AEZ v3: Authenticated Encryption by Enciphering

Viet Tung Hoang
University of Maryland
Georgetown University
tvhoang@umd.edu

Ted Krovetz
Sacramento State
ted@krovetz.net

Phillip Rogaway
UC Davis
ETH Zürich
rogaway@cs.ucdavis.edu

August 22, 2014
last edit: October 4, 2014

The named authors are both designers and submitters.

Abstract

AEZ encrypts by appending to the plaintext a fixed authentication block and then enciphering the resulting string with an arbitrary-input-length blockcipher, this tweaked by the nonce and AD. The approach results in strong security and usability properties, including nonce-reuse misuse resistance, automatic exploitation of decryption-verified redundancy, and arbitrary, user-selectable ciphertext expansion. AEZ is parallelizable and its computational cost is close to that of AES-CTR. On a recent Intel processor (Haswell), our C implementation achieves a peak speed of about 0.7 cpb.

The latest version of this document, and all related materials, can always be found on the AEZ homepage: http://www.cs.ucdavis.edu/~rogaway/aez

Contents

0 Introduction **1**

1 Specification **3**

1.1 Notation 3

1.2 Arguments and Parameters 3

1.3 AEZ Extensions 4

1.4 Pseudocode 4

1.5 Usage cap 9

2 Security Goals **9**

3 Security Analysis **13**

4 Features **15**

5 Design Rationale **18**

6 Intellectual Property **19**

7 Consent **19**

8 Revision History **19**

References **20**

0 Introduction

This document describes AEZ, which we view as both an enciphering scheme and an authenticated-encryption scheme. Before specifying it we provide a brief overview.

Authenticated encryption by enciphering. When we speak of an *enciphering scheme* we mean an object that is like a conventional blockcipher except that the plaintext's length is arbitrary and variable, and, additionally, there's a tweak. Regarding AEZ in this way, enciphering maps a key K, plaintext X, and tweak T to a ciphertext $Y = \text{Encipher}(K, T, X)$ having the same length as X. Going backwards, one can recover $X = \text{Decipher}(K, T, Y)$. The security property we seek is that of a tweakable, strong-PRP (pseudorandom permutation): for a random key K it should be hard to distinguish oracles $(\text{Encipher}(K, \cdot, \cdot), \text{Decipher}(K, \cdot, \cdot))$ from oracles $(\pi(\cdot, \cdot), \pi^{-1}(\cdot, \cdot))$ that realize a family of independent, uniformly random permutations and their inverse.

When we instead regard AEZ as an *authenticated-encryption* (AE) *scheme*, encryption maps key K, plaintext M, nonce N (also called a "public nonce" or "public message number"), associated data A, and an authenticator length ABYTES to a ciphertext $C = \text{Encrypt}(K, N, A, \tau, M)$ that is $\tau = 8 \cdot \text{ABYTES}$ bits longer than M. Calling $\text{Decrypt}(K, N, A, \tau, C)$ returns either a string M or an indication of invalidity. The security property we seek is that of a *robust* authenticated-encryption scheme, a new and very strong notion that implies protection of the privacy and authenticity of M and the authenticity of N and A, and must do so to the maximal extent possible even if nonces get reused ("misuse resistance" [37]), the authenticator length is small (including zero), or if, on decryption, invalid plaintexts get prematurely released.

Why speak of enciphering when CAESAR is a competition for AE schemes? Because an enciphering scheme determines an AE scheme by a simple and generic transformation—the *encode-then-encipher* method—and the AE scheme one gets in this way has attractive security and usability properties.

Encode-then-encipher encrypts the string M by enciphering a string X that encodes both M and a block of ABYTES zero bytes, doing so using a tweak T that encodes N, A, and ABYTES. Decryption works by deciphering the presented string (again using the tweak determined from N, A, ABYTES) and verifying the presence of the anticipated zero bytes. See Figure 1.

What are these "attractive security and usability properties" to which we allude? (1) If plaintexts are known *a priori* not to repeat, no nonce is needed to ensure semantic security. (2) If there's arbitrary redundancy in plaintexts whose presence is verified on decryption, this augments authenticity. (3) Any authenticator length can be selected, achieving best-possible authenticity for this amount of expansion. (4) Because of the last two properties, one can minimize length-expansion for low-energy or bandwidth-constrained applications. (5) If what's supposed to be a nonce should accidentally get repeated, the privacy loss is limited to revealing repetitions in (N, A, M) tuples, while authenticity is not damaged at all. (6) If a decrypting party leaks some or all of a putative plaintext that was supposed to be squelched because of an authenticity-check failure, this won't compromise privacy or authenticity.

The authors believe that the properties just enumerated would sometimes be worth a considerable computational price. Yet the overhead we pay is modest: AEZ is about as fast as OCB.

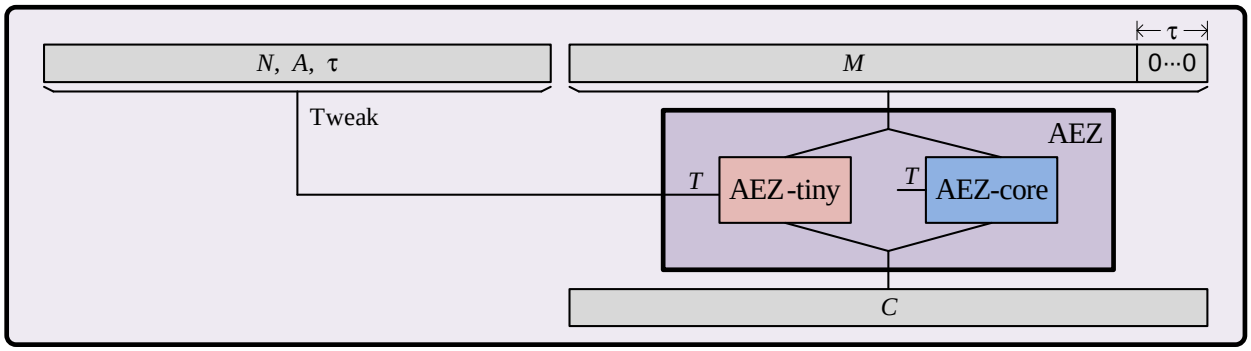

Figure 1: **High-level structure of AEZ**. After appending to the message a block of zeros we encipher it using a tweak comprising the nonce, associated data, and the ciphertext expansion $\tau = 8 \cdot \text{ABYTES}$. How this happens depends on the length of what's being enciphered: usually we use AEZ-core, but strings shorter than 32 bytes are enciphered by AEZ-tiny. Both depend on the underlying key K, which is not shown.

Realizing the enciphering. The way AEZ enciphers depends on the length of the plaintext. If it's fewer than 32 bytes we use AEZ-tiny, which builds on FFX [5, 12]. When it's 32 bytes or more we use AEZ-core, which builds on EME [15, 16] and OTR [24]. Again see Figure 1.

AEZ-tiny is a balanced-Feistel scheme. Its round function is based on AES4, a four-round version of AES. Guided by known attacks, more rounds are used for short strings than long ones.

AEZ-core resembles EME mode ("encipher-mask-encipher") but, for each of the two enciphering layers, consecutive pairs of blocks are processed together using a two-round Feistel network. The round function for this is based on AES4. The mask that is injected as the middle layer is determined, for each pair of blocks, using another AES4 call. The result is an enciphering scheme that employs five AES4 operations to encipher each consecutive pair of blocks, so 10 AES rounds per block. Thus our performance approaches that of AES-CTR.

The design of AEZ employs a paradigm we call *accelerated provable-security*. One begins by designing a cryptographic scheme based on a well-known tool: for AEZ, a tweakable blockcipher (TBC) [22]. One proves security assuming that the tool meets some standard assumption; here, that the TBC is secure as a tweakable PRP. Finally, the tool—our TBC—is selectively instantiated by a *scaled-down* primitive: we will mostly use AES4, a reduced-round version of AES. The *thesis* underlying this approach is that it can be instrumental in finding complex but highly efficient schemes; and that if the instantiation is done judiciously, then the scaled-down scheme retains at least some assurance benefit flowing from the approach.

The name. The name "AEZ" is not exactly an acronym. The AE prefix is meant to suggest *authenticated encryption* and the overlapping EZ suffix is meant to suggest *easy*, in the sense of ease of correct use. The AES-like name is also a nod to the fact that AEZ is based on AES and can likewise be considered a species of blockcipher. Finally, the name can be used to help identify individuals who can't distinguish an S from a Z.

1 Specification

1.1 Notation

Numbers and strings. A *number* means a nonnegative integer, $\mathbb{N} = \{0, 1, 2, \ldots\}$. A bit is 0 or 1 and a string is a finite sequence of bits. The length of a string X is written $|X|$. The empty string ε is the string of length zero. Concatenation of strings A and B is written AB or $A \parallel B$. When X is a string or a bit, X^n means X repeated n times; for example $0^3 = 000$ and $(01)^2 = 0101$. By $\mathcal{X}^*$ we denote the set of all strings over the alphabet $\mathcal{X}$, including ε. By $(\mathcal{X}^*)^*$ we denote the set of all vectors over $\mathcal{X}^*$, including the empty vector. The bitwise-and, bitwise-or, and bitwise-xor of strings A and B are denoted $A \wedge B$, $A \vee B$, and $A \oplus B$ respectively. For operations on unequal-length strings, first drop the necessary number of rightmost bits from the longer ($10 \oplus 0100 = 11$). For X a string, let $X0^* = X0^p$ with p the smallest number such that 128 divides $|X| + p$. If $|X| = n$ and $1 \le i \le j \le n$ then $X[i]$ is the ith bit of X (indexing from the left starting at 1), $\mathrm{msb}(X) = X[1]$, and $X[i..j] = X[i] \cdots X[j]$. Let $[n]_t$ be the t-bit string representing $n \bmod 2^t$ and let $[n]$ be shorthand for $[n]_8$; for example $[0]^{16} = ([0]_8)^{16} = 0^{128}$ and $[1]^{16} = (00000001)^{16}$. A byte is a string of eight bits. The set of all bytes is denoted $\textsc{Byte}$. A byte string is an element of $\textsc{Byte}^*$.

A block is 128 bits. Let $\mathbf{0} = 0^{128}$. If $X = a_1 \cdots a_{128}$ is a block ($a_i \in \{0,1\}$) then we define $X \ll 1 = a_2 \cdots a_{128} 0$. For $n \in \mathbb{N}$ and $X \in \{0,1\}^{128}$ define $n \cdot X$ by asserting that $0 \cdot X = \mathbf{0}$ and $1 \cdot X = X$ and $2 \cdot X = (X \ll 1) \oplus [135 \cdot \mathrm{msb}(X)]_{128}$ and $2n \cdot X = 2 \cdot (n \cdot X)$ and $(2n+1) \cdot X = (2n \cdot X) \oplus X$.

AES4 and AES10. We assume familiarity with AES. For $K, X \in \{0,1\}^{128}$ we write $\texttt{aesenc}(X, K)$ for a single round of AES: permute X by performing `SubBytes` then `ShiftRows` then `MixColumns`, then do an `AddRoundKey` with K. For $\boldsymbol{K} = (K_0, K_1, K_2, K_3, K_4)$ a list of five blocks let $\mathrm{AES4}_{\boldsymbol{K}}(X) = \mathrm{AES4}(\boldsymbol{K}, X)$ be

$$\texttt{aesenc}(\texttt{aesenc}(\texttt{aesenc}(\texttt{aesenc}(X \oplus K_0, K_1), K_2), K_3), K_4)\,.$$

For $\boldsymbol{K} = (K_0, K_1, \ldots, K_{10})$ a list of 11 blocks define $\mathrm{AES10}_{\boldsymbol{K}}(X) = \mathrm{AES10}(\boldsymbol{K}, X)$ like we defined AES4 but composed of ten rounds of `aesenc`. Note that we do *not* omit the final-round `MixColumns`, as does AES itself, for either AES4 or AES10.

1.2 Arguments and Parameters

By *parameter* we mean "a value on which AEZ encryption depends that, independent of any particular API, would usually to be held constant throughout some long-lived context." Under this interpretation of the word, AEZ has five arguments and one parameter. See Figure 2. In particular, we do not regard $\textsc{keybytes} = |K|$ as a parameter (we permit keys of any length), nor $\textsc{npubbytes} = |N|$ (we permit nonces to have varying lengths, even within a session). While these two values are omitted from the CAESAR-specified API, they could be specified in a different API.

The *authenticator length*, $\textsc{abytes}$, determines how much longer a ciphertext is than its plaintext. Arbitrary values are allowed, but values exceeding 16 are not expected to provide additional security. We do not insist that $\textsc{abytes}$ be held constant throughout a session; a user is free to change it with each encryption. Still, we expect most applications to fix $\textsc{abytes}$, and choose to regard it as a parameter. We will use $\tau = 8 \cdot \textsc{abytes}$ if we want to measure ciphertext expansion in bits.

symbol	comments
M	Plaintext. $M \in \text{BYTE}^*$
C	Ciphertext. $C \in \text{BYTE}^*$
K	Key. $K \in \text{BYTE}^*$. Will get turned into AEZ's internal keys. $\lvert K\rvert \geq 128$ recommended
N	Nonce (aka: public sequence number). $N \in \text{BYTE}^*$. $\lvert N\rvert \leq 128$ recommended
A	Associated data. $A \in (\text{BYTE}^*)^*$. String-valued AD is regarded as a one-element vector
ABYTES	Authenticator length. ABYTES $\in \mathbb{N}$. Default is 16. ABYTES ≤ 16 recommended.

Figure 2: **Arguments and parameters to AEZ**. One might consider ABYTES an argument *or* a parameter: its value is allowed to change during the uses of a key, but conventionally this would not be done.

An unusual aspect of AEZ encryption is that it permits vector-valued AD: $A \in (\text{BYTE}^*)^*$. To recover the usual setting (a string-valued AD) the user selects an AD A with a single component.

We provide a default value for the ABYTES parameter: ABYTES = 16. The only named parameter set, denoted `aez`, uses this value. A conforming AEZ implementation is free to select a different default. In any context where the key length or nonce length are *required* to be fixed, we select byte lengths for these of KEYBYTES = 16 and NPUBBYTES = 12.

1.3 AEZ Extensions

Early versions of AEZ included a parameter EXTNS, a string-valued *extensions directive*. The intent was that this would, in the future, unlock capabilities traditionally seen as outside the scope of an encryption scheme's functionality, including secret nonces (secret message numbers), plaintext-length obfuscation (via a specified padding regime), and encoding ciphertexts into a prescribed alphabet. These extensions are to be realized by a wrapper that keylessly transforms a plaintext, AEZ encrypts it, then keylessly transforms the result. We dropped the EXTNS parameter when we made the AD vector-valued, as the same effect can now be achieved using that feature. A document defining AEZ extensions will be released later.

1.4 Pseudocode

The definition of AEZ is provided in Figures 3 and 4. Let us explain some aspects of the pseudocode.

Encryption and decryption. To encrypt a string M we augment it with an *authenticator*—a block of ABYTES zero bytes—and encipher the resulting string, tweaking this with a tweak formed from A, N, and ABYTES. Next we encipher the augmented message. To decrypt a ciphertext C we reverse the process, verifying the presence of the all-zero authenticator.

Enciphering. Messages are enciphered by either of two methods. Strings of 1–31 bytes are enciphered using AEZ-tiny, while those of 32 bytes or more are enciphered using AEZ-core.

Roughly following FFX [5, 12], AEZ-tiny uses a balanced Feistel network. The number of rounds

```
100  algorithm Encrypt(K, N, A, τ, M)                          // AEZ authenticated encryption
101  X ← M ‖ 0^τ; (A_1, ..., A_m) ← A
102  T ← ([τ]_128, N, A_1, ..., A_m)
103  if M = ε then return AEZ-prf(K, T, τ)
104  return Encipher(K, T, X)
```

```
110  algorithm Decrypt(K, N, A, τ, C)                          // AEZ authenticated decryption
111  (A_1, ..., A_m) ← A;  T ← ([τ]_128, N, A_1, ..., A_m)
112  if |C| < τ then return ⊥
113  if |C| = τ then if C = AEZ-prf(K, T, τ) then return ε else return ⊥ fi fi
114  X ← Decipher(K, T, C)
115  M ‖ Z ← X where |Z| = τ
116  if (Z = 0^τ) then return M else return ⊥
```

```
200  algorithm Encipher(K, T, X)                               // AEZ enciphering
201  if |X| < 256 then return Encipher-AEZ-tiny(K, T, X)
202  if |X| ≥ 256 then return Encipher-AEZ-core(K, T, X)
```

```
210  algorithm Encipher-AEZ-tiny(K, T, M)                      // AEZ-tiny enciphering
211  m ← |M|;  n ← m/2;  Δ ← AEZ-hash(K, T)
212  if m = 8 then k ← 24 else if m = 16 then k ← 16 else if m < 128 then k ← 10 else k ← 8 fi
213  L ← M[1 .. n];  R ← M[n + 1 .. m];  if m ≥ 128 then j ← 6 else j ← 7 fi
214  for i ← 0 to k − 1 do R' ← L ⊕ ((E_K^{0,j}(Δ ⊕ R10* ⊕ [i]_128))[1 .. n]);  L ← R;  R ← R'  od
215  C ← R ‖ L;  if m < 128 then C ← C ⊕ (E_K^{0,3}(Δ ⊕ (C0* ∨ 10*)) ∧ 10*) fi
216  return C
```

```
220  algorithm Encipher-AEZ-core(K, T, M)                      // AEZ-core enciphering
221  Δ ← AEZ-hash(K, T)
222  M_1 M'_1 ··· M_m M'_m M_uv M_x M_y ← M where |M_1| = ··· = |M'_m| = |M_x| = |M_y| = 128 and |M_uv| < 256
223  d ← |M_uv|;  if d ≤ 127 then M_u ← M_uv; M_v ← ε else M_u ← M_uv[1..128]; M_v ← M_uv[129..|M_uv|] fi
224  for i ← 1 to m do W_i ← M_i ⊕ E_K^{1,i}(M'_i);  X_i ← M'_i ⊕ E_K^{0,0}(W_i) od
225  if d = 0 then X ← X_1 ⊕ ··· ⊕ X_m ⊕ 0 else if d ≤ 127 then X ← X_1 ⊕ ··· ⊕ X_m ⊕ E_K^{0,4}(M_u10*)
226  else X ← X_1 ⊕ ··· ⊕ X_m ⊕ E_K^{0,4}(M_u) ⊕ E_K^{0,5}(M_v10*) fi
227  S_x ← M_x ⊕ Δ ⊕ X ⊕ E_K^{0,1}(M_y);  S_y ← M_y ⊕ E_K^{-1,1}(S_x);  S ← S_x ⊕ S_y
228  for i←1 to m do S' ← E_K^{2,i}(S); Y_i ← W_i ⊕ S'; Z_i ← X_i ⊕ S'; C'_i ← Y_i ⊕ E_K^{0,0}(Z_i); C_i ← Z_i ⊕ E_K^{1,i}(C'_i) od
229  if d = 0 then C_u ← C_v ← ε;  Y ← Y_1 ⊕ ··· ⊕ Y_m ⊕ 0
230  else if d ≤ 127 then C_u ← M_u ⊕ E_K^{-1,4}(S);  C_v ← ε;  Y ← Y_1 ⊕ ··· ⊕ Y_m ⊕ E_K^{0,4}(C_u10*)
231  else C_u ← M_u ⊕ E_K^{-1,4}(S);  C_v ← M_v ⊕ E_K^{-1,5}(S);  Y ← Y_1 ⊕ ··· ⊕ Y_m ⊕ E_K^{0,4}(C_u) ⊕ E_K^{0,5}(C_v10*) fi
232  C_y ← S_x ⊕ E_K^{-1,2}(S_y);  C_x ← S_y ⊕ Δ ⊕ Y ⊕ E_K^{0,2}(C_y)
233  return C_1 C'_1 ··· C_m C'_m C_u C_v C_x C_y
```

Figure 3: **AEZ authenticated-encryption: main routines.** The tweakable blockcipher E, the hash AEZ-hash, and the PRF AEZ-prf are all defined in Figure 4. Requested ciphertext expansion is $\tau = 8 \cdot \text{ABYTES}$ bits. Algorithm Decipher(K, T, C), not shown, returns the unique M such that Encipher$(K, T, M) = C$. See the accompanying text for how this is computed.

```
algorithm AEZ-hash(K, T)                                   // AXU hash. T is a vector of strings
(T_1, ..., T_t) ← T
for i ← 1 to t do
    m ← max(1, ⌈|T_i|/128⌉);  X_1···X_m ← T_i where |X_1| = ··· = |X_{m−1}| = 128
    if |X_m| = 128 then Δ_i ← E_K^{2+i,1}(X_1) ⊕ ··· ⊕ E_K^{2+i,m}(X_m)
    if |X_m| < 128 then Δ_i ← E_K^{2+i,1}(X_1) ⊕ ··· ⊕ E_K^{2+i,m−1}(X_{m−1}) ⊕ E_K^{2+i,0}(X_m 10*)
return Δ_1 ⊕ ··· ⊕ Δ_t ⊕ 0
```

```
algorithm AEZ-prf(K, T, τ)                                 // PRF used when M = ε
Δ ← AEZ-hash(K, T)
return (E_K^{−1,3}(Δ) ‖ E_K^{−1,3}(Δ⊕[1]_128) ‖ E_K^{−1,3}(Δ⊕[2]_128) ‖ ···)[1..τ]
```

```
algorithm E_K^{i,j}(X)                                     // Scaled-down TBC
(I, J, L) ← Extract(K)
k_0 ← (0, I, J, L, 0);  k_1 ← (0, J, L, I, 0);  k_2 ← (0, L, I, J, I)
K ← (0, I, L, J, I, L, J, I, L, J, I)
if i = −1 and 0 ≤ j ≤ 7 then return AES10_K(X ⊕ jJ)
if i = 0 and 0 ≤ j ≤ 7 then return AES4_{k_0}(X ⊕ jJ)
if 1 ≤ i ≤ 2 and j ≥ 1 then return AES4_{k_i}(X ⊕ (j mod 8)J ⊕ 2^{⌊(j−1)/8⌋}L)
if i ≥ 3 and j ≥ 1 then return AES4_{k_0}(X ⊕ (j mod 8)J ⊕ 2^{⌊(j−1)/8⌋} · L ⊕ (i − 2)8J)
if i ≥ 3 and j = 0 then return AES4_{k_0}(X ⊕ (i − 2)8J)
```

```
algorithm Extract(K)                                       // Map key to subkeys
z ← [0][1]···[15];  m ← max(1, ⌈|K|/128⌉);  K_1···K_m ← K where |K_1| = ··· = |K_{m−1}| = 128
for i ← 1 to 3 do
   for j ← 0 to m do c_{i,j} ← AES4_{(z,z,z,z,z)}([i]_64[j]_64);  C_{i,j} ← (c_{i,j}, c_{i,j}, c_{i,j}, c_{i,j}, c_{i,j});  od
   if |K_m|=128 then X_i ← AES4_{C_{i,1}}(K_1)⊕AES4_{C_{i,2}}(K_2)⊕···⊕AES4_{C_{i,m−1}}(K_{m−1})⊕AES4_{C_{i,m}}(K_m)
   if |K_m|<128 then X_i ← AES4_{C_{i,1}}(K_1)⊕AES4_{C_{i,2}}(K_2)⊕···⊕AES4_{C_{i,m−1}}(K_{m−1})⊕AES4_{C_{i,0}}(K_m 10*)
return (X_1, X_2, X_3)
```

Figure 4: **AEZ's hash, PRF, and TBC.** The last carries out key processing that an implementation would normally do at session-setup. An alternative "scaled-up" algorithm, AEZ10, would redefine E by setting $\mathrm{E}_K^{i,j}(X) = \mathrm{AES}_K(X \oplus iI \oplus jJ)$ where $I = \mathrm{AES}_K(\mathbf{0})$ and $J = \mathrm{AES}_K(\mathbf{1})$, now restricting keys to $\{0,1\}^{128}$.

depends on the length of the plaintext: as few as eight, or as many as 24. The round function is based on AES4. This is embodied in the pseudocode by the fact that our tweakable PRP decides to use AES4 or AES10 based on the first component of the tweak, employing the AES10 only for tweaks beginning with a −1. The Encipher-AEZ-tiny routine is illustrated at the bottom-right if Figure 5 for the setting where messages have 16 or more bytes.

A novel feature of AEZ-tiny is the possible xoring of a bit into the ciphertext just before the algorithm's conclusion. This is done to avoid simple random-permutation distinguishing attacks, for very short strings, based on the fact that Feistel networks only generate even permutations. A similar trick, conditionally swapping two fixed points, has been used before [31]. Our approach has the benefit that the natural implementation is constant-time.

The heart of AEZ is AEZ-core. It melds EME [15, 16], OTR [24], and a variety of other ideas. Consider the case where we want to encipher a string $M = M_1 M_1' \cdots M_m M_m' \, M_{\mathsf{x}} M_{\mathsf{y}}$ having an even number of blocks, all of them full. We call the first $2m$ blocks of M the i-blocks. Refer to the top-left and top-right of Figure 5. Regard each rectangle with a pair of numbers as a TBC, the label being

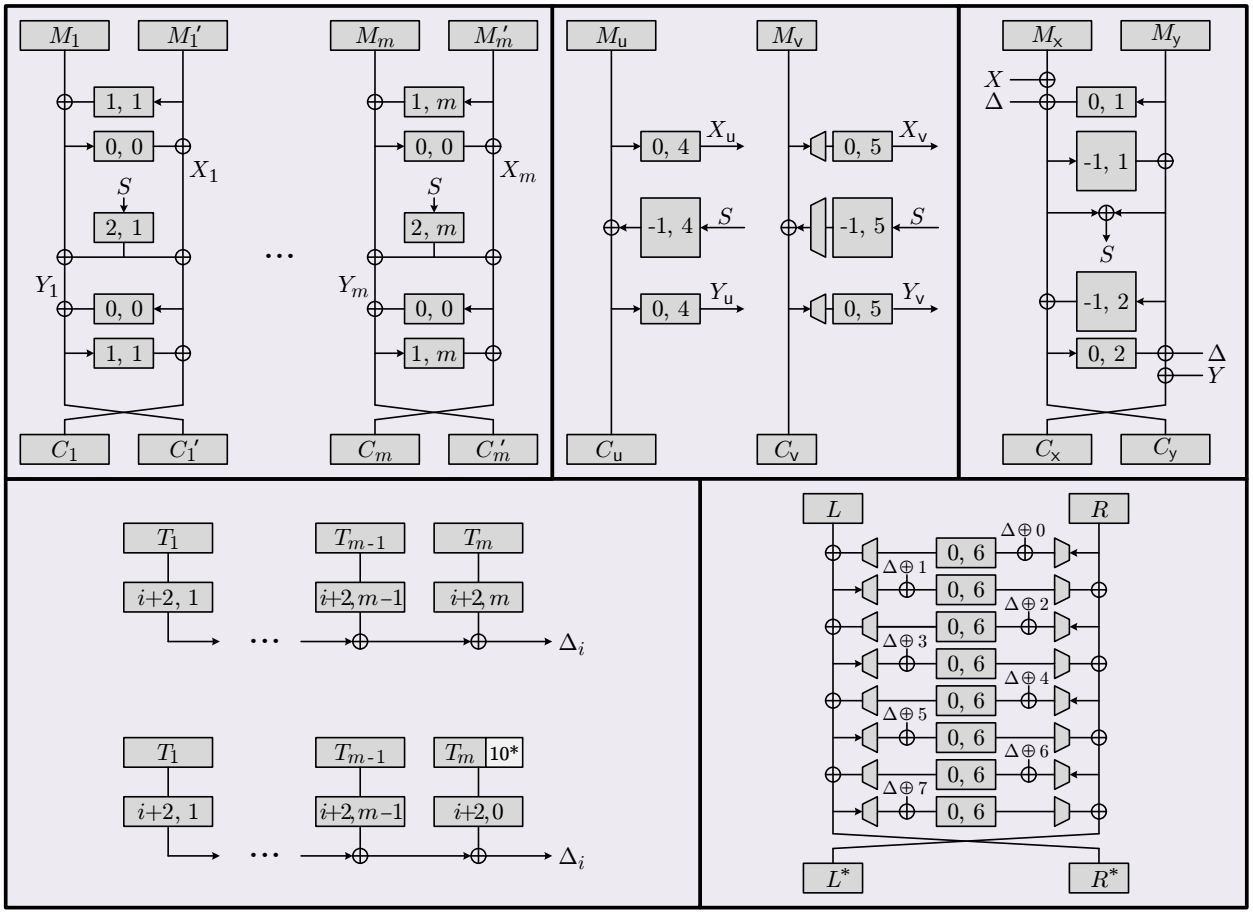

Figure 5: **Illustration of AEZ enciphering.** Rectangles with pairs of numbers are tweakable blockciphers, the pair being that tweak (the key, always K, is not shown). **Top row:** enciphering a message M of (32 or more bytes) with AEZ-core. The i-block (top left) is used for the bulk of the message, but the xy-block (top right) comprises the last 32 bytes, while the uv-block (top middle) comprises the prior 0–31 bytes. (The picture shows a uv-block of 17–31 bytes.) **Bottom left:** AEZ-hash computes $\Delta = \bigoplus \Delta_i$ from a vector-valued tweak encoding A, N, and ABYTES. Its i-th component $T_1 \cdots T_m$ is hashed as shown. **Bottom right:** AEZ-tiny, when operating on a string $M = L \parallel R$ of 16–31 bytes. More rounds are used if M has 1–15 bytes.

the tweak and the key K left implicit. Each pair of i-blocks $M_i M_i'$ is subjected to a two-round Feistel network. This both begins the scrambling of $M_i M_i'$ and yields a value $X = X_1 \oplus \cdots \oplus X_m$ that is a computational almost-xor-universal hash of $M_1 M_1' \cdots M_m M_m'$. The final pair of blocks $M_{\mathsf{x}} M_{\mathsf{y}}$ are now processed, but where X initially offsets one of them. This both begins the scrambling of $M_{\mathsf{x}} M_{\mathsf{y}}$ and yields the value S that is a computational almost-universal hash of all of M. The TBC calls of the middle row of the i-blocks now inject an (i, S)-dependent value. Two more Feistel rounds to each i-block gives $C_1 C_1' \cdots C_m C_m'$. To compute $C_{\mathsf{x}} C_{\mathsf{y}}$ we likewise employ two more Feistel rounds, the C_{x} value offset by a value $Y = Y_1 \oplus \cdots \oplus Y_m$ analogous to X.

The top-middle panel shows how to deal with messages having an even number of blocks, the last of these a fragment. Now the message is partitioned into $M_1 M_1' \cdots M_m M_m' M_{\mathsf{u}} M_{\mathsf{v}} M_{\mathsf{x}} M_{\mathsf{y}}$ with all blocks full except M_{v}, which will have 1–15 bytes. The Figure 5 trapezoids used to process M_{v} denote 10^*

padding (top and bottom) and truncation (middle). The value $X = X_1 \oplus \cdots \oplus X_m \oplus X_{\mathsf{v}} \oplus X_{\mathsf{u}}$ now includes contributions from X_{u} and X_{v}, and similarly for $Y = Y_1 \oplus \cdots \oplus Y_m \oplus Y_{\mathsf{u}} \oplus Y_{\mathsf{v}}$.

Messages with an odd number of blocks are handled similarly, with the v-column omitted and the padding and truncation, if needed, moved to the u-column. We say "if needed" because no padding or truncation is used if the u block is full (ie, an odd number of blocks, all of them full).

Let us call this construction just given AEZ-core[E]. It is the generalization of AEZ-core that employs an arbitrary tweakable blockcipher E. It should not be surprising that the construction is a strong-PRP under the assumption that the TBC used is secure as a tweakable-PRP. We prove this in the academic paper corresponding to this submission.

At this point we could instantiate E using a standard TBC based on AES: the XE method [22, 35] would do, yielding the scheme AEZ10 specified in the caption of Figure 4. We would then have a provably-secure enciphering scheme (for strings of 32 or more bytes) costing about five AES calls per pair of blocks, so 2.5 AES calls per block. The cost would be similar to EME [15, 16]: 0.5 more AES calls per block, but avoiding the repeated doubling and the use of AES-inverse.

But suppose we shatter our abstraction boundary and look at all that is going on to encipher M in AEZ10. Then the design starts to seem like major overkill: each block M_i is subjected to 30 rounds of AES (ten shared with a neighboring block), plus additional AES rounds to produce the unpredictable, M-dependent value S that gets injected into the process while 20 rounds yet remain.

In light of such overkill, AEZ-core selectively prunes some of the AES calls that AEZ10 would perform, using AES4 in their place. In particular, we prune invocations where we are trying to achieve computational xor-universal hashing. We leave enough AES rounds so that each block M_i is effectively processed with 12 AES rounds, eight of these subsequent to injection of the highly-unpredictable S and four of them shared with a neighboring block. The key steps in calculating S are not pruned, nor the TBC used to mask u- and v-blocks.

Tweak processing. So far we have not mentioned the processing of the tweak T built from N and A. This is shown in the bottom-left of Figure 5. First we compute a hash AEZ-hash on T to create a value Δ; and then Δ is injected into AEZ-tiny and AEZ-core processing as shown.

Deciphering. We define Decipher(K, T, Y) as the unique X such that Encipher$(K, T, X) = Y$. Logically, this is all we need say for the specification to be well-defined, so we omit writing out the implementing pseudocode. Still, that pseudocode is easy to describe. The reason this is so is that enciphering and deciphering are highly symmetric for both AEZ-tiny and AEZ-core.

AEZ-tiny deciphering is identical to AEZ-tiny enciphering except that we must count backwards instead of forwards, and we must do the only-even-cycles correction (line 215) at the beginning instead of the end. Specifically, a routine Decipher-AEZ-tiny(K, T, M) (the M now representing ciphertext) is identical to Encipher-AEZ-tiny(K, T, M) except that line 214 is changed to count from $k - 1$ down to 0, while for line 215 has each C replaced by M before moving the line up to just after line 212.

AEZ-core deciphering is identical to AEZ-core enciphering except that we must take the xy-tweaks in reverse order. Specifically, a routine Decipher-AEZ-core(K, T, M) (the M now representing ciphertext) is identical to Encipher-AEZ-core(K, T, M) except we swap tweaks $(0, 1)$ and $(0, 2)$, and we swap tweaks $(-1, 1)$ and $(-1, 2)$. These four tweaks appear at lines 227 and 232.

PRF. Using the Carter-Wegman approach, we also build a PRF AEZ-prf: counter-mode is employed to extend the output length if ABYTES > 16. The PRF is only used for the special case of enciphering the empty message. This is done for efficiency reasons—to make AEZ-prf$_K(X)$ = Encrypt$(K, \varepsilon, X, \tau, \varepsilon)$ an attractive PRF.

Key processing. For the users' convenience, AEZ allows keys of any length. Using procedure Extract, the provided key is processed into 48 bytes using an almost-universal hash function with a fixed but random-looking key. The approach is rooted in the leftover hash lemma [2, 10, 17]. The hash we use is simple and parallelizable, and in the spirit of AEZ-hash.

Tweakable Blockcipher. The TBC $\mathrm{E}_K^{i,j}(X)$ takes a tweak $(i, j) \in (\{-1, 0\} \times [0..7] \cup \{1, 2, 3\} \times \mathbb{N}$. The first component selects between AES10 (when $i = -1$) and AES4 (when $i \geq 0$). Either way, the construction is based on XE [22, 35]. A small amount of precomputation (to compute J, $2J$, $4J$, $8J$, and $16J$) will suffice: after each eight successive j-values we double L, add it in, and reset j.

1.5 Usage cap

We impose a limit that AEZ be used for at most 2^{48} bytes of data (about 280 TB); by that time, the user should rekey. For the purpose of this requirement, we say that, when encrypting (N, A, M) with a given key K, AEZ is acting on $\lceil |N|/8 \rceil + \lceil |A|/8 \rceil + \lceil |M|/8 \rceil$ bytes. The above requirement stems from the existence of birthday attacks on AEZ, as well as the use of AES4 to create a universal hash function.

2 Security Goals

Nonce-reuse security. AEZ achieves *nonce-reuse misuse-resistance* (MRAE), as previously defined by Rogaway and Shrimpton [37]. In an MRAE scheme, repeating a nonce will violate privacy only insofar as repetitions of (N, A, M) triples will be identified as such. It will not compromise authenticity at all. SIV [37] is the best-known MRAE scheme.

Some researchers call AE schemes nonce-reuse misuse-resistant more broadly, encompassing schemes that achieve much weaker notions, like those that leak the longest common block-aligned prefix (for some fixed and typically small blocksize). Such notions were invented to approximate best-possible security for online schemes, which they do rather inexactly. MRAE schemes can't be online.

Exploitation of embedded novelty. MRAE security implies automatic exploitation of randomness or sequence numbers present in messages: in any context where messages are known to be distinct (eg, a sequence number is embedded somewhere within) or are extremely unlikely to collide (eg, a freshly-generated session key is embedded somewhere within), use of a nonce unnecessary. In such settings, omission of a nonce does *not* represent misuse; it is a sound way to encrypt.

Exploitation of domain-specific redundancy. In many contexts, plaintexts have a certain expected structure. This might arise because the message was produced by or for a particular protocol. We intend that if the user checks for the anticipated structure and regards messages as inauthentic if they don't comply, then this check augments authenticity and correspondingly lessens

the need for the nominal redundancy that is inserted by AEZ before enciphering (that is, the extra ABYTES zero bytes). The concept of automatically exploiting redundancy present in plaintexts to achieve authenticity is well known in cryptographic folklore, where it has often been wrongly assumed, and demonstrably achieved for AE based on a strong-PRP [4].

Releasing unverified plaintext. When decrypting, an *unverified plaintext* is a string that will be released if the ciphertext is deemed authentic, but is supposed to be quashed otherwise. While not definitionally mandated, AE schemes routinely compute such a thing. One form of encryption-scheme misuse is to release some or all of the unverified plaintext despite the ciphertext's invalidity. This might happen because of an incremental decryption API or a more traditional side-channel.

Contemporaneous work by Andreeva *et. al* gives definitions to formalize an AE scheme's security against release of unverified plaintexts [1]. Our own definitional approach is different; we formalize *robust* AE, which incorporates the unverified-plaintext concern among its aspects. In claiming robust-AE security for AEZ the unverified plaintext is the value X computed at line 114. Achieving robust AE implies that no harm would come of returning $(X, \perp)$ instead of $\perp$ at line 116.

Per-message nonce-length and parameter authentication. No security problems result from employing nonces of varying lengths during a session, nor from changing the authenticator length ABYTES during a session. Of course accessing such capabilities would require an appropriate API.

Good security for low ciphertext expansion. Traditionally, AE security definitions "give up" when the adversary forges. This means that, at least definitionally, it's OK for a scheme to fail catastrophically when it first fails. A consequence is that authentication tags need to be so long that forgeries almost never occur. Yet there are applications where an occasional forgery is fine. For example, in some settings it ought to be fine to use a one-byte authenticator: while the adversary will have a 2^{-8} chance of forging a given message, we could still expect that, say, a reasonable adversary won't have much more than a 2^{-80} chance to forge ten consecutive messages.

AEZ permits short authentication tags, getting security as strong as possible given the selected authenticator length. This implies that we must use a new definition for AE, one that does *not* give up when a forgery occurs. It is described next.

Robust AE. Our new security definition for AE formalizes that one is doing *as good a job as possible for a given value τ of ciphertext expansion* ($\tau = 8 \cdot$ ABYTES). The statement is required to hold even in the face of decryption leaking some specified information. An upcoming academic paper defines and investigates this notion of *robust AE*. Here we sketch the idea.

We restrict attention to AE schemes $\Pi = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ that operate on strings of any length and that are *τ-expanding*, $|\mathcal{E}_K^{N,A}(M)| = |M| + \tau$, for a user-selectable $\tau \in [0..\tau_{\max}]$. We first consider an adversary that has access to one of two pairs of oracles. In the *real* setting the *encryption oracle* encrypts according to $\mathcal{E}$ and the *decryption oracle* decrypts according to $\mathcal{D}$. In the *ideal* setting the encryption oracle, asked (N, A, τ, M), returns $\pi_{N,A,\tau}(M)$ where, for each N, A, τ, the function $\pi_{N,A,\tau}$ is a uniformly selected random injection from m-bit strings to $(m+\tau)$-bit ones. All of these functions are chosen independently. The decryption oracle, given (N, A, τ, C), checks if there's an M such that $\pi_{N,A,\tau}(M) = C$. If so, it returns M. Otherwise it returns the distinguished value $\perp$.

The above notion coincides with that of a *pseudorandom injection* (PRI) that has been updated

to regard τ as in input. To arrive at the more general notion of an RAE scheme, we modify how decryption works in the ideal setting. This is unchanged when (N, A, τ, C) is valid (that is, when there is an M such that $\pi_{N,A,\tau}(M) = C$), but when it's not, a simulator S gets to return what it wants. The return value may be based only on N, A, C, τ and any saved state of S. The real decryption algorithm $\mathcal{D}$ can now be augmented to capture any desired leakage when the ciphertext is invalid: have algorithm $\mathcal{D}$ return what it wants, as long as it is recognizably invalid (eg, we can require that the length of the unverified plaintext not be $|C| - \tau$ bits). The notion is stronger than before insofar as not only must the scheme approximate a PRI with respect to valid ciphertexts, but, when they're invalid, the simulator must still be able to approximate that which $\mathcal{D}$ returns.

While the simulator S and invalid-message-returning $\mathcal{D}$ strengthen the RAE notion relative to the PRI notion, the key aspect, we think, is simply our insistence that encryption looks like a PRI even in the case that the ciphertext expansion is zero or small. In fact, when the ciphertext expansion is large, the PRI notion and the MRAE notion effectively coincide [37]. On the other hand, when ciphertext expansion is zero, the RAE (and PRI) notion coincides with that of a strong-PRP. RAE security can be thought of as a way to bridge strong-PRP security and MRAE security, coinciding with the former when τ is zero and the latter when τ is large.

Provable security. AEZ has been developed using the tools of provable security. The paradigm used is what we call *accelerated* provable-security. First, a scheme is designed and proven secure when its underlying cryptographic tool—a tweakable blockcipher (TBC), in the case of AEZ—meets some well-established security definition. At that point one could instantiate the primitive with a conventional tool—eg, using AES and the XE construction [22, 35], as we described for AEZ10. One would then have a scheme with a customary provable-security claim. Instead, to make our scheme faster, we choose to selectively instantiate *some* of the TBC calls with a construction based on AES4, a four-round version of AES. Insofar as AES4 is *not* secure as a PRP (and, additionally, our method of tweaking it is not always XE), this step is effectively heuristic.

We call the instantiation of a scheme using a mixture of full and downgraded primitives the *scaled-down* design. In contrast, using a conventional construction for the primitive would yield the usual, *scaled-up* design. AEZ is a scaled-down realization of $\widetilde{\text{AEZ}}$. It is a *thesis* underlying our design methodology that the approach is useful both to discover good schemes and to have some measure of assurance for them.

Quantitative security statements. For the scaled-up version of AEZ with default parameters, we expect that an adversary cannot be exhibited that violates RAE security with advantage exceeding $4s^2/2^{128} + t/2^{128}$ where s is the total number of 16-byte blocks of messages encrypted or authenticated (plus 3 blocks per message, by convention) and t is the time (including the description size) in which the adversary runs. The second addend is a stand-in for an advantage term associated to breaking the PRP security for the underlying blockcipher. Constants 3 and 4 are the result of ongoing analysis. The number of encryption and decryption queries does not appear in the formula above because we have folded them into s.

For `aez` itself, the formula should be replaced by $4s^2/2^{113} + t/2^{128}$ because of the higher maximal expected differential probability of AES4 [19] compared to an ideal hash or cipher.

Many authors prefer to think of security in terms of number-of-bits. We would summarize the $4s^2/2^{113} < s^2/2^{110}$ term of the last formula by saying that `aez` is expected to have 55 bits of

Security goal	**Query complexity**	**Time complexity**	**Approx formula**
Confidentiality of plaintext	55	128	$s^2/2^{110} + t/2^{128}$
Authenticity of plaintext	55	128	$s^2/2^{110} + t/2^{128}$
Authenticity of AD	55	128	$s^2/2^{110} + t/2^{128}$
Authenticity of the nonce	55	128	$s^2/2^{110} + t/2^{128}$
Robust AE	55	128	$s^2/2^{110} + t/2^{128}$

Figure 6: **Security goals for AEZ with default parameters (`aez`).** Query complexity is log base-2 of blocks queried: one needs about 2^{55} blocks before having a good chance to violate the goal. Time complexity is log base-2 of cycles: one needs about 2^{128} time to break the goal if one has only small amount of plaintext/ciphertext. The formula bounds adversarial advantage as a function of queried blocks (s) and time (t) by a known, modest-size adversary. The final row, RAE security, not only implies the other rows but also nonce-reuse misuse-resistance: AEZ provides maximum-possible robustness against nonce reuse.

security. We warn that when an author makes a claim like "GCM has 128 bits of security" the focus is *time* complexity, imagining a fixed and small amount of ciphertext. When saying that we have at least 55 bits of security we are speaking exclusively of query complexity: that an adversary must gather roughly 2^{55} blocks (2^{59} bytes) worth of ciphertext before it has a good chance to break RAE security (assuming an explicitly given attack of reasonable description size and time complexity). Recall our usage cap, that AEZ should be used for at most 2^{48} bytes. One might summarize targeted security goals for `aez` as shown in Figure 6.

Non-goals. We have not tried to achieve security beyond the birthday bound; like traditional modes of operation based on a 128-bit blockcipher, there certainly *are* easy distinguishing and forging attacks by the time the adversary queries AEZ with about 2^{64} blocks of message, AD, or nonce. Similarly, we are not targeting time-complexity security in excess of what is inherent in employing a 128-bit key. (That said, we avoid the obvious 2^{128}-time brute-force attack for keys in excess of 128 bits by processing arbitrary-length keys to 48-byte subkey material in our key extraction processing.)

Warning. Robust AE should not be understood as blanket permission to omit a nonce or allocate too few bits for ciphertext expansion. Let us elaborate.

In the context of AE, misuse resistance (MR) has, of late, been brandished far too liberally, with some authors going so far as to call their *online* AE schemes *misuse resistant*, or even *nonce free*. We disapprove. Online AE schemes can never be misuse-resistant in the sense originally defined [37], and, what is worse, it is unclear that they imply *any* useful guarantee when nonces repeat. The definitions here are deceptive [3, 14], sounding stronger than they are [33].

We worry that the expansion of the term "misuse resistance" to online schemes may wrongly signal that there are online AE schemes where nonces are effectively optional. We wish to emphasize that, even with RAE, nonces *still* should not be construed as optional in the absence of supporting analysis. Specifically, a nonce must be used unless one has certitude that, even in the presence of the adversary, all encrypted (A_i, M_i) pairs will be distinct; or else one has some other domain-specific reason to believe that, for the given context, leaking plaintext-equality is not problematic.

In a similar vein, AEZ allows little or no ciphertext expansion. But the adversary's per-message forging probability increases with decreasing redundancy. Applications should not reduce ABYTES to zero or some other small value without ensuring that, combining the ABYTES zero bytes with any decryption-verified redundancy, there remain enough total bits r of redundancy that forging each message with probability 2^{-r} is alright.

3 Security Analysis

An academic paper with the relevant security proofs for AEZ is in preparation. In the meantime, we summarize some of our results. All are in the provable-security tradition (as opposed to our making cryptanalytic claims).

Ciphertexts of at least one block. Let $\widetilde{\text{AEZ}}[\text{E}]$ be the generalization of AEZ where each E is a tweakable blockcipher (TBC) of the correct signature [22]. We can prove that $\widetilde{\text{AEZ}}[\text{E}]$ achieves RAE security as long as E is secure as a tweakable PRP. The claim assumes that $|M| + \tau \geq 128$ for each encryption query employing plaintext M, and $|C| \geq 128$ for each decryption query of a ciphertext C. These conditions hold automatically for the default choice of ABYTES = 16. With those provisos, RAE security can be proven along the following lines.

- AEZ-core provides a length-preserving, variable-input-length, strong PRP on $\text{Byte}^{\geq 32}$ (strings of 32 or more bytes) with birthday-bound distinguishing advantage. This statement requires only chosen-plaintext-attack PRP security for the underlying TBC.
- The tweak provided to Encipher-AEZ-core is incorporated by what can be regarded as the XEX construction [22, 35]. The underlying hash function, AEZ-hash, is almost-xor universal (AXU) when E is a PRP.
- The round functions of AEZ-tiny are derived from a tweakable blockcipher (TBC) with tweak space $\mathcal{T} = \{(i, 0) \mid i = 1, \ldots, 24\}$. We employ the XE construction [22, 35] to extend the tweak space to $\mathcal{T} \times \mathcal{N} \times \mathcal{A}$. One can then view that, for each (N, A), we use independent round functions. Since a 6-round Feistel network on $\{0,1\}^{2n}$ already yields a strong PRP with birthday-bound distinguishing advantage [21, 29, 30], AEZ-tiny gives a length-preserving, strong tweakable-PRP on $\text{Byte}^{\geq 16} \cap \text{Byte}^{\leq 31}$, with birthday-bound distinguishing advantage.
- Once one has shown that the Encipher procedure of $\widetilde{\text{AEZ}}$ provides a length-preserving strong tweakable-PRP then $\widetilde{\text{AEZ}}$ itself is a robust-AE scheme. This follows from a generic result that asserts that encode-then-encipher conversion gives RAE security.

The choice of our TBC is heuristically justified as follows.

- The processing of the tweaks to compute the XE offsets only requires a universal hash, and four-round AES with independent, uniformly random subkeys is already known to be a good AXU hash [19]. Similarly, the AXU security for AEZ-hash can be justified by viewing AEZ-hash as an approximation of a variant in which the subkeys are chosen uniformly and independently from $\{0,1\}^{128 \cdot 4} \times \{0^{128}\}$. That variant of AEZ-hash is again AXU due to the fact that four-round AES with independent, uniformly random subkeys is an AXU hash [19].
- For AEZ-core, when processing each pair of blocks M_{x} and M_{y}, the first and last rounds only

need to be AXU, due to the classic result of Naor and Reingold [26]. Then, for the four-round Feistel networks that process M_i and M'_i with $i \geq 1$, we heuristically use AES4 for the round function, since, even then, each ciphertext block C_i is processed with 12 AES rounds (four of which are shared with a single neighboring block), eight of which are subsequent to full mixing, and all of which are subsequent to the position-dependent masking.

- For AEZ-tiny we are effectively using a minimum of $32 = 8 \cdot 4$ rounds of AES. While AES4 is not itself a good PRF, it would seem to be a stronger round function than those used by most conventional Feistel-based designs.

Let ϵ be the maximum expected differential probability of (independently-keyed) AES4; this is known to be at most $(52/2^{34})^4 \approx 2^{-113.088}$ [19]. While $\widetilde{\text{AEZ}}$ achieves RAE security with birthday-bound security in the blocksize, AEZ only achieves RAE security with advantage about $\sigma^2 \cdot \epsilon$, where σ is the number of blocks that the adversary queries. There are corresponding attacks. As a simple example, let ABYTES = 16 and have an adversary repeatedly ask to encrypt a fixed message M with a fixed nonce N but using AD values that consist of two random blocks. A collision in ciphertexts will be found in about $1/\sqrt{\epsilon}$ expected queries. Say it arose from AD values of $A = A_0 A_1$ and $A' = A'_0 A'_1$. Then test if one again gets a collision with M and N but with AD values of either $A \parallel \mathbf{0}$ or $A' \parallel \mathbf{0}$. If so, one almost certainly has a "real" encryption oracle.

Security of AEZ-prf. If AEZ-hash is an AXU hash and AES10 is a good PRF then AEZ-prf is a PRF, as AEZ-prf is constructed from the Carter-Wegman paradigm [7], with output length expanded via counter mode of AES10 for ABYTES > 16. Alternatively, for the default ABYTES = 16, one can view AEZ-prf as an approximation of an AES-based PMAC [6] in which all but the final blockcipher call have had the number of AES rounds reduced from 10 to 4, a heuristic employed in ALRED, MARVIN, and PELICAN [8, 9, 38, 39]. This gives another heuristic justification for the scaling down from AES10 to AES4 in AEZ-hash.

Ciphertexts of less than one block. The claim that Encipher-AEZ-tiny gives a tweakable, strong PRP over $\text{BYTE}^{\leq 15}$ is heuristically justified. Consider a collection of independent, ideal, k-round Feistel networks on $\{0,1\}^{2n}$; the round functions are all uniformly random and independent. The best attack known that distinguishes them from a family of independent, truly random even permutations, requires at least $2^{(k-4)n}$ plaintext/ciphertext pairs [28]. From our choice of the number of rounds, this attack needs at least 2^{72} plaintext/ciphertext pairs, and thus doesn't violate our security goals.

There are of course many provable-security results on balanced Feistel as well, but proven bounds for a fixed-round Feistel network operating on an m-bit string vanish at about $2^{m/2}$ queries, and we are looking at settings with m as small as 8.

Key scheduling and AES4/AES10 details. For the analysis above we sometimes pretended that the subkeys for AES4 (excluding the XE offsets) are independent of other keys. In the implementation, to reduce context size, all subkeys are made from three blocks: I, J, and L. Associated to this choice, we elect to determine these from the underlying key K in a more conservative (and also more parallelizable) manner than with the AES key-scheduling algorithm, which we do not employ at all.

operation	$m \geq 2$ even $d=128$	$m \geq 2$ even $d<128$	$m \geq 3$ odd $d=128$	$m \geq 3$ odd $d<128$	$m=1$ $d=8$	$m=1$ $d=16$	$m=1$ $d \geq 24$	$m=2$ $d<128$
encipher or decipher [a]	m+0.8 (3.6)	m+2.4 (3.6)	m+1.6 (3.6)	m+1.6 (3.6)	10 (10)	6.8 (6.8)	4.4 (4.4)	3.2 (3.2)
encrypt or decrypt [b]	m+3 (3.6)	m+3 (3.6)	m+2.2 (3.6)	m+3.8 (3.6)	3.6 (3.6)	3.6 (3.6)	3.6 (3.6)	5 (4)
reject invalid ciphertext [b]	0.4m+2.4 (2.8)	0.4m+2 (2.8)	0.4m+2.4 (2.8)	0.4m+2 (2.8)	0 (0)	0 (0)	0 (0)	3.6 (3.6)

Figure 7: **Efficiency of AEZ.** Worst-case computational work (and, parenthesized, latency) measured in AES-equivalents, defined as ten AES rounds. The nonempty string X being operated on has $m = \lceil |X|/128 \rceil$ blocks, the possibly-fragmentary last one having $1 \leq d \leq 128$ bits. Assumptions: (a) Key already setup, nonce and AD already processed. (b) Key already setup, AD already processed, nonce has 16 or fewer bytes, ABYTES = 16. Other tasks: **Key setup**: 1.2m (0.4), assuming all needed constants have been precomputed. **Process string-valued AD**: 0.4m (0.4) (key already setup, nonce of 16 or fewer bytes).

In defining AES4 and AES10 subkeys, the initial and final subkeys are sometimes taken to be zero. In most cases arguments can be given that the simplification is without adverse consequence. Cascading a post-whitened permutation with a pre-whitened one seems redundant. If prewhitening is used from the XE construction then little benefit is gained from an initial round key. If one is aiming to construct an AXU hash function, postwhitening is pointless. Pleasantly, using zero as a final AES4 round key frees up the xor included in the `aesenc` instruction to do the other computational work needed for Feistel.

AES4 and AES10 do not omit the final-round `MixColumns`, as AES itself does. In the context of repeated AES4 applications, omission of the final `MixColumns` would likely decrease security. See Dunkelman and Keller for some work in this direction [11]. And the AES designers' motivation for removing the `MixColumns` step from the last round of AES is for us moot: the inverse AES cipher is never used.

We emphasize that the E construction is *not* secure as a tweakable-PRP; four AES rounds of AES is not sufficient for that purpose. This where the scaling-down has occurred. One only gets a tweakable PRP by moving to the construction described for AEZ10.

At present we view the entropy extraction procedure Extract as essentially heuristic, although some weak provable-security claims about it can be made from the leftover hash lemma [2, 10, 17].

4 Features

See Figure 7 for a table summarizing computational costs and Figure 8 for a table summarizing algorithmic features. Below we enumerate additional features and restate some key ones.

1) Strings of any byte length m can be encrypted into strings of $m + \text{ABYTES}$ bytes for any (user-selectable) value ABYTES. One achieves the maximal privacy and authenticity consistent with ABYTES (assuming this value is not excessively large, whence its increase adds nothing). The value ABYTES is authenticated and may change as often as a user likes.
2) Computational cost is close to that of AES-CTR mode: roughly 1 AES-equivalent per block.

Objective	**Robust-AE**, a goal that implies **MRAE** (nonce-reuse misuse resistance) [37].
Type	**Blockcipher-based** scheme, based on **AES4** and **AES10**.
Intended for	**sw/hw/lw**. Intended to do well where AES does, in SW or HW, and on low-power devices where ciphertext length should be minimized.
Key length	**Arbitrary**. Subkeys are extracted from a parallelizable universal hash.
Nonce length	**Arbitrary**. May vary during a session.
Auth length	**Arbitrary**. Expansion beyond 16 bytes does not enhance security. Expansion by 0 bytes gives a strong, tweakable, VIL blockcipher.
Nonce reuse	**Yes**. Secure against nonce-reuse in the strongest sense of the phrase [37].
Online	**No**. MRAE schemes can't be online (neither encryption nor decryption).
Unverified plaintext	**Yes**. It is fine to release unverified plaintext (a recovered but inauthentic plaintext). This is one aspect of our notion of a robust AE (when $\mathcal{D}$ is appropriately defined).
Parallelizable	**Yes**. Two passes must be made to encrypt or decrypt, but both are parallelizable. Processing of the AD is also parallelizable.
Incremental	**No**. MRAE schemes can't be incremental. Use as a deterministic MAC is incremental with respect to block replacement or appending-on-the-right.
Inverse free	**Yes**. The inverse direction of AES4 or AES10 is never used.
Context size	**144 bytes** for $I, J, 2J, 4J, 8J, 16J, L, \Delta_1, \Delta_3$ or **128 bytes** for $I, J, 2J, 4J, 8J, 16J, L, \Delta_3$ (0.4 extra AES per AD) or **64 bytes** for I, J, L, Δ_3 (2 extra doublings per msg, 4 per AD) or **48 bytes** for I, J, L (can't preprocess the AD).
Static AD	**Yes.** Static AD values can be preprocessed and used thereafter at near-zero cost.
Fast reject	**Yes.** Invalid ciphertexts can be rejected far more quickly than valid ones decrypted.
Performance	About the cost of OCB or AES-CTR, approaching 1.0 AES-equivalents per block
Proofs	Either: **Yes**, there are proofs, but then a heuristic optimization is applied to a provably-secure scheme to get a nice speedup; or **No**, there are no proofs for AEZ itself, although the authors employ provable-security to motivate and justify some design choices.
Further features	▸ Can exploit arbitrary redundancy in messages for authenticity ▸ Can be used as an efficient, parallelizable MAC (encrypt the empty string). ▸ Can be used to encipher short strings and to encrypt strings with low expansion. ▸ ABYTES is authenticated may vary during a session. ▸ Extensions (not AEZ itself) will support secret nonces, plaintext-length obfuscation, and radix64url output encoding. ▸ No patents.

Figure 8: **Table of properties for AEZ**. The choice of properties to list as rows evolved from slides prepared by Bart Preneel during a Dagstuhl workshop [32].

And an implementation only needs to employ the forward direction of AES.

3) Nonces can have any length. Their lengths can vary during a session.
4) The AD can be an arbitrary list of arbitrary strings. This obviates the need for users to encode multiple strings into a string-valued AD.
5) It is fine to omit nonces (select $N = \varepsilon$) if one is sure that (A, M) pairs will not repeat. If they do repeat, damage is limited to divulging equality among (A_i, M_i), (A_j, M_j) pairs.
6) Can be used as an arbitrary-output-length pseudorandom generator (PRG) or pseudorandom function (PRF): define $f_K(X, \ell)$ as the result encrypting $M = \varepsilon$ with $A = X$ and ABYTES $= \ell$. Here ℓ is number of bytes requested.

7) Keys can have any length. A user may, for example, use a passphrase or DH ephemeral key. (Note: some features one might want for mapping a passphrase to a 128-bit key, like salting and an intentionally slow mapping to slow password guessing, are not natively provided.)
8) AEZ functions well as a stand-alone MAC and as a stand-alone enciphering scheme. The first use needs only 0.4 AES operations per block.
9) Verification of plaintext redundancy enhances authenticity, as we have already explained.
10) Short authenticators provide the security one would hope for. Our security notion doesn't "give up" when the adversary forges. This is part of the robust-AE notion.
11) Release of unverified plaintext is fine. This is another part of the robust-AE notion.
12) The security properties achieved by AEZ enable support for secret message numbers as a simple add-on. This will be accomplished as an AEZ extension. Further AEZ extensions will handle plaintext-length obfuscation, password salting, password guess-throttling, and encoding ciphertexts into a target alphabet.
13) An encryption implementation can make one left-to-right pass over the plaintext, writing an intermediate string as long as the plaintext; and then a second left-to-right, constant-memory pass over the intermediate string, this time outputting the ciphertext online. Decryption can be similarly realized. If one does not want to write out an intermediate string, which must not be released to the adversary, the cost increases from 1 AES-equivalent per block to 1.4 AES-equivalents per block.
14) Some AE schemes need that the AD be available before the message is processed. AEZ only requires AD processing be completed by the end of pass-1.
15) It is possible to accelerate the rejection of invalid ciphertexts by having decryption compute the final ciphertext block prior to computing the remainder of the plaintext. The cost for early-rejection is about 0.4 AES-equivalents per block.
16) AEZ is fully parallelizable in the processing of plaintext, ciphertext, and AD.
17) Static AD can be preprocessed so that one doesn't have to subsequently pay a per-message $|A|$-dependent cost. Of course realizing this benefit requires an API that decouples provisioning of the AD and provisioning of other inputs.)
18) Word alignment of the message and AD are not disrupted (for example, one never prepends a byte to the message or AD, and then processes it.
19) The context size has been kept quite small: retaining "everything" one might want gives a context size is 144 bytes, and an implementation can make due with as little as 48 bytes without incurring an excessive computational price.
20) No AEZ-related patents have been or will be requested.

On an Intel Haswell CPU, our current implementation of AEZ, written in C with SSE intrinsics, encrypts or decrypts 16 KB strings at 0.72 cpb. Strings of 1500 bytes take 0.89 cpb. On decryption, invalid messages are more quickly rejected; for example, an invalid 1500 byte string is rejected at a cost of 0.38 cpb. Associated data processing speed, which also determines MAC speed, is also 0.38 cpb for 1500 bytes. All of these figures assume 12-byte nonces and 16-byte tags.

Advantages over GCM. AEZ has much stronger security properties than GCM. The later is not nonce-reuse secure, cannot safely generate short tags [13], and is not secure with respect to

disclosure of unverified plaintext. GCM does not achieve the RAE security definition. AEZ avoids GF(2^{128}) multiplies (apart from the finite-field "doubling" that it uses).

A closer match to AEZ in terms of high-level aims is SIV, which is at least nonce-reuse secure [37]. But SIV has to output 128-bits more than its input; it is not RAE secure; and it is not parallelizable (although the last issue could easily be fixed).

5 Design Rationale

Enciphering-based AE. An old result had already shown that enciphering with a strong PRP provides a versatile route to AE [4]. We recently came to understand just how attractive this route might be. On the one hand, we kept hearing requests for stronger AE security properties, like nonce-reuse misuse-resistance, authenticity without minimal ciphertext expansions, and security if unverified plaintexts are disclosed. Enciphering-based AE could deliver such aims. On the other hand, enciphering schemes that worked on either long or short strings were steadily becoming better-known objects. While they didn't have the efficiency of OCB, say, neither were they computationally exorbitant. And there was the hope of doing better.

Developing the enciphering scheme. With AES support increasingly embedded into devices, we wanted to base our enciphering scheme on the AES round function. A wide body of work had made abundantly clear that the best techniques for AES-based enciphering were going to depend on the length of the plaintext. When the plaintext was short, we would want a simple, `aesenc`-based design. For long strings we would want a more conventional mode. To cover all strings we'd have to glue the two together.

For enciphering short strings, some version of FFX [5] was the obvious choice. It was already in a draft standard [12], and the long history of Feistel networks made the choice seem safe (even if security bounds for balanced Feistel networks become disappointing when the input gets too short).

For enciphering longer strings, there were a great many off-the-shelf alternatives we could turn to (see [36] for a list). The best-known was EME2 [15, 18]. But its treatment of final fragments and long messages seemed complex, and it needed two AES calls per block and lots of doubling. Most alternatives traded a blockcipher calls for a potentially expensive finite-field operation, a direction we didn't want to go. We decided that no off-the-shelf solution would do.

AEZ-core builds on EME [15, 16] and OTR [24], but uses tweakable blockciphers [22] to arrive at an analyzable design. It makes strong use of what we have called accelerated provable-security. The scaled-down design, with a per-block cost of just 1.0 times that of AES and no use of inverse-AES, was cheaper than we initially imagined to be possible. While it has long been understood that stream ciphers could be faster than blockciphers, it was not anticipated, at least by us, that a wide-blocksize blockcipher could be about as cheap as a conventional blockcipher.

No hidden weaknesses. The designers have not hidden any weaknesses in this cipher. The authors do not know any technical means by which one *could* intentionally weaken the design of a scheme like AEZ. The authors excoriate intelligence-agency efforts to subvert security standards and mass-market implementations.

6 Intellectual Property

The submitters have not applied for any patents in connection with this submission and have no intention to do so. As far as the inventors know, AEZ may be used in an application or context without IP-related restrictions. If any of this information changes, the submitters will promptly (and within at most one month) announce these changes on the crypto-competitions mailing list.

7 Consent

The 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 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 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 submitters acknowledge that the committee decisions reflect the collective expert judgments of the committee members and are not subject to appeal. The 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 submitters understand that this statement is required as a condition of consideration of this submission by the CAESAR selection committee.

8 Revision History

Below we record each public version of AEZ since its inception.

- **AEZ v1** (2014.03.15): Initial definition. Submitted to the CAESAR competition.
- **AEZ v1.1** (2014.04.29): A minor revision to correct various v1 typos.
- **AEZ v2** (2014.08.17): A major revision, the enciphering algorithm MEM was replaced by EME4. While no problems were ever found with MEM, the move facilitates two major gains: (a) the cost is reduced from from 1.8 times that of AES to 1.0 times that of AES, while (b) all use of the AES-inverse operation is removed from AEZ. Also, EME4 is simpler, and the entire spec was correspondingly simplified.
- **AEZ v3** (2014.09.22): To simplify implementations the $(M_{\mathsf{x}}, M_{\mathsf{y}})$ pair of blocks is now taken from the end of the string being enciphered/deciphered instead of the beginning. Round keys were simplified. To minimize latency and facilitate fast rejection of invalid ciphertexts, both X and Δ are added to M_{x}. Support is added for vector valued-AD, which entailed enriching AEZ-hash and removing Format(), EXTNS, and the upper limit on ABYTES. Functions (FF0, EME4, AHash, AMac) were renamed to (AEZ-tiny, AEZ-core, AEZ-hash, AEZ-prf).

Acknowledgments

Phil thanks Dustin Boswell and René Struik for their unwitting role motivating the creation of AEZ. In an April 2013 email to Mihir Bellare, Boswell wrote of wanting an easier-to-use encryption scheme, and he sketched how such a scheme might look to its user. Bellare kindly passed the note on. Struik later gave a presentation that emphasized the importance of minimizing length expansion in low-energy environments [40]. Rogaway also thanks Stefan Lucks for a Jan 2012 discussion in which Lucks advocated developing a solution to the problem of leaking unverified plaintexts.

Thanks to Terence Spies for feedback during the AEZ design. Work with him on FPE and FFX helped motivate linking up enciphering and AE. Thanks to Danilo Gligoroski, Liden Mu, Chris Patton, Tom Ristenpart, and Yusi (James) Zhang for comments and corrections. Liden and Chris did AEZ implementations that helped validate our own.

Hoang was supported by NSF grants CNS-0904380, CCF-0915675, CNS-1116800 and CNS-1228890; Krovetz was supported by NSF grant CNS-1314592; and Rogaway was supported by NSF grants CNS-1228828 and CNS-1314885. Many thanks to the NSF for their continuing support. Part of this work was done when Hoang was working at UC San Diego.

No animals were harmed in conducting our research.

References

[1] E. Andreeva, A. Bogdanov, A. Luykx, B. Mennink, N. Mouha, and K. Yasuda. How to securely release unverified plaintext in authenticated encryption. Cryptology ePrint Archive, Report 2014/144. Feb 25, 2014.

[2] B. Barak, Y. Dodis, H. Krawczyk, O. Pereira, K. Pietrzak, F. Standaert, and Y. Yu. Leftover hash lemma, revisited. *CRYPTO 2011*, LNCS 6841, pp. 1–20, 2011.

[3] M. Bellare, A. Boldyreva, L. Knudsen, and C. Namprempre. On-line ciphers and the hash-CBC constructions. *CRYPTO 2001*, LNCS vol. 2139, Springer, pp. 292–309, 2001. Full version as Cryptology ePrint report 2007/197, June 29, 2007.

[4] M. Bellare and P. Rogaway. Encode-then-encipher encryption: How to exploit nonces or redundancy in plaintexts for efficient cryptography. *ASIACRYPT 2000*, LNCS 1976, Springer, pp. 317–330, 2000.

[5] M. Bellare, P. Rogaway, and T. Spies. The FFX mode of operation for format-preserving encryption. Draft 1.1. Submission to NIST, available from their website. Feb 20, 2010.

[6] J. Black and P. Rogaway. A block-cipher mode of operation for parallelizable message authentication. *EUROCRYPT 2002*, LNCS 2332, Springer, pp. 384–397, 2002.

[7] L. Carter and M. Wegman. Universal classes of hash functions. *Journal of computer and system sciences*, 18(2), pp. 143–154, 1979.

[8] J. Daemen and V. Rijmen. The Pelican MAC function. Cryptology ePrint Archive: Report 2005/088. 2005.

[9] J. Daemen and V. Rijmen. A new MAC construction ALRED and a specific instance ALPHA-MAC. *Fast Software Encryption*. LNCS 3557, pp. 1–17, 2005.

[10] Y. Dodis, R. Gennaro, J. Håstad, H. Krawczyk, and T. Rabin. Randomness extraction and key derivation using the CBC, cascade and HMAC Modes. *CRYPTO 2004*, LNCS 3152, Springer, pp. 494–510, 2004.

[11] O. Dunkelman and N. Keller. The effects of the omission of last round's MixColumns on AES. *Information Processing Letters*, 110, pp. 304–308, 2010.

[12] M. Dworkin. Recommendation for block cipher modes of operation: methods for format-preserving encryption. NIST Special Publication 800-38G: Draft. Jul 2013.

[13] N. Ferguson. Authentication weaknesses in GCM. Manuscript. May 2005.

[14] E. Fleischmann, C. Forler, S. Lucks, and J. Wenzel. McOE: A foolproof on-line authenticated encryption scheme. IACR Cryptology ePrint Archive 2011, 644 (2011)

[15] S. Halevi. EME*: Extending EME to handle arbitrary-length messages with associated data. INDOCRYPT 2004. pp. 315–327, 2004.

[16] S. Halevi and P. Rogaway. A parallelizable enciphering mode. *CT-RSA 2004*, LNCS 2964, Springer, pp. 292–304, 2004.

[17] J. Håstad, R. Impagliazzo, L. Levin, and M. Luby. Construction of a pseudo-random generator from any one-way function. *SIAM Journal on Computing*, 28(4), pp. 1364-1396, 1999.

[18] IEEE P1619.2. Draft standard architecture for wide-block encryption for shared storage media. 2008. Available from https://siswg.net.

[19] L. Keliher and J. Sui. Exact maximum expected differential and linear probability for two-round Advanced Encryption Standard. *IET Information Security*, 1(2), pp. 53–57, 2007.

[20] T. Krovetz and P. Rogaway. The software performance of authenticated-encryption modes. *FSE 2011*, LNCS 6733, Springer, pp. 306–327, 2011.

[21] R. Lampe and J. Patarin. Composition theorems for CCA cryptographic security. Cryptology ePrint report 2012/131. May 2012.

[22] M. Liskov, R. Rivest, and D. Wagner. Tweakable block ciphers. *CRYPTO 2002*, LNCS 2442, Springer, pp. 31–46, 2002

[23] D. McGrew. An interface and algorithms for authenticated encryption. RFC 5116. Jan 2008

[24] K. Minematsu. Parallelizable rate-1 authenticated encryption from pseudorandom functions. *EUROCRYPT 2014*, LNCS 8441, Springer, pp. 275–292, 2014.

[25] K. Minematsu and Y. Tsunoo. Provably secure MACs from differentially-uniform permutations and AES-based implementations. FSE 2006, LNCS 4047, Springer, pp. 226–241, 2006.

[26] M. Naor and O. Reingold. On the construction of pseudo-random permutations: Luby-Rackoff revisited. *Journal of Cryptology*, 12(1), pp. 29-66, 1999.

[27] M. Naor and O. Reingold. The NR mode of operation. Undated manuscript realizing the mechanism of [26].

[28] J. Patarin. Generic attacks on Feistel schemes. *ASIACRYPT 2001*, LNCS 2248, Springer, pp. 222–238, 2001. Also see Cryptology ePrint report 2008/036.

[29] J. Patarin. Security of balanced and unbalanced Feistel schemes with linear non equalities. Cryptology ePrint report 2010/293. May 2010.

[30] J. Patarin. Security of random Feistel schemes with 5 or more rounds. *CRYPTO 2004*, LNCS 3152, Springer, pp. 106–122, 2004.

[31] J. Patarin, B. Gittins, and J. Treger. Increasing block sizes using Feistel networks: the example of the AES. *Cryptography and Security: From Theory to Applications*, LNCS 6805, Springer, pp. 67–82, 2012.

[32] B. Preneel. Personal communications, via D. Bernstein. CAESAR competition: partial status of submissions (draft output of Dagstuhl discussion session 9 Jan 2014). Set of slides.

[33] R. Reyhanitabar and D. Vizár. Careful with misuse resistance of online AEAD. Undated manuscript (5 pages). Posted to the CAESAR mailing list Aug 24, 2014.

[34] P. Rogaway. Authenticated-encryption with associated-data. *ACM Conference on Computer and Communications Security*, ACM Press, pp. 98–107, 2002.

[35] P. Rogaway. Efficient instantiations of tweakable blockciphers and refinements to modes OCB and PMAC. *ASIACRYPT 2004*, LNCS 3329, Springer, pp. 16–31, 2004.

[36] P. Rogaway. A synopsis of format-preserving encryption. Unpublished manuscript, available from the author's webpage. Mar 2010.

[37] P. Rogaway and T. Shrimpton. A provable-security treatment of the key-wrap problem. *EUROCRYPT 2006*, LNCS 4004, Springer, pp. 373–390, 2006. Also Cryptology ePrint Report 2006/221, retitled, Deterministic authenticated-encryption: a provable-security treatment of the key-wrap problem. 2006.

[38] M. Simplício, P. Barbuda, P. Barreto, T. Carvalho, and C. Margi. The MARVIN message authentication code and the LETTERSOUP authenticated encryption scheme. *Security and Communication Networks*, 2(2), pp. 165–180, 2009.

[39] M. Simplício and P. Barreto. Revisiting the security of the ALRED Design and Two of Its Variants: Marvin and LetterSoup. *IEEE Transactions on Information Theory*, 58(9), pp. 6223–6238, 2012.

[40] R. Struik. AEAD ciphers for highly constrained networks. DIAC 2013 presentation. Chicago, Illinois, USA. Aug 13, 2013.