBU v2 has large security margin against the differential forgery attack.

Saha et al. also analyzed the linear property of the permutation of Tiny-JAMBU [19], which is close to our analysis. TinyJAMBU has huge security margin against the linear cryptanalysis.

Chapter 8

The Performance of TinyJAMBU

8.1 Hardware Performance

To evaluate the hardware performance of TinyJAMBU, we implemented TinyJAMBU-128 in VHDL using the NIST LWC Hardware API, which was developed by Kaps et al. from George Mason University [12]. We synthesis our implementations with the Synopsys Design Compiler for an ASIC using 180 nm UMC technology. The results are summarized in Table 8.1. In Table 8.1, the area of the cipher core excludes the area of API.

Table 8.1: TinyJAMBU-128 v2 on ASIC 180nm UMC Technology (* indicates that a fixed key is embedded in the device)

Rounds per clock cycle	Area (GE)	Crypto Core Area	Throughput AD (Mbps)	Throughput Plaintext (Mbps)
8 rounds*	3223	1610	216	135
32 rounds*	4430	2296	737	460
8 rounds	4352	2708	206	128
32 rounds	5527	3328	698	436
128 rounds	6688	4026	3422	2139

Key schedule is not used in TinyJAMBU. If a fixed key is embedded in the device, about 1100 GE can be reduced in the implementation of TinyJAMBU-128, as shown in Table 8.1.

Gaj et al. from George Mason University gave the comprehensive bench-

marking of the lightweight AEAD ciphers on FPGA [1, 15]. Zidaric and Aagaard provided the detailed benchmarking of the lightweight AEAD ciphers on ASIC [23]. The area of TinyJAMBU is much smaller than all the other candidates in these benchmarkings.

8.2   Software Performance

At the NIST Lightweight Cryptography Workshop 2020, Çalık, Hasan and Kang implemented the round 2 candidates on various microcontrollers [7]. Based on their results, the code size of TinyJAMBU v2 is expected to be 872 bytes on ARM Cortex-M4F, which is the smallest among the round 2 candidates. The code size of TinyJAMBU v2 is expected to be identical to that of the original TinyJAMBU since we only increased the round number of one permutation. The state of TinyJAMBU is small, so the state can be updated efficiently to allow fast processing of short messages. The speed of the original tinyJAMBU-128 on ARM Cortex-M4F is 159 cycles/byte for 2048B plaintext. The estimated speed of TinyJAMBU v2 on ARM Cortex-M4F are given in Table 8.2.

Table 8.2: The estimated speed (clock cycles/byte) of TinyJAMBU-128 v2 on Microcontroller ARM Cortex-M4F (derived from [7])

Input Size	Speed	Input Size	Speed
16B AD	275	16B Plaintext	394
32B AD	187	32B Plaintext	276
128B AD	121	128B Plaintext	187

In addition to the microcontroller performance data presented at NIST Lightweight Cryptography Workshop 2020 [7], Weatherley [21] and Renner et al. [18] also give the benchmarking of the lightweight AEAD ciphers on microcontrollers.

Chapter 9

Features

- Lightweight, efficient authenticated encryption mode. It is lightweight since TinyJAMBU mode is based on a 128-bit keyed permutation. It is efficient since the message block size is 32 bits for 128-bit permutation.

- Strong protection of the secret key when nonce is reused. When nonce is reused in TinyJAMBU, and each key is used to process less than 2^{50} bytes of data, the secret key cannot be recovered with less than 2^{112} computations.

- Strong authentication security when nonce is reused. When nonce is reused in TinyJAMBU, and each key is used to process less than 2^{50} bytes of data, the attacker's advantage of forgery is less than 2^{-15}.

- The associated data is part of the nonce in TinyJAMBU, i.e., the combination of nonce and associated data is the effective nonce of the cipher.

- Lightweight Permutation. The keyed permutation is based on a 128-bit nonlinear feedback shift register with only 5 taps. It is thus efficient to implement the keyed permutation in hardware.

- No key schedule in the keyed permutation, so the hardware area is significantly reduced when a constant fixed key is used in TinyJAMBU (or when a key must be stored in the devices for protecting multiple messages).

- Lightweight Input Loading. The nonce, associated data and the plaintext/ciphertext can be loaded into the state bit-by-bit when the nonlinear feedback shift register of the keyed permutation is clocked. It is thus efficient to load the input into the cipher in hardware.

- Parallel Computation. In TinyJAMBU, 32 steps can be computed in parallel. This parallel feature benefits fast hardware and software implementations.

Chapter 10

Design Rationale

- Design goal

  We aim to design a lightweight authenticated encryption algorithm which is optimized for the devices in which a secret key is stored. In the applications in which a secret key is used to protect multiple messages, it is reasonable to store a secret key in the devices, either store the key temporarily as a session key or store the key permanently as a fixed key.

  When a key is already stored in the device, the state of the cipher could be very small since we can use the keyed permutation to prevent an attacker from computing the states offline and launching the state collision attacks using the computed states.

- Design a nonce-misuse resistant cipher

  Since we aim to design a cipher for a stored key (especially a fixed key), we design the cipher to resist the key recovery attack when the attacker is able to control the nonce, associated data and plaintext.

  We also design the cipher to provide strong authentication security when nonce is misused.

  To provide better nonce-misuse resistance, the associated data also plays the role of nonce. In case when nonce is repeated due to some error, but associated data is different, the security of the cipher would not be affected by the repeated nonce.

  Feeding the plaintext twice into the cipher (one for tag generation, another for encryption using the tag as part of the nonce) can provide high security for plaintext when nonce is misused. But such cost is high for a lightweight cipher, so we do not provide strong protection of plaintext when both nonce and associated data are repeated.

- Design the TinyJAMBU authenticated encryption mode

  JAMBU mode is the smallest block cipher authenticated encryption mode in the CAESAR competition. The message block size and the state size

of the JAMBU mode is 0.5 and 1.5 times of the block size of the block cipher, respectively.

If we reduce the message block size in the JAMBU mode, the state size of JAMBU can be reduced significantly, and we can get better authentication security when nonce is reused. So we designed TinyJAMBU which is a small variant of JAMBU.

In the TinyJAMBU mode, the attacker cannot control the state bits easily, so when nonce is reused, TinyJAMBU mode achieves better authentication security than the Duplex mode for the same keyed permutation and the same message block size. In th Duplex mode, part of the state can be set to any arbitrary value when nonce is reused. So we do not use he Duplex mode in our design.

- Different processing of associated data and plaintext

  In TinyJAMBU, we use a strong permutation for encrypting plaintext, while we use a relatively weak permutation for processing associated data for better efficiency. The reason is that when we are processing the associated data, there is no information leakage as long as there is no state collision. The permutation for associated data only needs to resist the forgery attack, while the permutation for plaintext should resist not only the forgery attack, but also the the key recovery attack when the nonce is reused.

- Design the keyed permutation

  The keyed permutation is based on a simple nonlinear feedback shift register with only 5 taps. There are three reasons for using the nonlinear feedback shift register to update the state:

  - The hardware cost of nonlinear feedback shift register is low.
  - A number of steps of the nonlinear feedback shift registers can be computed in parallel for efficient hardware and software implementation.
  - The stream input data can be easily loaded into the state when the state gets updated.

  There is slide property in the keyed permutation based on a feedback shift register. We use the FrameBits for each permutation to destroy the slide property.

  In the keyed permutation, there is no key schedule (the round keys are simply generated by repeating the secret key). It is to reduce the hardware cost of key schedule when there is an embedded key in the device. This approach of designing lightweight block cipher has been been used in the design of KTANTAN [8].

- Design the nonlinear feedback shift register

  We use a 128-bit nonlinear feedback shift register with 5 taps to update the state of the permutation. The number of taps is not large since we aim at lightweight hardware implementation; the number of taps is not too small so that we can get reasonably good differential and linear characteristics.

  Among those five taps, two taps are NANDed together (nonlinear feedback), the other three taps are for linear feedback.

  To design a nonlinear feedback shift register with 5 taps, we need to choose four tap positions.

  – The tap positions are chosen to ensure that 32 steps can be computed in parallel.

  – There are 15 tap distances for the feedback register with 5 taps. The tap positions are chosen to ensure that most of those 15 tap distances are coprime to each other. When two tap distances are not coprime, the greatest common divisor should be small.

  – After the above filtering of tap positions, we choose the tap positions with fast diffusion (measured as the number of steps being used for one state bit affecting the whole state). We keep the top 20 sets of tap positions with fast diffusion.

  – After the above filtering, we tested the differential and linear property of the permutation (for 32-bit message block). We choose the tap positions that gives excellent differential and linear characteristics.

Bibliography

[1] ATHENA project of George Mason University. Hardware Benchmarking of Lightweight Cryptography. Available at `https://cryptography.gmu.edu/athena/index.php?id=LWC`

[2] Guido Bertoni, Joan Daemen, Michael Peeters, and Gilles Van Assche. Duplexing the sponge: Single-pass authenticated encryption and other applications. In *Selected Areas in Cryptography – SAC 2011*, pages 320–337.

[3] Eli Biham and Adi Shamir. Differential Cryptanalysis of DES-like Cryptosystems. *Journal of Cryptology*, 4(1):3–72, 1991.

[4] Alex Biryukov and David Wagner. Slide Attacks. International Workshop on Fast Software Encryption – FSE'99, pp. 245—259.

[5] Eli Biham and Adi Shamir. *Differential Cryptanalysis of the Data Encryption Standard*. Springer-Verlag, London, UK, 1993.

[6] CAESAR: Competition for Authenticated Encryption: Security, Applicability, and Robustness. Available at `https://competitions.cr.yp.to/caesar-submissions.html`

[7] Çağdaş Çalık, Munawar Hasan, and Jinkeon Kang. Benchmarking Round 2 Candidates on Microcontrollers. *NIST Lightweight Cryptography Workshop 2020*. Available at `https://csrc.nist.gov/CSRC/media/Presentations/benchmarking-round-2-candidates/images-media/session-1-calik-benchmarking-second-round-cadidates.pdf`

[8] Christophe De Cannière, Orr Dunkelman, Miroslav Knežević. KATAN and KTANTAN — A Family of Small and Efficient Hardware-Oriented Block Ciphers. International Workshop on Cryptographic Hardware and Embedded Systems – CHES 2009, pp 272–288.

[9] Itai Dinur and Adi Shamir. Cube attacks on tweakable black box polynomials. In Antoine Joux, editor, *Advances in Cryptology - EUROCRYPT 2009*, pages 278–299, 2009. Springer Berlin Heidelberg.

[10] Gurobi Optimizer. Available at `http://www.gurobi.com/`

[11] Philipp Jovanovic, Atul Luykx, Bart Mennink, Yu Sasaki, and Kan Yasuda. Beyond conventional security in sponge-based authenticated encryption modes. *Journal of Cryptology*, Jun 2018.

[12] Jens-Peter Kaps, William Diehl, Michael Tempelmeier, Ekawat Homsirikamol, and Kris Gaj Hardware API for Lightweight Cryptography. `https://cryptography.gmu.edu/athena/LWC/LWC_HW_API.pdf`

[13] Mitsuru Matsui. Linear Cryptanalysis Method for DES cipher. In *Advances in Cryptology–EUROCRYPT'93*, pages 386–397. Springer, 1994.

[14] Mitsuru Matsui and Atsuhiro Yamagishi. A New Method for Known Plaintext Attack of FEAL Cipher. In *Advances in Cryptology – EUROCRYPT'92*, pages 81–91. Springer, 1993.

[15] Kamyar Mohajerani, Richard Haeussler, Rishub Nagpal, Farnoud Farahmand, Abubakr Abdulgadir, Jens-Peter Kaps, and Kris Gaj. FPGA Benchmarking of Round 2 Candidates in the NIST Lightweight Cryptography Standardization Process: Methodology, Metrics, Tools, and Results. Available at `https://eprint.iacr.org/2020/1207`

[16] Nicky Mouha, Qingju Wang, Dawu Gu, Bart Preneel. Differential and linear cryptanalysis using Mixed-Integer Linear Programming. *Information security and cryptology – Inscrypt 2011*, pages 57–76.

[17] Yusuke Naito, Mitsuru Matsui, Takeshi Sugawara, and Daisuke Suzuki. SAEB: A lightweight blockcipher-based AEAD mode of operation. *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, 2018(2):192–217, 2018.

[18] Sebastian Renner, Enrico Pozzobon, and Jurgen Mottok. NIST LWC Software Performance Benchmarks on Microcontrollers. Available at `https://lwc.las3.de/`

[19] Dhiman Saha, Yu Sasaki, Danping Shi, Ferdinand Sibleyras, Siwei Sun, and Yingjie Zhang (2020). On the Security Margin of TinyJAMBU with Refined Differential and Linear Cryptanalysis. IACR Transactions on Symmetric Cryptology, 2020(3), 152-174. Available at `https://tosc.iacr.org/index.php/ToSC/article/view/8699/8291`.

[20] Dhiman Saha, Yu Sasaki, Danping Shi, Ferdinand Sibleyras, Siwei Sun, and Yingjie Zhang. The MILP code pertaining to the paper "On the Security Margin of TinyJAMBU with Refined Differential and Linear Cryptanalysis". Available at `https://github.com/c-i-p-h-e-r/refinedTrailsTinyJambu`

[21] Rhys Weatherley. Microcontroller Benchmarking of Lightweight Cryptography Primitives. Available at `https://rweather.github.io/lightweight-crypto/`

[22] Hongjun Wu and Tao Huang. JAMBU Lightweight Authenticated Encryption Mode and AES-JAMBU. *NIST Lightweight Cryptography Workshop 2015*. Available at `https://www.nist.gov/news-events/events/2015/07/lightweight-cryptography-workshop-2015`

[23] Nusa D. Zidaric and Mark Aagaard. ASIC Benchmarking of Round 2 Candidates in the NIST Lightweight Cryptography Standardization Process. Available at `https://eprint.iacr.org/2021/049`

Appendix A

Analysis of the Probability for Equation 6.2

Here we show how to derive the maximum probability for Equation 6.2 with minimum required message blocks. The equation is given below.

$$\Pr[Coll\mathcal{E} \text{ in } \Theta_x] = \sum_{i \neq j, 1 \leq i,j \leq \theta} \min\{\frac{\beta_i \beta_j}{2^v}, 1\} \cdot \frac{1}{2^c}$$

First, we ignore the min function, so the right hand side becomes:

$$\sum_{i \neq j, 1 \leq i,j \leq \theta} \frac{\beta_i \beta_j}{2^v} \cdot \frac{1}{2^c} = \frac{1}{2^{c+v}} \sum_{i \neq j, 1 \leq i,j \leq \theta} \beta_i \beta_j$$

Let $z = \sum_{i \neq j, 1 \leq i,j \leq \theta} \beta_i \beta_j$, we have

$$
\begin{aligned}
z &= \sum_{i \neq j, 1 \leq i,j \leq \theta} \beta_i \beta_j \\
&\leq \sum_{i \neq j, 1 \leq i,j \leq \theta} (\frac{\beta_i^2 + \beta_j^2}{2}) \\
&= \frac{(\theta - 1)(\sum_{i=1}^{\theta} \beta_i^2)}{2}.
\end{aligned}
$$

Then,

$$\sum_{i=1}^{\theta} \beta_i^2 \geq \frac{2z}{\theta - 1}$$

Add $2z$ to both size, we have

$$(\sum_{i=1}^{\theta} \beta_i)^2 = \sum_{i=1}^{\theta} \beta_i^2 + 2z \geq \frac{2z}{\theta - 1} + 2z.$$

Therefore,

$$z \leq \frac{(\theta - 1)}{2\theta} (\sum_{i=1}^{\theta} \beta_i)^2. \tag{A.1}$$

The equality holds when β_i are equal for all $1 \leq i \leq \theta$.

Now we take the min function into consideration. When $\beta_i = \beta_j \geq 2^{v/2}$, $\min\{\frac{\beta_i \beta_j}{2^v}, 1\} = 1$. Then, the Equation 6.2 obtains the maximum value which is $\sum_{i \neq j, 1 \leq i, j \leq \theta} \frac{1}{2^c} = \binom{\theta}{2} \cdot \frac{1}{2^c}$.

On the other hand, when $\beta_i = \beta_j < 2^{v/2}$, $\min\{\frac{\beta_i \beta_j}{2^v}, 1\} = \frac{\beta_i \beta_j}{2^v} < 1$, the probability cannot reach the maximum value. Moreover, from inequality A.1, z is related to the $(\sum_{i=1}^{\theta} \beta_i)^2$. Any decrease in $\sum_{i=1}^{\theta} \beta_i$ will lead to the overall probability reduce even more.

Therefore, the $\beta_i = \beta_j \geq 2^{v/2}$ will obtain the maximum value of probability with least number of message blocks $\sum_{i=1}^{\theta} \beta_i$. Any decrease in the choice of $\sum_{i=1}^{\theta} \beta_i$ will decrease the probability in a squared manner.