

Automated Analysis and Synthesis of Authenticated Encryption Schemes

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

Jonathan Katz
University of Maryland
jkatz@cs.umd.edu

Alex J. Malozemoff
University of Maryland
amaloz@cs.umd.edu

Abstract

Authenticated encryption (AE) schemes are symmetric-key encryption schemes ensuring strong notions of confidentiality and integrity. Although various AE schemes are known, there remains significant interest in developing schemes that are more efficient, meet even stronger security notions (e.g., misuse-resistance), or satisfy certain non-cryptographic properties (e.g., being patent-free).

We present an *automated* approach for analyzing and synthesizing blockcipher-based AE schemes, significantly extending prior work by Malozemoff et al. (CSF 2014) who synthesize encryption schemes satisfying confidentiality only. Our main insight is to restrict attention to a certain class of schemes that is expressive enough to capture several known constructions yet also admits automated reasoning about security. We use our approach to generate thousands of AE schemes with provable security guarantees, both known (e.g., variants of OCB and CCM) and new. Implementing two of these new schemes, we find their performance competitive with state-of-the-art AE schemes.

1 Introduction

Historically, symmetric-key encryption schemes were designed only to ensure confidentiality. With the realization that practitioners were often (implicitly) assuming that such schemes also provided some form of integrity, however, researchers began explicit consideration and analysis of encryption schemes additionally satisfying that property [15, 7]. Since then, a tremendous amount of research has focused on the design of *authenticated encryption* (AE) schemes ensuring both confidentiality and integrity.

While a generic construction of an AE scheme based on any CPA-secure encryption scheme and message authentication code is possible [7], more efficient AE schemes can be devised. One example is OCB [24, 22, 16], which is online (i.e., requires only a single pass over the data), provably secure, and very fast. Unfortunately, due to patent restrictions, the scheme never gained widespread use. Other well-known AE schemes include CCM [10] and GCM [19]; however, these schemes are slower than OCB [16] and have other disadvantages as well.¹ Overall, the problem of designing AE schemes is still of interest, as evidenced by the ongoing CAESAR competition [9].

In designing highly efficient AE schemes, one might fail to realize opportunities to improve efficiency. For example, OCB was first introduced in 2001 [24] and its authors have been active in maintaining and optimizing the scheme, releasing OCB2 in 2004 [22] and OCB3 in 2011 [16]. However, recently, Minematsu [20] showed that a simple change to OCB’s design allows one to use only the *forward direction* of the underlying blockcipher, saving chip area in hardware realizations.

In this work, we propose an *automated* approach for analyzing and synthesizing AE schemes. Our approach builds on and extends the work of Malozemoff et al. [18], who explored a similar goal but limited to encryption schemes achieving confidentiality only. At a high level, as in their work, we view an encryption

¹For example, GCM is fairly complex and has a problematic security proof [13], whereas CCM is not online and cannot pre-process associated data.

scheme as being defined by a directed acyclic graph in which each node corresponds to an instruction (e.g., XORing two values) and is associated with an intermediate n -bit value. The graph defines how individual message blocks are processed; messages of arbitrary length are encrypted by iterating the computation defined by this graph over all the blocks of the message. (We actually consider processing *two* message blocks at a time, as this allows us to capture more AE schemes within our framework.) We develop a type system for the nodes of such graphs, and define constraints on how nodes can be typed based on their parents’ types. We then show that any “well-typed” graph defines a secure AE scheme. This allows us to automatically *analyze* a given scheme by checking whether the graph defining the scheme can be properly typed. Building on this, we can *synthesize* schemes by enumerating over valid graphs and analyzing each one to see if it is secure. Through a generic transformation, our work can handle messages of arbitrary length, whereas the prior work of Malozemoff et al. is limited to messages whose length is a multiple of the block length.

Although the high-level structure of our approach is similar to that of Malozemoff et al., the technical details differ greatly due to the added challenge of handling integrity. (Indeed, this was left as an explicit open question in their work.) We were unable to *directly* extend their work to deal with integrity; instead, we modify their approach and consider a restricted class of encryption schemes for which an automated analysis of integrity is tractable. Specifically, we focus on schemes constructed using *tweakable* blockciphers [17] in a particular way². Several existing AE schemes satisfy our requirements, indicating that our framework is not overly restrictive. The abstraction from using a tweakable blockcipher also significantly reduces the size of the graphs that we have to enumerate, making the synthesis feasible. Despite this simplification, our graphs are a lot more complex than those that Malozemoff et al. consider. For example, in the prior work, it is relatively easy to tell if a scheme is decryptable, as there is only one path from a node representing a plaintext block to a node representing a ciphertext block. In our case, this no longer holds—the graphs of schemes like OTR [20] have multiple paths between such pairs of nodes, and we have to find a nontrivial algorithm to explicitly construct a graph of the encryption scheme, given a graph of the decryption scheme.

Using our approach, we are able to synthesize thousands of secure AE schemes, hundreds of which are “optimal” in the sense that they use only *one* tweakable blockcipher call per block, on par with OCB. These schemes are provably secure, as verified by our analysis tool, with *concrete* security bounds. In contrast, the prior work of Malozemoff et al. [18] only gives asymptotic analyses. We also employ a simple algorithm to find *fully parallelizable* constructions among the “optimal” schemes, and discover seventeen such schemes, five of which use the same number of instructions as OCB. We implement two of these schemes and find that the running times are comparable to those of OCB. Thus, these schemes may be of interest to practitioners looking for efficient, simple, and patent-free AE schemes. Finally, in Appendix C, we devise a method for automatically finding attacks on schemes that our approach cannot claim secure, and we find that most of those schemes indeed are susceptible to concrete attacks.

Related work. Recently there has been a growing interest in applying automated techniques to the analysis and design of cryptographic primitives. In the public-key setting, Barthe et al. [4] introduced an approach applicable to RSA-based encryption schemes. More recently, Tiwari et al. [25] developed a unified technique for synthesizing both RSA-based encryption schemes and modes of operation, among other cryptographic primitives. Other work has looked at automated analysis of assumptions in generic groups [5] with applications to automated synthesis of signature schemes having certain properties [5, 6]. Finally, Akinyele et al. [1, 2] developed tools for analyzing signature and encryption schemes to determine when (and how) known secure transformations can be applied.

2 Preliminaries

Notation. Let \mathbb{Z} denote the set of all integers, and let \mathbb{N} denote the set of positive integers. Let $\{0, 1\}^*$ denote the set of all binary strings, including the empty string. For a string M , let $|M|$ denote the length of M . For $M \in \{0, 1\}^*$ and $1 \leq i \leq j \leq |M|$, let $M[i]$ denote the i -th bit of M , and $M[i, j]$ the substring

²Roughly, a tweakable blockcipher accepts a “tweak” in addition to a key and a regular input; for a fixed key, different tweaks should produce “independent-looking” permutations. See the following section for a formal definition.

of M from the i th to the j th bit, inclusive. For two strings X and Y , we write XY or $X \parallel Y$ to denote the concatenation of X and Y .

We write $x \leftarrow S$ to denote uniform sampling of x from finite set S . For finite sets S_1, S_2 , and random variables $X, Y \in S_1, Z \in S_2$, define $\|X - Y \mid Z\|$, the statistical distance between X and Y given Z , as

$$\frac{1}{2} \sum_{v \in S_1, z \in S_2} \Pr[Z = z] \cdot \left| \Pr[X = v \mid Z = z] - \Pr[Y = v \mid Z = z] \right|.$$

Games. We use the code-based, game-playing framework of Bellare and Rogaway [8]. Due to lack of space, we assume the reader is familiar with this framework.

Tweakable blockciphers [17]. Let $n \in \mathbb{N}$. A tweakable blockcipher on n -bit strings with tweak space \mathcal{T} and key space \mathcal{K} is a map $E : \mathcal{K} \times \mathcal{T} \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ such that $E_K(T, \cdot)$ is a permutation on $\{0, 1\}^n$ for any $K \in \mathcal{K}$ and $T \in \mathcal{T}$. Let E^{-1} denote the inverse of E , meaning $E_K^{-1}(T, E_K(T, x)) = x$ for $K \in \mathcal{K}, T \in \mathcal{T}$, and $x \in \{0, 1\}^n$. For brevity we sometimes write $E_K^T(x)$ for $E_K(T, x)$. Define the *strong tweakable-PRP advantage* of an adversary \mathcal{A} against E as

$$\mathbf{Adv}_E^{\pm \widetilde{\text{PRP}}}(\mathcal{A}) = \left| \Pr[K \leftarrow \mathcal{K} : \mathcal{A}^{E_K(\cdot, \cdot), E_K^{-1}(\cdot, \cdot)} \Rightarrow 1] - \Pr[\pi \leftarrow \text{Perm}(\mathcal{T}, n) : \mathcal{A}^{\pi(\cdot, \cdot), \pi^{-1}(\cdot, \cdot)} \Rightarrow 1] \right|,$$

where $\text{Perm}(\mathcal{T}, n)$ is the set of all \mathcal{T} -indexed families of permutations on $\{0, 1\}^n$. (I.e., $\text{Perm}(\mathcal{T}, n)$ is the set of all functions $\pi : \mathcal{T} \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ with the property that for each $T \in \mathcal{T}$, the reduced function $\pi(T, \cdot)$ is a permutation on $\{0, 1\}^n$.) If the adversary is prohibited from making queries to the second oracle, we drop the word “strong” and write $\mathbf{Adv}_E^{\text{PRP}}(\mathcal{A})$ instead.

Authenticated encryption. Rather than view encryption schemes as being randomized or stateful, we follow Rogaway [23] in viewing them as deterministic transformations that take as input a message along with some associated data (which need not be kept secret) as well as a user-supplied nonce. Security is then required to hold as long as the same nonce is never used twice.

Formally, an authenticated encryption (AE) scheme [7, 15, 21, 23] is a tuple $\Pi = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ with key space \mathcal{K} , nonce space \mathcal{N} , associated data space \mathcal{A} , message space \mathcal{M} , and tag length $\tau \in \mathbb{N}$. Both algorithms \mathcal{E} and \mathcal{D} are deterministic. The encryption algorithm \mathcal{E} maps an input tuple $(K, N, A, M) \in \mathcal{K} \times \mathcal{N} \times \mathcal{A} \times \mathcal{M}$ to a ciphertext $C \in \{0, 1\}^*$. Decryption \mathcal{D} reverses encryption, mapping an input tuple $(K, N, A, C) \in \mathcal{K} \times \mathcal{N} \times \mathcal{A} \times \{0, 1\}^*$ to either a message $M \in \mathcal{M}$ or a distinguished error symbol \perp . The correctness requirement demands that $\mathcal{D}_K^{N, A}(\mathcal{E}_K^{N, A}(M)) = M$ for every $(K, N, A, M) \in \mathcal{K} \times \mathcal{N} \times \mathcal{A} \times \mathcal{M}$.

We define the privacy advantage of an adversary \mathcal{A} against an AE scheme Π as

$$\mathbf{Adv}_{\Pi}^{\text{priv}}(\mathcal{A}) = \Pr[K \leftarrow \mathcal{K} : \mathcal{A}^{\mathcal{E}_K(\cdot, \cdot, \cdot)} \Rightarrow 1] - \Pr[\mathcal{A}^{\mathcal{S}(\cdot, \cdot, \cdot)} \Rightarrow 1],$$

where $\mathcal{S}(\cdot, \cdot, \cdot)$ is an oracle that, on any input (N, A, M) , outputs a fresh, uniform $(|M| + \tau)$ -bit answer. We require here that the adversary never uses the same nonce twice as input to its oracle. Informally, a scheme satisfies privacy if the privacy advantage of any efficient adversary is small. Nonces used by the honest party during encryption need only be unique³, not uniform.

For authenticity, the adversary is again given access to an encryption oracle $\mathcal{E}_K(\cdot, \cdot, \cdot)$, and as before must not use the same nonce twice. We say that \mathcal{A} *outputs a forgery* if it outputs (N, A, C) such that $\mathcal{D}_K(N, A, C) \neq \perp$ and C was not the result of a prior oracle query $\mathcal{E}_K(N, A, M)$ for some message M . We define the authenticity advantage of \mathcal{A} as $\mathbf{Adv}_{\Pi}^{\text{auth}}(\mathcal{A}) = \Pr[K \leftarrow \mathcal{K} : \mathcal{A}^{\mathcal{E}_K(\cdot, \cdot, \cdot)} \text{ outputs a forgery}]$. Informally, a scheme satisfies authenticity if the authenticity advantage of any efficient adversary is small.

³The requirement that nonces be unique is necessary, since repeating (N, A, M) will repeat the corresponding ciphertext. For real schemes such as OCB, reusing a nonce is devastating, damaging the privacy and authenticity of not just past queries, but also future ones. It is the responsibility of the implementation to ensure that nonces are unique.

3 Automated Security Analysis

We now describe our approach to the automated analysis of AE schemes constructed from tweakable blockciphers following a particular template (cf. Section 3.1). Although this template does not capture all known AE schemes, it is expressive enough to include simplified variants of, e.g., OCB [22], XCBC [12], COPA [3], OTR [20], and CCM [10].⁴

As discussed in the Introduction, we view an encryption scheme as being defined by a directed acyclic graph in which each node is associated with an instruction and carries an n -bit intermediate value. In Section 3.2 we describe a type system for the nodes of such graphs, and show how to use these types for reasoning about properties of the intermediate values that those nodes carry. Then, in Section 3.3, we show how this reasoning enables us to automatically verify whether an AE scheme, given by its graph representation, satisfies privacy and authenticity.

3.1 A Template for AE Schemes

Fix associated data space \mathcal{A} , and let $\mathcal{N} = \{0, 1\}^n$. Let $\mathcal{T} = \mathcal{N} \times \mathcal{A} \times \mathbb{Z}$ and let $E : \mathcal{K} \times \mathcal{T} \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ be a tweakable blockcipher.⁵ We consider AE schemes $\Pi[E] = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ that use E as an oracle. The schemes we consider have message space⁶ $\mathcal{M} = (\{0, 1\}^{2n})^*$ and are built from algorithms $(\text{Enc}, \text{Dec}, \text{Tag})$ having the following form:

- Enc^{E_K} takes as input tweak $T = (N, A, v) \in \mathcal{T}$, an initial state $X \in \{0, 1\}^{2n}$, and a (double-length) message block $M \in \{0, 1\}^{2n}$. It outputs a (double-length) ciphertext block $C \in \{0, 1\}^{2n}$ and final state $Y \in \{0, 1\}^{2n}$. This algorithm makes a fixed number of queries to E_K , denoted by $\text{Cost}(\Pi)$, and we require that the tweak in the i th such query is $(N, A, v + i - 1)$.
- $\text{Dec}^{E_K, E_K^{-1}}$ “inverts” algorithm Enc in the following sense: if $\text{Enc}^{E_K}(T, X, M) = (Y, C)$ then it holds that $\text{Dec}^{E_K, E_K^{-1}}(T, X, C) = (Y, M)$.
- Tag^{E_K} takes as input tweak $T \in \mathcal{T}$ and initial state $X \in \{0, 1\}^{2n}$, and produces a tag $V \in \{0, 1\}^n$. It makes a single query to E_K using tweak T .

The encryption/decryption algorithms $(\mathcal{E}, \mathcal{D})$ of Π are then defined as in Figure 3.1, where we require $\tau \leq n$. Roughly, to encrypt a message $M = M_1 M_2 \cdots M_{2m}$ using nonce N and associated data A , set the initial state $X = 0^{2n}$ and set $T = (N, A, 1)$. Then, iteratively process two message blocks at a time using Enc , each time updating the initial state and outputting the next two ciphertext blocks. After processing the entire message, Tag is used to compute a tag based on the final state output by Enc and a designated tweak that depends on the message length; the (truncated) tag is appended to the ciphertext.

Graph representation. As in the work of Malozemoff et al. [18], we represent algorithms Enc , Dec , and Tag as directed acyclic graphs, where each node is associated with an instruction and carries an n -bit value. The n -bit value on each node is determined by applying the instruction at that node to the values at the parent nodes. In the next section we introduce a system for “typing” the nodes of such graphs; our main theorem states that AE schemes built from Enc , Dec , and Tag algorithms whose graphs can be correctly typed are secure.

The main instructions we support are **XOR**, which computes the XOR of two n -bit strings, and **TBC**, which invokes the tweakable blockcipher or its inverse. We also have an instruction **DUP** that duplicates a

⁴For efficiency, the real-world variants are often built directly from a blockcipher instead of a tweakable one, and employ a scheme-specific way to handle fragmentary data. The real-world CCM is not online due to its treatment of fragmentary data, whereas our variant is online. OCB is built from a tweakable blockcipher, but the tweaks are (N, i) instead of (N, A, i) . To handle associated data, OCB employs an XOR-universal hash (based on a tweakable blockcipher), and XORs the hash image to the tag.

⁵One can extend the XEX construction [22] of a tweakable blockcipher from any blockcipher to tweak space $\mathcal{N} \times \{0, 1\}^* \times \mathbb{Z}$ as follows: On tweak (N, A, i) , one applies a (keyed) universal hash to (N, A) to derive a synthetic nonce N' , and apply the XEX construction on (N', i) . By buffering L , hashing need only be done once per message.

⁶Messages of arbitrary length can be handled by naive padding in the usual way. In Appendix B we describe a more efficient approach for handling messages of arbitrary length.

```

 $\mathcal{E}_K(N, A, M)$ 
 $X := 0^{2n}; v := 1; M_1 \cdots M_{2m} := M \quad // |M_i| = n$ 
for  $i = 1$  to  $m$  do
   $T := (N, A, v)$ 
   $(Y, C_{2i-1}C_{2i}) := \text{Enc}^{E_K}(T, X, M_{2i-1}M_{2i}) \quad // |C_j| = n$ 
   $v := v + \text{Cost}(\Pi); X := Y$ 
 $T := (N, A, 1 - v); V := \text{Tag}^{E_K}(T, X)$ 
return  $C_1 \cdots C_{2m} \parallel V[1, \tau]$ 

 $\mathcal{D}_K(N, A, C)$ 
if  $|C| \not\equiv \tau \pmod{2n}$  then return  $\perp$ 
 $C_1 \cdots C_{2m} \parallel \text{tag} := C \quad // |C_i| = n \text{ and } |\text{tag}| = \tau$ 
 $X := 0^{2n}; v := 1$ 
for  $i = 1$  to  $m$  do
   $T := (N, A, v)$ 
   $(Y, M_{2i-1}M_{2i}) := \text{Dec}^{E_K, E_K^{-1}}(T, X, C_{2i-1}C_{2i})$ 
   $v := v + \text{Cost}(\Pi); X := Y$ 
 $T := (N, A, 1 - v); V := \text{Tag}^{E_K}(T, X)$ 
if  $\text{tag} \neq V[1, \tau]$  then return  $\perp$  else return  $M_1 \cdots M_{2m}$ 

```

Figure 3.1: Code of an AE scheme $\Pi = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ following our template. The scheme is based on a tweakable blockcipher E and a triple of deterministic algorithms ($\text{Enc}, \text{Dec}, \text{Tag}$).

value. Nodes corresponding to input blocks are labeled IN, those corresponding to output blocks are labeled OUT, those corresponding to the initial state are labeled INI, and those corresponding to the final state are labeled FIN. These labels, along with their in-/out-degree, are summarized for convenience next:

Name	In-deg	Out-deg	Meaning
IN	0	1	Input block
INI	0	1	Initial state
FIN	1	0	Final state
OUT	1	0	Output block
DUP	1	2	Duplicate
XOR	2	1	XOR operation
TBC	1	1	Tweakable blockcipher

Figure 3.2 illustrates the OCB scheme [22], Figure 3.3 shows the corresponding Enc, Dec , and Tag algorithms, and Figure 3.4 shows the corresponding graphs. (In OCB, only the first n bits of the state are used, so we treat the state as an element of $\{0, 1\}^n$.) Note that Figure 3.4 is informal and omits information needed to fully specify OCB; see next for formal details of how graphs are specified.

Formally, we denote a graph G by a tuple (d, r, F, P, L) , where $d \in \{2, 4\}$ is the total number of IN and INI nodes (Enc and Dec graphs have $d = 4$; Tag graphs have $d = 2$), and $r \in \mathbb{N}$ is the total number of nodes. Each node in the graph is numbered from 1 to r , and we require that if node i is a parent of node j then $i < j$. (This ensures that G is acyclic.) Let Nodes denote the power set of $\{1, \dots, r\}$, and let $\text{Inst} = \{\text{IN}, \dots, \text{TBC}\}$ be the set of instructions. Then $F : \{1, \dots, r\} \rightarrow \text{Inst}$ gives the instruction of each node and $P : \{1, \dots, r\} \rightarrow \text{Nodes}$ gives the set of parents for each node. We require that $F(1) = F(2) = \text{INI}$ and $F(r) = \text{OUT}$. For Enc and Dec graphs, we additionally require that $F(3) = F(4) = \text{IN}$, $F(r-2) = F(r-3) = \text{FIN}$, and $F(r-1) = \text{OUT}$.

For Enc and Dec graphs, let $S \subset \{1, \dots, r\}$ be the set of all nodes corresponding to a TBC instruction. Function $L : S \rightarrow \mathbb{Z}$ specifies, for each such node i , whether the tweakable blockcipher is computed in the forward direction (if $L(i) \geq 0$) or the reverse direction (if $L(i) < 0$) at that node.⁷ (Note that $L(i) \geq 0$ for

⁷While it might be conceptually simpler to use two different instructions for E_K and E_K^{-1} , instead of just a single TBC instruction with positive/negative labels, our approach is an optimization that prunes the search space when synthesizing




Figure 3.2: The OCB scheme illustrated for a four-block message M_1, \dots, M_4 , where Σ is the checksum $M_1 \oplus \dots \oplus M_4$.

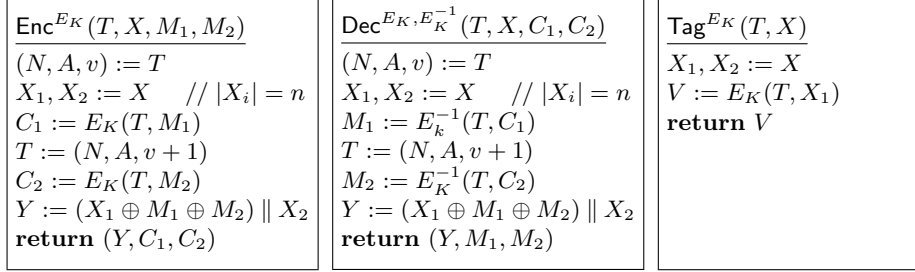


Figure 3.3: The algorithms (Enc, Dec, Tag) corresponding to OCB. We have $\text{Cost}(\text{OCB}) = 2$.




Figure 3.4: Graph representations for algorithms Enc (left), Dec (middle), and Tag (right) of OCB.

Enc graphs.) Moreover, $|L(i)|$ determines the tweak at node i ; i.e., on input tweak (N, A, v) , the tweak at node i is $(N, A, v + |L(i)|)$.

Let $G^- = (d, r, F, P)$ denote the *unlabeled graph* corresponding to a graph G . In Section 3.2, we introduce a type system and show that one can reason about the security of an AE scheme by evaluating the scheme’s unlabeled graphs.

Fix some graph G , tweakable blockcipher E , and key K . Given a tweak T and n -bit values for all INI/IN nodes in G , we can naturally define an n -bit value Z_i associated with each node i in the graph. We describe this formally as procedure Eval in Figure 3.5, which shows how to compute Z_i given values $Z_1, \dots, Z_d \in \{0, 1\}^n$ and tweak T .

3.2 A Type System for AE Schemes

Let $\text{Types} = \{\$, \perp, 0, 1\}$ be a set of “types” we can assign to nodes. Intuitively, ‘\$’ indicates a node whose output value is (pseudo)random (when the key K for E is random and secret), whereas ‘ \perp ’ indicates a node whose output value is arbitrary (i.e., potentially controlled by an attacker). Looking ahead, types ‘0’ and ‘1’ will be used to compare values on the same node in two different decryption queries using the same nonce and associated data; ‘0’ means the corresponding values are the same, and ‘1’ means they are different.

schemes (cf. Section 4).

```

proc EvalEK, EK-1(G, T, Z1, ..., Zd)
(d, r, F, P, L) := G
for i = d + 1 to r do
  if F(i) ∈ {DUP, OUT, FIN} then {p} := P(i); Zi := Zp
  elseif F(i) = XOR then {p1, p2} := P(i); Zi := Zp1 ⊕ Zp2
  else // F(i) = TBC
    ℓ := L(i); (N, A, v) := T
    T* := (N, A, v + |ℓ|); {p} := P(i)
    if ℓ > 0 then Zi := EK(T*, Zp) else Zi := EK-1(T*, Zp)
return (Z1, ..., Zr)

```

Figure 3.5: Procedure to compute the value Z_i of each node i in a graph G , given input Z_1, \dots, Z_d and tweak T .

```

proc Map( $G^-$ , type1, ..., typed, rand)
(d, r, F, P) :=  $G^-$ ; maxCtr := 0
for i = 1 to d do
  if typei = $ then R(i) := ($, 1); maxCtr := 1
  else R(i) := (typei, 0)
for i = d + 1 to r do
  if F(i) ∈ {FIN, OUT, DUP} then
    {p} := P(i); R(i) := R(p)
  elseif F(i) = TBC then
    {p} := P(i); (x, ctr) := R(p)
    if x ∈ {1, $} or (rand = true) then
      maxCtr := maxCtr + 1; R(i) := ($, maxCtr)
    else R(i) := (x, ctr)
  else // F(i) = XOR
    {p1, p2} := P(i); (x, ctr) := R(p1); (y, ctr') := R(p2)
    // Assume that ctr ≥ ctr'
    if (x, y) ∈ {(0, 0), (0, 1), (1, 0)} then R(i) := (x ⊕ y, ctr)
    elseif x = $ and ctr > ctr' then R(i) := ($, ctr)
    else R(i) := (⊥, ctr)
return R

```

Figure 3.6: A procedure for generating a mapping $R : \{1, \dots, r\} \rightarrow Types \times \mathbb{N}$ for a given unlabeled graph.

In Figure 3.6, we define a deterministic procedure `Map` that takes as input an unlabeled graph G^- , pre-assigned types $type_1, \dots, type_d \in Types$ for all the INI/IN nodes, and a boolean flag $rand$, and returns a map R that associates each node i with a pair $(type_i, ctr_i) \in Types \times \mathbb{N}$. The procedure `Map` traverses the graph in topological order and assigns types to each node of the graph based on the instruction associated with that node and the types of its parents. These types are used for probabilistic reasoning about the underlying n -bit values on that node; e.g., we show that if a node has type $\$$ then the n -bit value of that node is (pseudo)random. The ctr values are used as “timestamps” for values output by TBC nodes in order to determine independence among values of type $\$$. Finally, the $rand$ flag denotes whether the nonce/associated data are fresh.

For FIN, OUT, and DUP nodes, `Map` simply propagates the type of the parent node. For TBC nodes, if the nonce or input is fresh then the output is (pseudo)random and independent of any prior random values, and so the node gets type $\$$; otherwise, we propagate the type of the parent node.

For XOR nodes, we have several cases. If the two input nodes x and y are typed $type_x$ and $type_y$, respectively, with $(type_x, type_y) \in \{(0, 0), (0, 1), (1, 0)\}$, then we type the XOR node as $type_x \oplus type_y$. We briefly explain this reasoning. The fact that $type_x, type_y \in \{0, 1\}$ means there is a prior query using the same nonce and associated data. If the two parents have type 0, indicating that the values computed at

those nodes are equal in the two queries, then clearly the value computed at the XOR node is also equal in the two queries, and thus that node gets type 0. On the other hand, if one parent is typed 0 and the other is typed 1, then the value computed at the XOR node will be different from the corresponding value in the prior query, and thus the XOR node is assigned type 1. If $(\text{type}_x, \text{type}_y) = (1, 1)$ then we cannot say anything definitive and thus Map assigns type \perp to the XOR node. Finally, suppose input node x has type $\$$. Here we utilize the ctr values. If the ctr value at x is different from the ctr value of y , then the (random) value of x is independent of the value of y , and hence we assign the XOR node type $\$$.

In the next two lemmas we show how determining the types for an *unlabeled* graph can be used to reason about the values that one obtains when evaluating the *labeled* graph. We first show that all values typed $\$$ by Map (when inputs are typed \perp and hence may be under arbitrary control of the adversary) are indeed random when computed using Eval and a truly random tweakable permutation.

Lemma 3.1. Let $G = (d, r, F, P, L)$ and let $n \geq 1$ be an integer. Set $R := \text{Map}(G^-, \perp, \dots, \perp, \text{true})$. Fix arbitrary $Z_1, \dots, Z_d \in \{0, 1\}^n$ and $T \in \mathcal{T}$, and consider the following probabilistic experiment:

1. Choose $f \leftarrow_s \text{Perm}(\mathcal{T}, n)$.
2. Run $(Z_1, \dots, Z_r) := \text{Eval}^{f, f^{-1}}(G, T, Z_1, \dots, Z_d)$.

Then for any j with $R(j) = (\$, \text{ctr}_j)$, the random variable Z_j is uniform and independent of $\{Z_i \mid \text{ctr}_i < \text{ctr}_j\}$.

Proof. First note that for any node i and its parent p , we have $\text{ctr}_i \geq \text{ctr}_p$. Thus, there is a topological ordering s_1, \dots, s_r of the nodes such that the sequence $\text{ctr}_{s_1}, \dots, \text{ctr}_{s_r}$ is non-decreasing, and $s_i = i$ for $i \leq d$. Write $i \prec j$ if node i precedes node j in this topological order. We prove by induction (with respect to \prec) that for all j we have (i) $\text{type}_j \in \{\perp, \$\}$ and (ii) if $\text{type}_j = \$$ then Z_j is uniform and independent of $\{Z_i \mid \text{ctr}_i < \text{ctr}_j\}$. Note that these claims are trivially true when $j \leq d$, because then $\text{type}_j = \perp$.

Suppose both claims hold for all $i \prec j$. If $F(j) \in \{\text{FIN}, \text{OUT}, \text{DUP}\}$ then let $p \prec j$ be the parent of j . Since $(\text{type}_j, \text{ctr}_j) = (\text{type}_p, \text{ctr}_p)$, the claims follow for j . If $F(j) = \text{TBC}$ then $\text{type}_j = \$$, proving claim (i). Since $f \leftarrow_s \text{Perm}(\mathcal{T}, n)$, and Eval never repeats a tweak in querying f , we see that random variable Z_j is uniform and independent of $\{Z_i \mid \text{ctr}_i < \text{ctr}_j\}$, justifying claim (ii). Finally, say $F(j) = \text{XOR}$. Let i and t be the parents of j , and assume $t \prec i$. Then $\text{type}_t, \text{type}_i \in \{\perp, \$\}$, and thus so is type_j , proving claim (i). For claim (ii), note that $\text{type}_j = \$$ only if $\text{type}_i = \$$ and $\text{ctr}_i > \text{ctr}_t$. Since $\text{ctr}_j = \text{ctr}_i$, the claim follows. \square

The next lemma proves a similar property as above for *pairs* of queries. Consider the query $(Y_1, \dots, Y_r) := \text{Eval}^{f, f^{-1}}(G, T, Y_1, \dots, Y_d)$ followed by query $(Z_1, \dots, Z_r) := \text{Eval}^{f, f^{-1}}(G, T, Z_1, \dots, Z_d)$, where each Z_i is either chosen equal to Y_i (and thus $\text{type}_i = 0$), distinct from Y_i (and thus $\text{type}_i = 1$), or uniformly (and thus $\text{type}_i = \$$). We show that for all nodes j of type $\$$ assigned by $\text{Map}(G^-, \text{type}_1, \dots, \text{type}_d, \text{false})$, the statistical difference between Z_j and uniform is small, even conditioned on all the $\{Y_i\}$.

Lemma 3.2. Let $G = (d, r, F, P, L)$ and let $n \geq 1$ be an integer. Fix arbitrary $Y_1, \dots, Y_r \in \{0, 1\}^n$ and $T \in \mathcal{T}$ such that the set $S = \{f \in \text{Perm}(\mathcal{T}, n) \mid (Y_1, \dots, Y_r) = \text{Eval}^{f, f^{-1}}(G, T, Y_1, \dots, Y_d)\}$ is non-empty. For each $i \leq d$, choose Z_i and type_i in one of the following ways: (i) $Z_i = Y_i$ and $\text{type}_i = 0$, (ii) $Z_i \neq Y_i$ and $\text{type}_i = 1$, or (iii) $Z_i \leftarrow_s \{0, 1\}^n$ and $\text{type}_i = \$$. Let $R = \text{Map}(G^-, \text{type}_1, \dots, \text{type}_d, \text{false})$. Consider the following probabilistic experiment:

1. Choose $f \leftarrow_s S$.
2. Run $(Z_1, \dots, Z_r) := \text{Eval}^{f, f^{-1}}(G, T, Z_1, \dots, Z_d)$.

Then for any j with $R(j) = (\$, \text{ctr}_j)$, the statistical difference between the random variable Z_j and uniform, conditioned on $\{Z_i \mid \text{ctr}_i < \text{ctr}_j\}$ and all the $\{Y_i\}$, is at most $2 \cdot \text{ctr}_j / 2^n$.

Proof. As in the previous lemma, there is a topological ordering s_1, \dots, s_r of the nodes such that $\text{ctr}_{s_1}, \dots, \text{ctr}_{s_r}$ is non-decreasing and $s_i = i$ for $i \leq d$. Write $i \prec j$ if i precedes j in this topological order. We prove by induction (with respect to \prec) that for all j : (i) if $\text{type}_j = 0$ then $Z_j = Y_j$, (ii) if $\text{type}_j = 1$ then $Z_j \neq Y_j$, and (iii) if $\text{type}_j = \$$ then the statement of the lemma holds. These claims all trivially hold when $j \leq d$.

Suppose that all three of the claims hold for all $i \prec j$. If $F(j) \in \{\text{FIN}, \text{OUT}, \text{DUP}\}$ then let $p \prec j$ be the parent of j . Since $(\text{type}_j, \text{ctr}_j) = (\text{type}_p, \text{ctr}_p)$, the claims follow easily in this case. If $F(j) = \text{XOR}$, let $t, i \prec j$ be the parents of j , and assume $t \prec i$. Note that $\text{type}_j \in \{0, 1\}$ only if $\text{type}_i, \text{type}_t \in \{0, 1\}$ and at most one of these values is 1, in which case the claims all hold. On the other hand, $\text{type}_j = \$$ only if $\text{type}_i = \$$ and $\text{ctr}_i > \text{ctr}_t$, in which case the claims also follow. Finally, if $F(j) = \text{TBC}$ then let i be the parent of j . Let $\ell = L(j)$, and let $T = (N, A, v)$. Then $Y_j = f(T^*, Y_p)$ and $Z_j = f(T^*, Z_p)$, where $T^* = (N, A, v + |\ell|)$. Consider the following cases:

Case 1. $\text{type}_i \in \{0, \perp\}$. Then $(\text{type}_j, \text{ctr}_j) = (\text{type}_i, \text{ctr}_i)$ and the claims follow.

Case 2. $\text{type}_i = 1$. Then $\text{type}_j = \$$ and $\text{ctr}_j \geq 1$. First, since $\text{ctr}_t \geq \text{ctr}_j$ when t is a descendant of j , we see that no node t with $\text{ctr}_t < \text{ctr}_j$ is a descendant of j . Next, since $Z_i \neq Y_i$ and we use a different tweak for each TBC node, $Z_j \leftarrow_s \{0, 1\}^n \setminus \{Y_j\}$ is independent of $\{Y_t \mid t \leq r\}$ and $\{Z_t \mid \text{ctr}_t < \text{ctr}_j\}$. Hence the statement of the lemma follows.

Case 3. $\text{type}_i = \$$. Then $\text{type}_j = \$$. By the induction hypothesis, Z_i is $(2\text{ctr}_i/2^n)$ -close to uniform (even conditioned on $\{Y_t \mid t \leq r\}$ and $\{Z_t \mid \text{ctr}_t < \text{ctr}_i\}$). If $Z_i \neq Y_i$, which occurs except with probability at most $(2\text{ctr}_i+1)/2^n$, then Z_j is 2^{-n} -close to uniform (even conditioned on all the $\{Y_t\}$ values and $\{Z_t \mid \text{ctr}_t < \text{ctr}_j\}$). Hence, overall, Z_j is $(2\text{ctr}_i+2)/2^n$ -close to uniform (conditioned on $\{Y_t \mid t \leq r\}$ and $\{Z_t \mid \text{ctr}_t < \text{ctr}_j\}$) and the statement of the lemma follows since we have $\text{ctr}_j \geq \text{ctr}_i + 1$. \square

3.3 Verifying Privacy and Authenticity

We use Lemmas 3.1 and 3.2 to automatically check if a candidate AE scheme is secure in the sense of both privacy and authenticity. Specifically, Figure 3.7 shows procedures `Priv` and `Auth` to check for privacy and authenticity, respectively, of an AE scheme Π .

Intuitively, for privacy we verify that the tag and all the ciphertext blocks output by the scheme are random and independent (namely, have type $\$$ and distinct counter values) even when the inputs—that is, the message blocks—are controlled by the adversary (namely, have type \perp). We remark that the values of ctr assigned to nodes by the map R output by `Map` depend on the topological order in which `Map` traverses the input graph; see Figure 3.8 for an example. Thus, there *are* schemes which, depending on the order in which the graph is traversed, are accepted or (incorrectly) rejected by `Priv` (due to the ctr values for the `OUT` nodes being equal). This shows that the test is *sound* but not *complete*.⁸

The authenticity check for a scheme $(\text{Enc}, \text{Dec}, \text{Tag})$ is more complicated. We now argue informally that if a scheme passes the checks of algorithm `Auth` (cf. Figure 3.7), then the scheme satisfies authenticity. To see this, consider a candidate forgery (N, A, C) output by an adversary. First suppose there was no prior query (N, A, \star) to the encryption oracle. `Auth` verifies that the `Tag` algorithm outputs a random tag when the tweak for the TBC node in `Tag` was not used previously; thus, the candidate forgery will be invalid except with probability $2^{-\tau}$. (Recall that τ is the tag length.) Next, consider the case that there was a prior encryption query (N, A, M) , and let C' be the corresponding ciphertext. Then $C \neq C'$; otherwise (N, A, C) is not a valid forgery. If C and C' only differ in their tags, the candidate forgery must be invalid because the tag is uniquely determined by N, A , and the rest of the ciphertext. Otherwise, consider the first pair of blocks in which C and C' differ. `Auth` verifies that (i) the first half of the final state produced by `Dec` when run on those blocks is random, (ii) `Dec` has the property that if the first half of its initial state is random, then the first half of the final state it outputs is random, and (iii) `Tag` has the property that if the first half of its initial state is random, then the tag it outputs is random⁹. Taken together, these imply that the tag will be random, and hence the candidate forgery will be invalid except with probability $2^{-\tau}$.

⁸In Appendix C we describe a technique for generating attacks given a scheme which fails the tests of Figure 3.7. Looking ahead, we find only a handful of schemes which we can neither prove secure nor find concrete attacks for; see Section 4.

⁹Although here we are considering just the first half of the final/initial state, if one switches to the second half then one will get the *same* set of synthesized schemes: if one changes the topological ordering in the graphs so that the first `FIN/INI` node becomes the second one, and vice versa, then the scheme remains the same.

```

proc Priv( $G_1^-, G_2^-$ )
  //  $G_1^-$  and  $G_2^-$  are unlabeled graphs of Enc and Tag, respectively.
01 ( $d_1, r_1, F_1, P_1$ ) :=  $G_1^-$ ; ( $d_2, r_2, F_2, P_2$ ) :=  $G_2^-$ 
  // Check that output of Tag is random
02  $R := \text{Map}(G_2^-, \perp, \perp, \text{true})$ ; ( $type, ctr$ ) :=  $R(r_2)$ 
03 if  $type \neq \$$  then return false
  // Check that output blocks of Enc are random and independent
04  $R := \text{Map}(G_1^-, \perp, \perp, \perp, \perp, \text{true})$ 
05 ( $type_1, ctr_1$ ) :=  $R(r_1 - 1)$ ; ( $type_2, ctr_2$ ) :=  $R(r_1)$ 
06 return ( $(type_1 = \$) \wedge (type_2 = \$) \wedge (ctr_1 \neq ctr_2)$ )

proc Auth( $G_1^-, G_2^-$ )
  //  $G_1^-$  and  $G_2^-$  are unlabeled graphs of Dec and Tag, respectively.
11 ( $d_1, r_1, F_1, P_1$ ) :=  $G_1^-$ ; ( $d_2, r_2, F_2, P_2$ ) :=  $G_2^-$ 
  // Check that output of Tag is random when the nonce/associated data are fresh
12  $R := \text{Map}(G_2^-, \perp, \perp, \text{true})$ ; ( $type, ctr$ ) :=  $R(r_2)$ 
13 if  $type \neq \$$  then return false
  // Check that if there are two executions of Dec with the same initial state
  // but different input blocks, then the first half of the final state is random
14 for  $(x, y) \in \{(0, 1), (1, 0), (1, 1)\}$  do
15  $R := \text{Map}(G_1^-, 0, 0, x, y, \text{false})$ ; ( $type, ctr$ ) :=  $R(r_1 - 3)$ 
16 if  $type \neq \$$  then return false
  // Check that if the first half of the initial state input to Dec is random,
  // then the first half of the final state output by Dec is random
17 for  $x, y, z \in \{0, 1\}$  do
18  $R := \text{Map}(G_1^-, \$, x, y, z, \text{false})$ ; ( $type, ctr$ ) :=  $R(r_1 - 3)$ 
19 if  $type \neq \$$  then return false
  // Check that if there are two executions of Tag in which the first halves of the
  // initial states are different, then the resulting tags are random and independent
20 for  $x \in \{0, 1\}$  do
21  $R := \text{Map}(G_2^-, 1, x, \text{false})$ ; ( $type, ctr$ ) :=  $R(r_2)$ 
22 if  $type \neq \$$  then return false
23 return true

```

Figure 3.7: Tests to determine if a scheme Π satisfies privacy and authenticity, respectively.

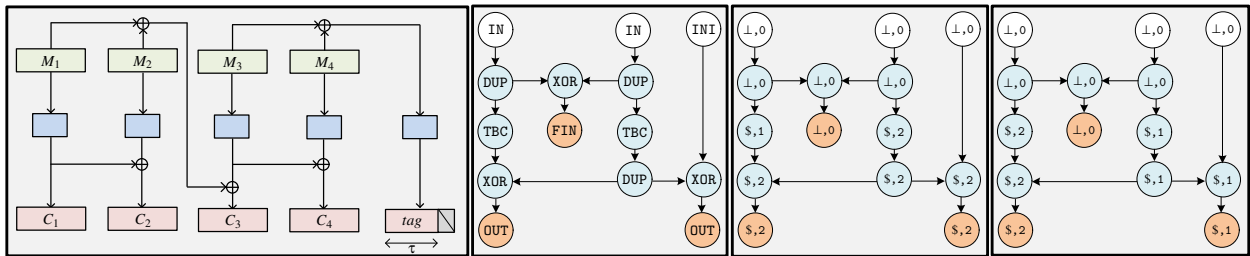


Figure 3.8: **Left:** A scheme that can be accepted or (incorrectly) rejected by Priv, depending on the topological ordering of the nodes. **Middle left:** The corresponding Enc graph. **Middle right:** The $(type, ctr)$ pairs in each node of the Enc graph if the left TBC node is visited first. The graph is (incorrectly) rejected because the two OUT nodes both have $ctr = 2$. **Right:** The $(type, ctr)$ pairs in each node of the Enc graph if the right TBC is visited first. This time, the graph is accepted because the two OUT nodes have different ctr values.

To demonstrate the strength of our approach, consider a modified version of the OTR scheme [20]. The original OTR scheme (cf. Figure 3.9) is secure, which our automated tests confirm. If, however, the scheme

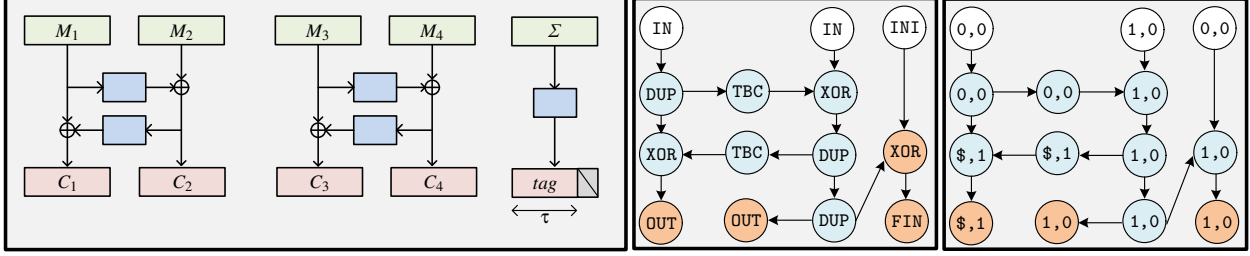


Figure 3.9: **Left:** The OTR scheme, illustrated for a four-block message $M_1 \cdots M_4$. Here, Σ is the checksum of the even blocks $M_2 \oplus M_4$. **Middle:** The Dec graph of the *insecure* variant of OTR, where Σ is instead the checksum of the odd blocks $M_1 \oplus M_3$. **Right:** The demonstration that the Dec graph of the insecure OTR variant does not pass the test at lines 14–16 of Figure 3.7. Each node is shown with its $(type, ctr)$ pair. The graph fails the test because the FIN node has type 1 instead of \$.

is changed so that Σ is computed as the checksum of the *odd* blocks $M_1 \oplus M_3 \oplus \cdots$, rather than the even blocks, then it becomes insecure. And, indeed, the modified scheme does not pass our tests. Namely, on input $(0, 0, 1, 0)$ to Map we find that the required FIN node is typed 1 instead of \$.

Proofs of correctness. We now prove that schemes that pass our tests are secure. We first show that if Priv returns true when given the (unlabeled) graphs corresponding to the Enc and Tag components of some AE scheme, then that scheme satisfies privacy when instantiated with a secure tweakable blockcipher.

Theorem 3.3. Let $\Pi[E] = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ be an AE scheme for which $\text{Priv}(G_1^-, G_2^-) = \text{true}$, where G_1^- and G_2^- are the unlabeled graphs for algorithms Enc and Tag of Π , respectively. Then for any adversary \mathcal{A} , there is an adversary \mathcal{B} with $\text{Adv}_{\Pi[E]}^{\text{priv}}(\mathcal{A}) \leq \text{Adv}_E^{\text{prp}}(\mathcal{B})$. Adversary \mathcal{B} has the same running time as \mathcal{A} and makes at most $(\text{Cost}(\Pi) + 1) \cdot \sigma/2$ queries, where σ is the number of message blocks in the queries of \mathcal{A} .

Proof. Adversary \mathcal{B} runs \mathcal{A} . For each of \mathcal{A} 's queries (N, A, M) , adversary \mathcal{B} runs the encryption scheme $\Pi[E]$ on (N, A, M) with each call to E_K replaced by a query to \mathcal{B} 's oracle, and returns the ciphertext to \mathcal{A} . Finally, \mathcal{B} outputs the same guess as \mathcal{A} . Let $\Pi[\pi]$ be the ideal variant of $\Pi[E]$, where calls to E_K are replaced by corresponding queries to π , with $\pi \leftarrow_s \text{Perm}(\mathcal{T}, n)$. It suffices to show that $\text{Adv}_{\Pi[\pi]}^{\text{priv}}(\mathcal{A}) = 0$.

Consider experiments H_1 – H_4 in Figure 3.10. The adversary has oracle access to the encryption scheme of $\Pi[\pi]$ in experiment H_1 , and oracle access to (\cdot, \cdot, \cdot) in experiment H_4 . Experiment H_2 is identical to H_1 , except that we re-sample $\pi \leftarrow_s \text{Perm}(\mathcal{T}, n)$ each time we use Enc or Tag. Since a tweak to π is never repeated, $\Pr[H_1^A \Rightarrow \text{true}] = \Pr[H_2^A \Rightarrow \text{true}]$. In experiment H_3 , instead of calling $\text{Tag}^\pi(T, X)$ to get the tag, we sample the tag at random. Considering lines 02–03 of Priv (and the fact that $\text{Priv}(G_1^-, G_2^-) = \text{true}$) in conjunction with Lemma 3.1 shows that the string $V := \text{Tag}^\pi(T, X)$ is uniform and so experiments H_2 and H_3 are identical. Finally, experiment H_4 is identical to H_3 , except that instead of calling $\text{Enc}^\pi(T, X, M_{2i-1}M_{2i})$ to get the blocks $C_{2i-1}C_{2i}$ of the ciphertext, we sample them at random. Considering lines 04–05 of Priv (and the fact that $\text{Priv}(G_1^-, G_2^-) = \text{true}$) in conjunction with Lemma 3.1 shows that the output blocks of $\text{Enc}^\pi(T, X, M_{2i-1}M_{2i})$ are uniform and independent (and this is true even conditioned on all prior ciphertext blocks). Hence H_3 and H_4 are identical, and $\text{Adv}_{\Pi[\pi]}^{\text{priv}}(\mathcal{A}) = \Pr[H_1^A \Rightarrow \text{true}] - \Pr[H_4^A \Rightarrow \text{true}] = 0$. \square

Next, in Theorem 3.4, we show that if Auth in Figure 3.7 returns true when given graphs corresponding to the Dec and Tag components of some AE scheme, then that scheme satisfies authenticity when instantiated with a secure tweakable blockcipher. (Examination of the proof shows that if algorithm Dec does not use E_K^{-1} , as in the case of OTR, then the term $\text{Adv}_E^{\pm \text{prp}}(\mathcal{B})$ in Theorem 3.4 can be weakened to $\text{Adv}_E^{\text{prp}}(\mathcal{B})$.)

Theorem 3.4. Let $\Pi[E] = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ be an AE scheme such that $\text{Auth}(G_1^-, G_2^-) = \text{true}$, where G_1^- , G_2^- are the unlabeled graphs for algorithms Dec and Tag of Π , respectively. Then for any adversary \mathcal{A} , there is an adversary \mathcal{B} with $\text{Adv}_{\Pi[E]}^{\text{auth}}(\mathcal{A}) \leq 2^{-\tau} + \ell \cdot (\text{Cost}(\Pi) + 2)/2^n + \text{Adv}_E^{\pm \text{prp}}(\mathcal{B})$, where ℓ is the number

<pre> proc ENCRYPT[π](N, A, M) // Experiments H_1, H_2 $M_1 \cdots M_{2m} := M; X := 0^{2n}; v := 1$ // $M_i = n$ for $i = 1$ to m do $T := (N, A, v); \pi \leftarrow \text{Perm}(\mathcal{T}, n)$ $(Y, C_{2i-1}C_{2i}) := \text{Enc}^\pi(T, X, M_{2i-1}M_{2i})$ $v := v + \text{Cost}(\Pi); X := Y$ $\pi \leftarrow \text{Perm}(\mathcal{T}, n); T := (N, A, 1 - v); V := \text{Tag}^\pi(T, X)$ return $C_1 \cdots C_{2m} \parallel V[1, \tau]$ </pre>
<pre> proc ENCRYPT[π](N, A, M) // Experiments H_3, H_4 $M_1 \cdots M_{2m} := M; X := 0^{2n}; v := 1$ // $M_i = n$ for $i = 1$ to m do $T := (N, A, v); \pi \leftarrow \text{Perm}(\mathcal{T}, n)$ $(Y, C_{2i-1}C_{2i}) := \text{Enc}^\pi(T, X, M_{2i-1}M_{2i})$ $C_{2i-1}C_{2i} \leftarrow \{0, 1\}^{2n}$ $v := v + \text{Cost}(\Pi); X := Y$ $T := (N, A, 1 - v); V \leftarrow \{0, 1\}^n$ return $C_1 \cdots C_{2m} \parallel V[1, \tau]$ </pre>

Figure 3.10: Experiments H_1 – H_4 in the proof of Theorem 3.3. Experiments H_2 and H_4 include the corresponding boxed statements, but H_1 and H_3 do not.

of blocks in the forgery output by \mathcal{A} . Adversary \mathcal{B} has the same running time as \mathcal{A} and makes at most $(\text{Cost}(\Pi) + 1) \cdot \sigma/2$ queries, where σ is the total number of message blocks in the queries of \mathcal{A} .

Proof. Adversary \mathcal{B} runs \mathcal{A} . For each of \mathcal{A} 's encryption queries, \mathcal{B} runs the encryption scheme of $\Pi[E]$ but with each call to E_K replaced by a query to \mathcal{B} 's first oracle, and returns the ciphertext to \mathcal{A} . When \mathcal{A} outputs a forgery (N, A, C) , adversary \mathcal{B} runs the decryption scheme of $\Pi[E]$ on (N, A, C) , but with each call to E_K/E_K^{-1} replaced by a query to \mathcal{B} 's oracles. Adversary \mathcal{B} returns 1 if \mathcal{A} output a valid forgery, and returns 0 otherwise. Let $\Pi[\pi]$ be the ideal variant of $\Pi[E]$, where calls to E_K/E_K^{-1} are replaced by corresponding queries to π/π^{-1} , with $\pi \leftarrow \text{Perm}(\mathcal{T}, n)$. It suffices to show that $\text{Adv}_{\Pi[\pi]}^{\text{auth}}(\mathcal{A}) \leq 2^{-\tau} + \ell(\text{Cost}(\Pi) + 2)/2^n$.

Consider experiments H_1 – H_3 in Figure 3.11. In H_1 , the adversary has oracle access to the encryption and decryption schemes of $\Pi[\pi]$. Experiment H_2 is identical to H_1 , except that when running the decryption algorithm, we re-sample $\pi \leftarrow \text{Perm}(\mathcal{T}, n)$ before using it in Tag . Experiment H_3 is identical to H_2 , except that instead of using Tag to generate the tag, we sample the tag uniformly.

Let (N, A, C) be the forgery output by \mathcal{A} . Suppose there is no encryption query (N, A, M') with $|M'| = |C| - \tau$. Since decryption of the forgery query involves calling Tag with a tweak that has never been used before, we have $\Pr[\mathcal{A} \text{ forges in } H_1] = \Pr[\mathcal{A} \text{ forges in } H_2]$. Considering lines 12–13 of Auth (and the fact that $\text{Auth}(G_1^-, G_2^-) = \text{true}$) in conjunction with Lemma 3.1 shows that the string $V := \text{Tag}^\pi(T, X)$ is uniform. Thus $\Pr[\mathcal{A} \text{ forges in } H_2] = \Pr[\mathcal{A} \text{ forges in } H_3]$. The probability that \mathcal{A} can forge in H_3 is at most $2^{-\tau}$. Hence $\text{Adv}_{\Pi[\pi]}^{\text{auth}}(\mathcal{A}) \leq 2^{-\tau}$ in this case.

Now, suppose that there is an encryption query (N, A, M') such that $|M'| = |C| - \tau$. (Note that there can be at most one such query, since the attacker is not allowed to re-use a nonce value in two encryption queries.) Let C' be the corresponding ciphertext output by this encryption query, and let $C = C_1 \cdots C_{2m} \parallel \text{tag}$ and $C' = C'_1 \cdots C'_{2m} \parallel \text{tag}'$. If $C_j = C'_j$ for every $j \leq 2m$ then tag and tag' must be different and thus, since Tag is deterministic, the forgery is invalid. Otherwise, take the least index $r \leq m$ such that $C_{2r-1}C_{2r} \neq C'_{2r-1}C'_{2r}$. Consider experiments P_1, \dots, P_{m-r+4} in Figure 3.12. In P_1 , the adversary has two oracles: ENCRYPT and DECRYPT . The first implements the encryption scheme of $\Pi[\pi]$, and the second implements the decryption scheme of $\Pi[\pi]$ but returns false if the decrypted value is \perp and returns true otherwise.

Let S be the subset of $\text{Perm}(\mathcal{T}, n)$ such that for any $f \in S$ and query (T, X) that $\text{ENCRYPT}[\pi](N, A, M')$

```

proc DECRYPT[ $\pi$ ]( $N, A, C$ ) // Experiments  $H_1, \boxed{H_2}$ 
if  $|C| \not\equiv \tau \pmod{2n}$  then return  $\perp$ 
 $C_1 \cdots C_{2m} \parallel \text{tag} := C$  //  $|C_i| = n$  and  $|\text{tag}| = \tau$ 
 $X := 0^{2n}; v := 1$ 
for  $i = 1$  to  $m$  do
   $T := (N, A, v)$ 
   $(Y, M_{2i-1}M_{2i}) := \text{Dec}^{\pi, \pi^{-1}}(T, X, C_{2i-1}C_{2i})$ 
   $v := v + \text{Cost}(\Pi); X := Y$ 
   $\pi \leftarrow_s \text{Perm}(\mathcal{T}, n)$ 
 $T := (N, A, 1 - v); V := \text{Tag}^\pi(T, X)$ 
if  $\text{tag} \neq V[1, \tau]$  then return  $\perp$ 
return  $M_1 \cdots M_{2m}$ 

proc DECRYPT[ $\pi$ ]( $N, A, C$ ) // Experiment  $H_3$ 
if  $|C| \not\equiv \tau \pmod{2n}$  then return  $\perp$ 
 $C_1 \cdots C_{2m} \parallel \text{tag} := C$  //  $|C_i| = n$  and  $|\text{tag}| = \tau$ 
 $X := 0^{2n}; v := 1$ 
for  $i = 1$  to  $m$  do
   $T := (N, A, v)$ 
   $(Y, M_{2i-1}M_{2i}) := \text{Dec}^{\pi, \pi^{-1}}(T, X, C_{2i-1}C_{2i})$ 
   $v := v + \text{Cost}(\Pi); X := Y$ 
 $V \leftarrow_s \{0, 1\}^n$ 
if  $\text{tag} \neq V[1, \tau]$  then return  $\perp$ 
return  $M_1 \cdots M_{2m}$ 

```

Figure 3.11: Experiments H_1 – H_3 in the proof of Theorem 3.4. Experiment H_2 includes the corresponding boxed statement, but experiment H_1 does not. Each experiment also has a procedure $\text{ENCRYPT}[\pi]$, implementing the encryption algorithm of $\Pi[\pi]$, that is not shown for simplicity.

makes to π , we have $f(T, X) = \pi(T, X)$. Experiment P_2 is identical to P_1 , except that in procedure DECRYPT , each time we call Dec or Tag we resample $\pi \leftarrow_s S$. Since in the forgery query we do not repeat the tweak of any encryption query other than (N, A, M') , and π and π^{-1} are called with distinct tweaks, we have $\Pr[\mathcal{A} \text{ forges in } P_1] = \Pr[\mathcal{A} \text{ forges in } P_2]$. In experiment P_3 we sample Y uniformly instead of computing $Y := \text{Dec}^{\pi, \pi^{-1}}(T, X, C_{2r-1}C_{2r})$. Considering lines 14–16 of Auth (and the fact that $\text{Auth}(G_1^-, G_2^-) = \text{true}$) in conjunction with Lemma 3.2, we have $\Pr[\mathcal{A} \text{ forges in } P_2] - \Pr[\mathcal{A} \text{ forges in } P_3] \leq \frac{2\text{Cost}(\Pi)+2}{2^n}$.

For $j = 1, \dots, m - r$, experiment P_{3+j} is identical to P_{2+j} , except that we sample Y uniformly instead of computing $Y := \text{Dec}^{\pi, \pi^{-1}}(T, X, C_{2r+2j-1}C_{2r+2j})$. Considering lines 17–19 of Auth (and the fact that $\text{Auth}(G_1^-, G_2^-) = \text{true}$) in conjunction with Lemma 3.2, we conclude that $\Pr[\mathcal{A} \text{ forges in } P_{2+j}] - \Pr[\mathcal{A} \text{ forges in } P_{3+j}] \leq \frac{2\text{Cost}(\Pi)+2}{2^n}$.

Experiment P_{m-r+4} is identical to P_{m-r+3} except that we sample V uniformly when checking the validity of the forgery instead of computing $V := \text{Tag}^\pi(T, X)$. Let X' be the state used by Tag in $\text{ENCRYPT}[\pi](N, A, M')$. If $X[1, n] \neq X'[1, n]$, which happens with probability at least $1 - 2^{-n}$, then applying Lemma 3.2 to lines 20–22 of procedure Auth , we have $\Pr[\mathcal{A} \text{ forges in } P_{m-r+3}] - \Pr[\mathcal{A} \text{ forges in } P_{m-r+4}] \leq \frac{2}{2^n}$. Finally, $\Pr[\mathcal{A} \text{ forges in } P_{m-r+4}] \leq 2^{-\tau}$. Summing up,

$$\text{Adv}_{\Pi[\pi]}^{\text{auth}}(\mathcal{A}) \leq 2^{-\tau} + \frac{2(m-r+1)(\text{Cost}(\Pi) + 1) + 3}{2^n} \leq 2^{-\tau} + \frac{\ell(\text{Cost}(\Pi) + 2)}{2^n}. \quad \square$$

To summarize, Theorems 3.3 and 3.4 show that if the graphs induced by a given scheme Π satisfy Priv and Auth as defined in Figure 3.7, then Π is a secure AE scheme.

<pre> proc DECRYPT[π](N, A, C) // Experiments $P_1, \boxed{P_2}$ if $C \not\equiv \tau \pmod{2n}$ then return \perp $C_1 \cdots C_{2m} \parallel tag := C$ // $C_i = n$ and $tag = \tau$ $X := 0^{2n}; v := 1$ for $i = 1$ to m do $T := (N, A, v); \boxed{\pi \leftarrow S}$ $(Y, M_{2i-1}M_{2i}) := \text{Dec}^{\pi, \pi^{-1}}(T, X, C_{2i-1}C_{2i})$ $v := v + \text{Cost}(\Pi); X := Y$ $T := (N, A, 1 - v); \boxed{\pi \leftarrow S}; V := \text{Tag}^\pi(T, X)$ if $tag \neq V[1, \tau]$ then return false return true </pre>
<pre> // Experiments P_{3+j}, for $0 \leq j \leq m - r + 1$ proc DECRYPT[π](N, A, C) if $C \not\equiv \tau \pmod{2n}$ then return \perp $C_1 \cdots C_{2m} \parallel tag := C$ // $C_i = n$ and $tag = \tau$ $X := 0^{2n}; v := 1$ for $i = 1$ to m do $T := (N, A, v); \pi \leftarrow S$ $(Y, M_{2i-1}M_{2i}) := \text{Dec}^{\pi, \pi^{-1}}(T, X, C_{2i-1}C_{2i})$ if $r \leq i \leq r + j$ then $Y \leftarrow \{0, 1\}^n$ $v := v + \text{Cost}(\Pi); X := Y$ $\pi \leftarrow S; V := \text{Tag}^\pi(T, X)$ if $j = m - r + 1$ then $V \leftarrow \{0, 1\}^n$ if $tag \neq V[1, \tau]$ then return false return true </pre>

Figure 3.12: Experiments P_1, \dots, P_{4+m-r} in the proof of Theorem 3.4. Experiment P_2 includes the corresponding boxed statement, but P_1 does not. Each experiment also has a procedure $\text{ENCRYPT}[\pi]$, implementing the encryption algorithm of $\Pi[\pi]$, that is not shown for simplicity. Here S is the set of $f \in \text{Perm}(\mathcal{T}, n)$ such that for any query (T, X) that $\text{ENCRYPT}[\pi](N, A, M')$ makes to π , it holds that $f(T, X) = \pi(T, X)$.

4 Implementation and Results

We have implemented the Priv and Auth algorithms described in Section 3, and used them to synthesize AE schemes. The code is written in OCaml and available at <https://github.com/amaloz/ae-generator>.¹⁰ Our system has two modules: an *analysis module* that, given graphs corresponding to an AE scheme, verifies whether the scheme is secure, and a *synthesis module* that synthesizes AE schemes by enumerating candidate AE schemes and using the analysis module to see if they are secure. We describe these components below, where throughout this section, the term *graph* denotes an unlabeled graph.

Analyzer. The analysis module takes as input a representation (in a stack-based language) of the Dec and Tag graphs; the stack-based language makes it easy to both convert the inputs into their respective graphs as well as to synthesize schemes. We first derive a graph for the Enc algorithm given the graph for the Dec algorithm, as described below. Given graphs for the Enc, Dec, and Tag algorithms, we can then run the privacy and authenticity checks described in Figure 3.7 to check security of the scheme. Our analyzer is able to verify simplified variants of OCB [22], XCBC [12], COPA [3], OTR [20], and CCM [10], among others.

Deriving the Enc graph. We implement an algorithm Reverse that, given a Dec graph, computes a corresponding Enc graph if one exists. The basic idea is to swap the IN and OUT nodes of the input graph (recall that IN and OUT nodes in the Dec graph denote ciphertext blocks and plaintext blocks, respectively,

¹⁰All results in this section were computed using the code from commit [515959bcf9fa805cbc102aa8a2772cf3c35906f7](https://github.com/amaloz/ae-generator/commit/515959bcf9fa805cbc102aa8a2772cf3c35906f7).

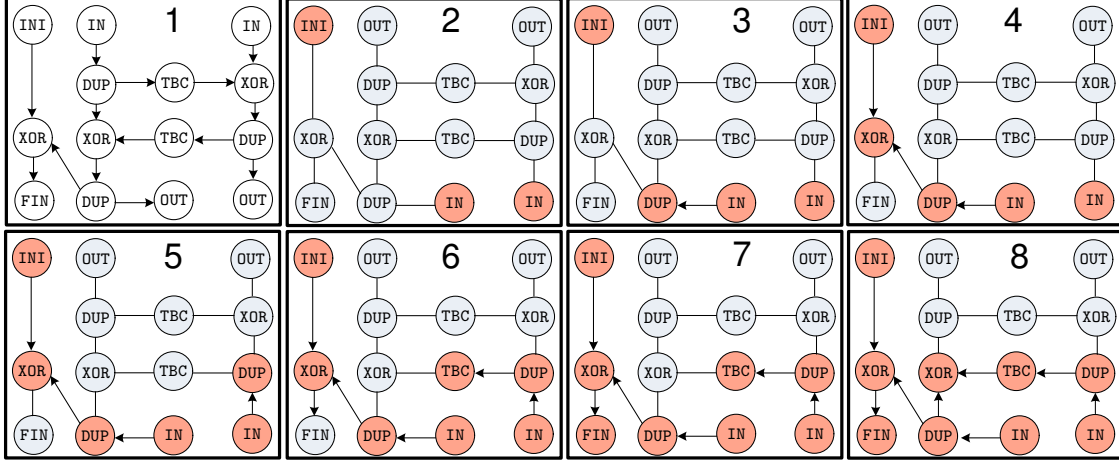


Figure 4.1: Illustration of the first few steps of running `Reverse` on the Dec graph of OTR. The first picture is the Dec graph of OTR.

whereas IN and OUT nodes in the Enc graph are flipped), and then selectively reverse the edges to ensure that each node has correct ingoing/outgoing degrees. Deriving the Enc graph is thus simple if there is at most one path from an IN node to an OUT node and these paths do not cross, as in the case of OCB. However, in other schemes, such as OTR, each IN node may have multiple paths to each OUT node. We handle this as described next.

On input G_1^- , let \mathcal{G}_1 be the *undirected* graph of G_1^- . In \mathcal{G}_1 , rename IN nodes as OUT nodes, and OUT nodes as IN nodes; let \mathcal{G}_2 be the resulting graph. `Reverse` then assigns direction to the edges of \mathcal{G}_2 such that each node has correct ingoing/outgoing degrees; the resulting graph G_2^- is output. (If no assignment is possible, then the output is \perp .)

To implement this idea efficiently, we color each node either “red” or “blue”, where red nodes denote nodes that have already been processed, and blue nodes denote unprocessed nodes. Starting from \mathcal{G}_2 , we initially color IN and INI nodes red and all other nodes blue. We repeatedly iterate over the blue nodes until we reach a fixed point, where in each iteration we assign direction to some edges and re-color some nodes red. If a fixed point is reached before all nodes have been colored red, we return \perp ; otherwise, we return \mathcal{G}_2 , which represents the reversed graph. If the graph \mathcal{G}_2 has r nodes then we have at most r iterations with each iteration taking $O(r)$ time.

In each iteration, we process each blue node x as follows. Let $\text{ord}(x) = 2$ if x is an XOR node, and let $\text{ord}(x) = 1$ otherwise. If there are *exactly* $\text{ord}(x)$ red neighbors of x then (1) for each such neighbor y , assign the direction $y \rightarrow x$, and (2) color x red. Note that in each step we ensure that the current node x has the correct ingoing degree if we color it red. We never assign an ingoing edge to x in any other step. Hence when there are no blue nodes, each node in the directed graph has the correct ingoing/outgoing degrees. See Figure 4.1 for an illustration of `Reverse` on OTR.

We prove in Appendix A that `Reverse` is sound; namely, that if running `Reverse` on a Dec graph produces an Enc graph, then Dec is a correct decryption algorithm for Enc.

As a side note, the `Reverse` algorithm allows us to easily check if a scheme is *inverse-free* (i.e., the scheme only uses the forward direction of the TBC), which is important when constructing hardware realizations of AE schemes due to the potential savings in chip space, among other benefits [14, 20]. After running `Reverse`, we can check if the parent nodes for all the TBC nodes in the Enc and Dec graph are the same; if so, the scheme is inverse-free.

Synthesizer. We synthesize schemes as follows. Fixing a Tag graph, we enumerate all possible Dec graphs of a given size, pruning out “uninteresting” schemes such as ones with two (or more) TBC nodes chained together, and feed each pair of (Dec, Tag) graphs to our analysis module. To generate the Dec graph, we

#Nodes	Unique	“Optimal”	WP	SP	Time
12	13 (0)	13	7	5	47 sec
13	142 (0)	0	0	0	4.3 min
14	582 (2)	171	48 (4)	5	24.2 min
15	2826 (54)	40	18	6	2.8 hours
16	3090 (—)	66	25 (4)	1	3 hours*
Total	6653	290	98 (8)	17	

Figure 4.2: Synthesis results. The first column shows the number of instructions in the Dec graph of the given scheme; the second column the number of secure (and unique) schemes, with the number in parentheses denoting the number of schemes in which the security check fails but we cannot automatically find a concrete attack; the third column the number of (secure) schemes that are “optimal”, i.e., having two TBC nodes per Dec graph; the fourth column the number of “optimal” weakly parallelizable schemes, with the number in parentheses denoting the weakly parallelizable schemes which only use the forward direction of the TBC; the fifth column the number of “optimal” *strongly* parallelizable schemes; and the final column the total synthesis time, where an asterisk indicates that we halted execution after the given time.

start from a graph containing just the IN and INI nodes, and add nodes and their corresponding edges until the given size bound is reached. If the resulting graph is “well-formed” (i.e., there are no “dangling” edges and no loops), we derive the corresponding Enc graph as discussed above and run the analysis module on the result. Unfortunately, this approach is prohibitively expensive as described, especially as the size bound increases. Thus, we use several optimizations to speed up the process.

Firstly, instead of synthesizing graphs with FIN and OUT nodes, we replace these with “terminal” nodes. Upon deriving a well-formed graph, we replace the “terminal” nodes with all possible permutations of FIN and OUT nodes and check security of each. Thus we no longer need to explore the search space for each FIN and OUT node; instead, we explore the search space *once* using a “terminal” node, and later replace the “terminal” node with all possible combinations of FIN and OUT nodes. Likewise, we can apply this same idea to INI and IN nodes by introducing a “start” node.

Secondly, we observe that AE schemes like OCB, COPA, and OTR do not utilize one of the INI nodes in the sense that they simply output the input value directly. Thus, we can remove two nodes from the synthesis by only synthesizing schemes containing one INI and FIN node. The drawback of this optimization is that it misses schemes such as XCBC and CCM which do in fact use both INI nodes; however, it greatly speeds up synthesis. All the results that follow use this optimization. (It would, of course, be possible to synthesize schemes without using this optimization.)

Results. Using the optimizations described above, we ran our synthesizer to find AE schemes with Dec and Enc graphs of sizes between twelve and sixteen (we found no AE schemes with size less than twelve). Note that our synthesizer does not remove duplicate schemes. In addition, there are many “equivalent” schemes in the sense that one is the same as another except with the outputs and/or inputs flipped. We thus developed a heuristic to remove duplicate and “equivalent” schemes as follows. Let $F(\cdot, \cdot, \cdot)$ be the encrypt operation of a given scheme, where the first argument is the INI input (recall we consider the simplified variant where we only use one INI node) and the other arguments are the IN inputs. Choosing arbitrary but fixed inputs X , M_1 , and M_2 , we compute $Y\|C_1\|C_2 := F(X, M_1, M_2)$ and $Y'\|C'_1\|C'_2 := F(X, M_2, M_1)$. We maintain a table of existing ciphertexts; if any of YC_1C_2 , YC_2C_1 , $Y'C'_1C'_2$, or $Y'C'_2C'_1$ exists in the table, we discard the scheme as a “duplicate”; otherwise, we add each of these to the table and continue.

Figure 4.2 shows the results. The experiments were run on a commodity laptop; because of the long running time for synthesizing schemes of size sixteen, we stopped the synthesis after three hours for this size. Due to the large number of discovered schemes, we developed two algorithms to prune the result space. The first simply filters out all schemes Π such that $\mathbf{Cost}(\Pi) > c$ for some integer c . In Figure 4.2 we set $c = 2$, thus pruning out all non-“rate-1” schemes; this removes 95% of the found schemes.

Our second algorithm checks whether a scheme is *parallelizable*, an important criterion for AE schemes.




Figure 4.3: Three of our synthesized schemes of size twelve, illustrated for a four-block message M_1, \dots, M_4 . In the second scheme, Σ is the checksum of the even blocks, i.e., $\Sigma = M_2 \oplus M_4$.

Scheme	Enc (cycles/byte)	Dec (cycles/byte)
OCB	0.7122 ± 0.0072	0.7650 ± 0.0025
1	0.7253 ± 0.0055	0.7485 ± 0.0047
2	0.7116 ± 0.0025	0.7643 ± 0.0023
3	0.8139 ± 0.0121	2.7566 ± 0.0010

Figure 4.4: Performance results of OCB and the three synthesized schemes in Figure 4.3 (Scheme 1 denotes the left scheme, Scheme 2 the middle scheme, and Scheme 3 the right scheme). We report the time for encryption and decryption when processing a 4096-bit message with empty associated data, along with the 95% confidence intervals over 100 runs of each scheme. The experiments were run on a 4-core 2.90 GHz Intel Core i5-4210H CPU with TurboBoost disabled.

Note that we can view the encryption of a message $M = M_1 \cdots M_{2m}$ as a single graph constructed from m Enc graphs G_1, \dots, G_m , where the FIN nodes of G_i coincide with the INI nodes of G_{i+1} . We can then assign a “depth” to each node in this graph as follows:

- The INI nodes in G_1 and the IN nodes in $\{G_i\}$ get a depth of 0.
- For each node x , let t be the maximum depth of x ’s parent(s). If x is a TBC node then $\text{depth}(x) = t + 1$; otherwise $\text{depth}(x) = t$.

(Intuitively, $\text{depth}(x)$ represents the *latency*, in terms of the number of TBC calls, of computing the value at node x .) We can use the same idea to compute a “depth” for decryption.

We now define our notion of parallelizability. We call an AE scheme Π *weakly parallelizable* if for any integer m and any node x in the graph described above we have $\text{depth}(x) \leq \mathbf{Cost}(\Pi)$ for both encryption and decryption. A scheme is *strongly parallelizable* if $\text{depth}(x) \leq 1$. Intuitively, weakly parallelizable schemes are ones where the TBC calls can be parallelized *across* two-block chunks (but not necessarily within the processing of a two-block chunk), and strongly parallelizable schemes are ones where the TBC calls can be parallelized even *within* a two-block chunk. As an example, OTR (cf. Figure 3.9) is weakly parallelizable while OCB (cf. Figure 3.2) is strongly parallelizable.

We can check these conditions efficiently by noting first that we only need to look at the OUT and FIN nodes, since depth is strictly increasing. Now, suppose we run the analysis on graph G_i . If the depth of the FIN nodes is zero, then it suffices to compute $t = \max\{\text{depth}(\text{OUT}_1), \text{depth}(\text{OUT}_2)\}$ (where OUT_1 and OUT_2 are the two OUT nodes) and check whether $t \leq \mathbf{Cost}(\Pi)$ or $t \leq 1$. If the FIN nodes have depth greater than zero (say, c), we need to rerun the analysis, this time setting the depths of the INI nodes to c (rather than zero). We can then compute $t' = \max\{\text{depth}(\text{OUT}_1), \text{depth}(\text{OUT}_2)\}$. If $t' \neq t$ then this implies that depth grows with m (and thus the scheme is not parallelizable); otherwise we can check whether $t' \leq \mathbf{Cost}(\Pi)$ or $t' \leq 1$ to determine whether the scheme is parallelizable.

Looking at the results of Figure 4.2, we found thirteen secure AE schemes of size twelve, five of which are strongly parallelizable. Of these schemes, as far as we know, only OCB exists in the literature. In Figure 4.3 we show two of these newly synthesized schemes, along with one which is *not* parallelizable. (For all the schemes in the figure, encryption is strongly parallelizable; for the third scheme, however, decryption cannot be parallelized.) We implemented all three schemes and compared their performance with that of

an optimized implementation of OCB by Krovetz¹¹ using AES-NI; see Figure 4.4. (Note that the results in Figure 4.4 are preliminary timing numbers; the purpose of these experiments is to show that our schemes are competitive with, not necessarily better than, OCB.) We find that the encryption procedure for all four schemes is comparable. However, the decryption procedure of the third synthesized scheme is noticeably slower than the others. This is because decryption for this scheme is not parallelizable; namely, to decrypt ciphertext block C_i we need plaintext block M_{i-1} .

In addition, among the weakly parallelizable schemes, we found eight schemes which are inverse-free (we found no such schemes for strongly parallelizable schemes). The schemes of size fourteen that we found use one fewer XOR instruction than OTR, the fastest known inverse-free AE scheme we are aware of.

We also ran our attack generation algorithm (cf. Appendix C) over schemes of size 12–15 and found that the number of schemes where no attack could be found closely matched the number of schemes our analysis found secure, thus pointing to the fact that while our analysis is not sound, it appears to capture most secure schemes. In particular, we found a total of 56 schemes which we could neither prove secure nor find a concrete attack for sizes between 12 and 15.

We remark that our tool currently takes a given bound S and enumerates all schemes in which decryption can be implemented using at most S instructions. In future work one could consider assigning a cost to different instructions (e.g., letting DUP have cost 0, and letting TBC have cost some fixed multiple of XOR) and enumerating all schemes having at most some given cost.

5 Conclusion

In this work, we present a methodology for automatically proving the security of a large class of authenticated encryption (AE) schemes. Using our approach, we are able to synthesize thousands of schemes, most of which have never been studied in the literature. Among these, we discovered five new schemes which are as “compact” (in terms of the number of instructions per message block), as “efficient” (in terms of the number of blockcipher calls per message block), and as parallelizable as OCB, with competitive performance.

There are several interesting avenues for future work. Further optimizing the synthesis procedure would allow us to generate more schemes. Some of these schemes may have additional properties of interest, such as misuse-resistance [11]; developing techniques for automatically checking schemes for these additional properties would be very useful. Taking a different approach, it would be interesting to see if similar techniques can be applied to more general classes of AE schemes.

Acknowledgments

We thank Samuel Neves for pointing out that on Haswell CPUs, TurboBoost must be disabled to measure the timing accurately. We also thank the anonymous reviewers of CCS 2015 for their suggestion of automatic generation of attacks on schemes that fail our tests.

Work of Jonathan Katz was done for Exelis under contract number N00173-11-C-2045 to NRL. Work of Alex J. Malozemoff was conducted with Government support awarded by DoD, Air Force Office of Scientific Research, National Defense Science and Engineering Graduate (NDSEG) Fellowship, 32 CFR 168a.

References

- [1] Joseph A. Akinyele, Matthew Green, and Susan Hohenberger. Using SMT solvers to automate design tasks for encryption and signature schemes. In *20th ACM Conference on Computer and Communications Security (CCS 2013)*, pp. 399–410. ACM Press, 2013.
- [2] Joseph A. Akinyele, Matthew Green, Susan Hohenberger, and Matthew W. Pagano. Machine-generated algorithms, proofs and software for the batch verification of digital signature schemes. In *19th ACM*

¹¹See <http://web.cs.ucdavis.edu/~rogaway/ocb/news>.

- Conference on Computer and Communications Security (CCS 2012)*, pp. 474–487. ACM Press, 2012. Full version available at <https://eprint.iacr.org/2013/175>.
- [3] Elena Andreeva, Andrey Bogdanov, Atul Luykx, Bart Mennink, Elmar Tischhauser, and Kan Yasuda. Parallelizable and authenticated online ciphers. In *Advances in Cryptology—Asiacrypt 2013, Part I*, volume 8269 of *LNCS*, pp. 424–443. Springer, 2013. Full version available at <https://eprint.iacr.org/2013/790>.
 - [4] Gilles Barthe, Juan Manuel Crespo, Benjamin Grégoire, César Kunz, Yassine Lakhnech, Benedikt Schmidt, and Santiago Zanella Béguelin. Fully automated analysis of padding-based encryption in the computational model. In *20th ACM Conference on Computer and Communications Security (CCS 2013)*, pp. 1247–1260. ACM Press, 2013. Full version available at <https://eprint.iacr.org/2012/695>.
 - [5] Gilles Barthe, Edvard Fagerholm, Dario Fiore, John C. Mitchell, Andre Scedrov, and Benedikt Schmidt. Automated analysis of cryptographic assumptions in generic group models. In *Advances in Cryptology—Crypto 2014, Part I*, volume 8616 of *LNCS*, pp. 95–112. Springer, 2014. Full version available at <https://eprint.iacr.org/2014/458>.
 - [6] Gilles Barthe, Edvard Fagerholm, Dario Fiore, Andre Scedrov, Benedikt Schmidt, and Mehdi Tibouchi. Strongly-optimal structure preserving signatures from type II pairings: Synthesis and lower bounds. In *18th International Conference on Theory and Practice of Public Key Cryptography (PKC 2015)*, volume 9020 of *LNCS*, pp. 355–376. Springer, 2015. Full version available at <https://eprint.iacr.org/2015/019>.
 - [7] Mihir Bellare and Chanathip Namprempre. Authenticated encryption: Relations among notions and analysis of the generic composition paradigm. *Journal of Cryptology*, 21(4):469–491, October 2008.
 - [8] Mihir Bellare and Phillip Rogaway. The security of triple encryption and a framework for code-based game-playing proofs. In *Advances in Cryptology—Eurocrypt 2006*, volume 4004 of *LNCS*, pp. 409–426. Springer 2006. Full version available at <https://eprint.iacr.org/2004/331>.
 - [9] Dan Bernstein. Cryptographic competitions: CAESAR call for submissions, final (2014.01.27). <http://competitions.cr.yt.to/caesar-call.html>.
 - [10] Morris Dworkin. Recommendations for block cipher modes of operation: The CCM mode for authentication and confidentiality. NIST Special Publication 800-38C, July 2007.
 - [11] Ewan Fleischmann, Christian Forler, and Stefan Lucks. McOE: A family of almost foolproof on-line authenticated encryption schemes. In *Fast Software Encryption (FSE) 2012*, volume 7549 of *LNCS*, pp. 196–215. Springer, 2012. Full version available at <https://eprint.iacr.org/2011/644>.
 - [12] Virgil D. Gligor and Pompiliu Donescu. Fast encryption and authentication: XCBC encryption and XECB authentication modes. In *Fast Software Encryption (FSE) 2001*, volume 2355 of *LNCS*, pp. 92–108. Springer, 2002.
 - [13] Tetsu Iwata, Keisuke Ohashi, and Kazuhiko Minematsu. Breaking and repairing GCM security proofs. In *Advances in Cryptology—Crypto 2012*, volume 7417 of *LNCS*, pp. 31–49. Springer, 2012. Full version available at <https://eprint.iacr.org/2012/438>.
 - [14] Tetsu Iwata and Kan Yasuda. BTM: A single-key, inverse-cipher-free mode for deterministic authenticated encryption. In *SAC 2009: 16th Annual International Workshop on Selected Areas in Cryptography*, volume 5867 of *LNCS*, pp. 313–330. Springer, 2009.
 - [15] Jonathan Katz and Moti Yung. Unforgeable encryption and chosen ciphertext secure modes of operation. In *Fast Software Encryption (FSE) 2000*, volume 1978 of *LNCS*, pp. 284–299. Springer, 2001.

- [16] Ted Krovetz and Phillip Rogaway. The software performance of authenticated-encryption modes. In *Fast Software Encryption (FSE) 2011*, volume 6733 of *LNCS*, pp 306–327. Springer, 2011. Full version available at <http://web.cs.ucdavis.edu/~rogaway/papers/ae.pdf>.
- [17] Moses Liskov, Ronald L. Rivest, and David Wagner. Tweakable block ciphers. In *Advances in Cryptology—Crypto 2002*, volume 2442 of *LNCS*, pp. 31–46. Springer, 2002.
- [18] Alex J. Malozemoff, Jonathan Katz, and Matthew D. Green. Automated analysis and synthesis of block-cipher modes of operation. In *IEEE CSF 2014*, pp. 140–152, 2014.
- [19] David A. McGrew and John Viega. The security and performance of the Galois/counter mode (GCM) of operation. In *Progress in Cryptology—Indocrypt 2004*, volume 3348 of *LNCS*, pp. 343–355. Springer, 2014. Full version available at <https://eprint.iacr.org/2004/193>.
- [20] Kazuhiko Minematsu. Parallelizable rate-1 authenticated encryption from pseudorandom functions. In *Advances in Cryptology—Eurocrypt 2014*, volume 8441 of *LNCS*, pp. 275–292. Springer, 2014. Full version available at <https://eprint.iacr.org/2013/628>.
- [21] Phillip Rogaway. Authenticated-encryption with associated-data. In *9th ACM Conference on Computer and Communications Security (CCS 2002)*, pp. 98–107. ACM Press, 2002. Full version available at <http://web.cs.ucdavis.edu/~rogaway/papers/ad.pdf>.
- [22] Phillip Rogaway. Efficient instantiations of tweakable blockciphers and refinements to modes OCB and PMAC. In *Advances in Cryptology—Asiacrypt 2004*, volume 3329 of *LNCS*, pp. 16–31. Springer 2004. Full version available at <http://web.cs.ucdavis.edu/~rogaway/papers/offsets.pdf>.
- [23] Phillip Rogaway. Nonce-based symmetric encryption. In *Fast Software Encryption (FSE) 2004*, volume 3017 of *LNCS*, pp. 348–359, Springer 2004. Full version available at <http://web.cs.ucdavis.edu/~rogaway/papers/nonce.pdf>.
- [24] Phillip Rogaway, Mihir Bellare, John Black, and Ted Krovetz. OCB: A block-cipher mode of operation for efficient authenticated encryption. In *8th ACM Conference on Computer and Communications Security (CCS)*, pp. 196–205. ACM Press, 2001. Full version available at <http://web.cs.ucdavis.edu/~rogaway/papers/ocb-full.pdf>.
- [25] Ashish Tiwari, Adrià Gascón, and Bruno Dutertre. Program synthesis using dual interpretation. In *CADE 2015*, 2015.

A Correctness of Reverse

Recall that `Dec` is supposed to reverse `Enc`: namely, for any $f \in \text{Perm}(\mathcal{T}, n)$, every $T \in \mathcal{T}$, every initial state X , and any input M , if $(Y, C) = \text{Enc}^f(T, X, M)$ then $(Y, M) = \text{Dec}^{f, f^{-1}}(T, X, C)$. We call this the *reversal condition*. We now prove that `Reverse` produces a “correct” `Enc` graph given a `Dec` graph as input, by showing that the input graph G_1^- is a *reversal* of the output graph G_2^- . That is, for any labeling of G_2^- such that the labels on TBC nodes are distinct positive integers, one can assign a corresponding labeling to G_1^- such that the resulting algorithms `Enc` and `Dec`—that are derived from the corresponding labeled graphs G_2 and G_1 , respectively—satisfy the reversal condition.

Lemma A.1. *If `Reverse`(G_1^-) outputs $G_2^- \neq \perp$, then G_1^- is a reversal of G_2^- .*

Proof. We first describe how to label G_1^- given the labeling for G_2^- . For convenience, we use the same symbol to refer to corresponding nodes in the two graphs. Recall that if a node x is a TBC node in G_2^- then x is also a TBC node in G_1^- . Let ℓ be the label of x in G_2^- and let y be the parent of x in G_2^- . If y is also the parent of x in G_1^- , then assign label ℓ to x in G_1^- ; otherwise assign x label $-\ell$.

We prove our result for general graphs, meaning that the number of INI, FIN, IN, and OUT nodes can be arbitrary. Our proof is by induction on the total number s of TBC, XOR, and DUP nodes in G_1^- . The base case $s = 0$ is trivial.

Let $G_2^- \neq \perp$ denote the output of $\text{Reverse}(G_1^-)$. Fix some $f \in \text{Perm}(\mathcal{T}, n)$, and a labeling on G_2^- such that the labels on TBC nodes are distinct, positive integers. Fix a tweak $T \in \mathcal{T}$, an initial state X and input M for the resulting labeled graph G_2 . Evaluate G_2 on input (T, X, M) using f , and let $\text{val}(x)$ be the value on node x . Let (Y, C) be the result of this evaluation. Among the DUP, TBC, and XOR nodes of G_1^- , let x_0 denote the first node that $\text{Reverse}(G_1^-)$ re-colors red. We consider the case that x_0 is an XOR node; the other cases are similar.

Let x_1 and x_2 be the red neighbors of x_0 , and let x_3 be the other (blue) neighbor of x_0 . Note that x_1 can't be an IN node in G_1^- (and thus an OUT node in G_2^-), since such a node will be re-colored red only after its unique neighbor turns red, but here x_1 is already red while x_0 is still blue. Likewise, x_2 can't be an IN node in G_1^- . Then x_1 and x_2 must be INI or OUT nodes in G_1^- , because x_0 is the first among the DUP, XOR, and TBC nodes of G_1^- that is re-colored red.

Let H_1^- be the graph obtained from G_1^- by merging x_0 , x_1 , and x_2 into a single node y . If both x_1 and x_2 are INI nodes in G_1^- then let y be an INI node; otherwise, let y be an OUT node. Each node in H_1^- has correct ingoing/outgoing degrees and there are $s - 1$ TBC, XOR, and DUP nodes in H_1^- . Note that G_2^- can be obtained by (i) running $\text{Reverse}(H_1^-)$ to get a graph H_2^- , (ii) splitting the node y in H_2^- into the nodes x_0, x_1, x_2 as in G_1^- , and (iii) changing OUT nodes to IN nodes among the reinstated nodes, and reversing the direction of their attached edges.

By the induction hypothesis, H_1^- is a reversal of H_2^- . Use the same labeling of G_1^- and G_2^- for H_1^- and H_2^- , respectively, and let H_1 and H_2 be the corresponding labeled graphs. In H_2 , assign the value $\text{val}(z)$ to each INI/IN node z , with $\text{val}(y) := \text{val}(x_1) \oplus \text{val}(x_2)$, which is also $\text{val}(x_3)$. If we evaluate H_2 on tweak T and the assigned input using f , then the value at each node x is also $\text{val}(x)$. Since H_1^- is a reversal of H_2^- , if we assign value $\text{val}(x)$ on each INI/IN node x of H_1 , and evaluate H_1 on tweak T and the assigned input using f and f^{-1} , then the value of each FIN/OUT node x is also $\text{val}(x)$. Next, evaluate G_1 on (T, X, C) using f and f^{-1} . We consider the following cases.

Case 1: Both x_1 and x_2 are INI nodes in G_1 . Then we are repeating the evaluation of H_1 , with the three nodes x_0 , x_1 , and x_2 simulating node y , and assigning value $\text{val}(x_1) \oplus \text{val}(x_2) = \text{val}(y)$ to x_3 . Thus, in the evaluation of G_1 , the value at each FIN/OUT node x is also $\text{val}(x)$, leading to the result (Y, M) .

Case 2: One of x_1 or x_2 is an OUT node in G_1 . Without loss of generality, suppose that x_1 is an OUT node. Hence in G_1 , the XOR node x_0 is the parent of x_1 , and so x_2 must be an INI node. (Recall that x_2 must be either INI or OUT, but if it is an OUT node then the XOR node x_0 has outgoing degree at least two, a contradiction.) Then we are repeating the evaluation of H_1 to get $\text{val}(y) = \text{val}(x_1) \oplus \text{val}(x_2)$ on node y , inserting node x_2 with value $\text{val}(x_2)$, and computing $\text{val}(x_2) \oplus \text{val}(y) = \text{val}(x_1) = \text{val}(x_0)$ at nodes x_1 and x_0 . Thus, for the evaluation on G_1 , the value at each FIN/OUT node x is also $\text{val}(x)$, leading to result (Y, M) . \square

B Arbitrary Message Lengths

Our template in Section 3.1 only handles messages in $(\{0, 1\}^{2n})^*$, yet in practice schemes should support messages of arbitrary length. While arbitrary length messages can be handled by reversibly padding out to a multiple of $2n$ bits, we show here how to extend our template to handle messages in $\{0, 1\}^*$ more efficiently.

The key for an AE scheme will now be a pair (K, Δ) . The first component K will be used as the key for a tweakable blockcipher as before, while Δ is an n -bit string used to handle the final (fragmentary) block. In Figure B.1, we show algorithms $\sharp\text{Enc}$ and $\sharp\text{Dec}$ that operate on strings of length less than $2n$ bits (see Figure B.2 for a graphical depiction); they are based on how OTR [20] handles fragmentary blocks. Figure B.3 shows how to use $\sharp\text{Enc}$ and $\sharp\text{Dec}$ (along with Enc , Dec , and Tag) to construct encryption and decryption algorithms for arbitrary length messages. Note that the code in Figure 3.1 is a special case of

```

 $\#Enc_{\Delta}^{EK}(T, X, M) \quad // |M| < 2n$ 
 $X_1X_2 := X; Y_2 := X_2 \quad // |X_i| = n$ 
 $(N, A, i) := T; L := (N, A, i + 1)$ 
if  $M = \varepsilon$  then  $C := \varepsilon; Y_1 := X_1$ 
elseif  $|M| > n$  then
   $M_1M_2 := M \quad // |M_1| = n$ 
   $V := E_K^T(M_1); C_2 := V[1, |M_2|] \oplus M_2; Y_1 := X_1 \oplus V \oplus \Delta$ 
   $C_1 := E_K^L(C_210^*) \oplus M_1; C := C_1C_2$ 
elseif  $|M| = n$  then  $C := E_K^T(0^n) \oplus M; Y_1 := X_1 \oplus M$ 
else  $C := E_K^T(0^n)[1, |M|] \oplus M; Y_1 := X_1 \oplus M10^* \oplus \Delta$ 
return  $(Y_1Y_2, C)$ 

 $\#Dec_{\Delta}^{EK}(T, X, C) \quad // |C| < 2n$ 
 $X_1X_2 := X; Y_2 := X_2 \quad // |X_i| = n$ 
 $(N, A, i) := T; L := (N, A, i + 1)$ 
if  $C = \varepsilon$  then  $M := \varepsilon; Y_1 := X_1$ 
elseif  $|C| > n$  then
   $C_1C_2 := C \quad // |C_1| = n$ 
   $M_1 := C_1 \oplus E_K^L(C_210^*); V := E_K^T(M_1)$ 
   $M_2 := C_2 \oplus V[1, |C_2|]; M := M_1M_2; Y_1 := X_1 \oplus V \oplus \Delta$ 
elseif  $|C| = n$  then  $M := E_K^T(0^n) \oplus C; Y_1 := X_1 \oplus M$ 
else  $M := E_K^T(0^n)[1, |C|] \oplus C; Y_1 := X_1 \oplus M10^* \oplus \Delta$ 
return  $(Y_1Y_2, M)$ 

```

Figure B.1: Algorithms $\#Enc$ and $\#Dec$ for handling a final (fragmentary) message block. For a string X with $|X| < n$, we let $X10^*$ denote $X10^{n-|X|-1}$.

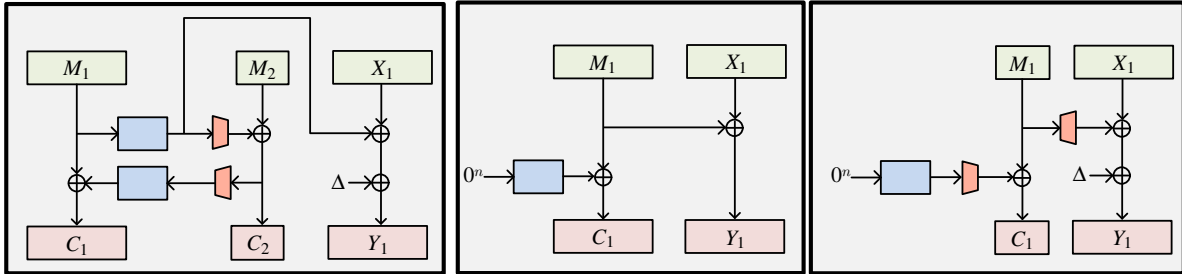


Figure B.2: Illustration of algorithm $\#Enc$ when $|M| > n$ (left), $|M| = n$ (middle), or $0 < |M| < n$ (right). The incoming state $X = X_1X_2$ and outgoing state $Y = Y_1Y_2$, with $|X_i| = |Y_i| = n$, will satisfy $Y_2 = X_2$, so we show only how to process X_1 to get Y_1 . The trapezoids represent truncation and padding with 10^* .

that in Figure B.3 when the message length is a multiple of $2n$ bits. Below, we show that the extended template indeed leads to secure schemes.

Theorem B.1. Let $\Pi[E] = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ be an AE scheme on $\{0, 1\}^*$ for which $\text{Priv}(G_1^-, G_2^-) = \text{true}$, where G_1^-, G_2^- are the unlabeled graphs for Enc and Tag of Π , respectively. Then for any adversary \mathcal{A} , there is an adversary \mathcal{B} with $\text{Adv}_{\Pi[E]}^{\text{priv}}(\mathcal{A}) \leq \text{Adv}_E^{\text{pp}}(\mathcal{B})$. Adversary \mathcal{B} has the same running time as \mathcal{A} and makes at most $(\text{Cost}(\Pi) + 1) \cdot \sigma/2$ queries, where σ is the number of (full) message blocks in the queries of \mathcal{A} .

Proof. Adversary \mathcal{B} runs \mathcal{A} as follows. For each of \mathcal{A} 's queries (N, A, M) , adversary \mathcal{B} runs the encryption scheme $\Pi[E]$ on (N, A, M) with each call to E_K replaced by a query to \mathcal{B} 's oracle, and returns the ciphertext to \mathcal{A} . Finally, \mathcal{B} outputs the same guess as \mathcal{A} . To compute the the number of queries of \mathcal{B} , note that

```

 $\mathcal{E}_{K,\Delta}(N, A, M)$ 
 $X := 0^{2n}; v := 1; M_1 \cdots M_{2m} M_{2m+1} := M$ 
//  $|M_i| = n$  for  $i \leq 2m$ , and  $|M_{2m+1}| < 2n$ 
for  $i = 1$  to  $m$  do
   $T := (N, A, v); (Y, C_{2i-1}, C_{2i}) := \text{Enc}^{E_K}(T, X, M_{2i-1}, M_{2i})$ 
   $v := v + \text{Cost}(\Pi); X := Y$ 
 $(X, C_{2m+1}) := \# \text{Enc}_{\Delta}^{E_K}((N, A, v), X, M_{2m+1})$ 
 $v := v + \lceil |M_{2m+1}|/n \rceil - 1$   $T := (N, A, -v); V := \text{Tag}^{E_K}(T, X)$ 
return  $C_1 \cdots C_{2m} \parallel V[1, \tau]$ 

 $\mathcal{D}_{K,\Delta}(N, A, C)$ 
if  $|C| \not\equiv \tau \pmod{2n}$  then return  $\perp$ 
 $C_1 \cdots C_{2m} C_{2m+1} \parallel \text{tag} := C; X := 0^{2n}; v := 1$ 
//  $|C_i| = n$  for  $i \leq 2m$ ,  $|C_{2m+1}| < 2n$ , and  $|\text{tag}| = \tau$ 
for  $i = 1$  to  $m$  do
   $T := (N, A, v); (Y, M_{2i-1}, M_{2i}) := \text{Dec}^{E_K, E_K^{-1}}(T, X, C_{2i-1}, C_{2i})$ 
   $v := v + \text{Cost}(\Pi); X := Y$ 
 $(X, M_{2m+1}) := \# \text{Dec}_{\Delta}^{E_K}((N, A, v), X, C_{2m+1})$ 
 $v := v + \lceil |C_{2m+1}|/n \rceil$   $T := (N, A, -v); V := \text{Tag}^{E_K}(T, X)$ 
if  $\text{tag} \neq V[1, \tau]$  then return  $\perp$ 
else return  $M_1 \cdots M_{2m}$ 

```

Figure B.3: Code of an AE scheme $\Pi = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ handling arbitrary length messages, based on a tweakable blockcipher E and a triple of deterministic algorithms $(\text{Enc}, \text{Dec}, \text{Tag})$.

$\text{Cost}(\Pi) \geq 2$ and $\# \text{Enc}/\# \text{Dec}$, on message M , makes $\lceil M/n \rceil \leq \lceil M/n \rceil \cdot \text{Cost}(\Pi)/2$ calls to the tweakable blockcipher or its inverse.

Let $\Pi[\pi]$ be the ideal variant of $\Pi[E]$, where E_K calls are replaced by corresponding queries to π , with $\pi \leftarrow \$_s \text{Perm}(\mathcal{T}, n)$. It suffices to show that $\text{Adv}_{\Pi[\pi]}^{\text{priv}}(\mathcal{A}) = 0$. Consider experiments H_1 – H_5 in Figure B.4. The adversary has oracle access to the encryption scheme of $\Pi[\pi]$ in experiment H_1 , and oracle access to $\$(\cdot, \cdot, \cdot)$ in experiment H_5 . Experiment H_2 is identical to H_1 , except that we re-sample $\pi \leftarrow \$_s \text{Perm}(\mathcal{T}, n)$ each time we use Enc or Tag . Since a tweak to π is never repeated, $\Pr[H_1^{\mathcal{A}} \Rightarrow \text{true}] = \Pr[H_2^{\mathcal{A}} \Rightarrow \text{true}]$. In experiment H_3 , instead of calling $\text{Tag}^{\pi}(T, X)$ to get the tag, we sample the tag at random. We then unroll procedure $\# \text{Enc}$ and remove the dead code that processes the state X . By applying Lemma 3.1 to lines 02–03 of procedure Priv , the string $V := \text{Tag}^{\pi}(T, X)$ is uniform and so experiments H_2 and H_3 are identical. Next, in experiment H_4 , we directly sample $C_{2m+1} \leftarrow \$_s \{0, 1\}^{|M_{2m+1}|}$ instead of processing via π . Since we never call π on the same tweak twice, each output of $\pi(\cdot, \cdot)$ is an independent, uniform string, and thus experiments H_3 and H_4 are identical. Finally, experiment H_5 is identical to H_4 , except that instead of calling $\text{Enc}^{\pi}(T, X, M_{2i-1} M_{2i})$ to get the blocks $C_{2i-1} C_{2i}$ of the ciphertext, we sample them uniformly. By applying Lemma 3.1 to lines 04–05 of procedure Priv , the output blocks of $\text{Enc}^{\pi}(T, X, M_{2i-1} M_{2i})$ are uniform and independent (even conditioned on all prior ciphertext blocks). Hence H_4 and H_5 are identical, and $\text{Adv}_{\Pi[\pi]}^{\text{priv}}(\mathcal{A}) = \Pr[H_1^{\mathcal{A}} \Rightarrow \text{true}] - \Pr[H_4^{\mathcal{A}} \Rightarrow \text{true}] = 0$. \square

Theorem B.2. Let $\Pi[E] = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ be an AE scheme on $\{0, 1\}^*$ such that $\text{Auth}(G_1^-, G_2^-) = \text{true}$, where G_1^-, G_2^- are the unlabeled graphs for algorithms Dec and Tag of Π , respectively. Then for any adversary \mathcal{A} , there is an adversary \mathcal{B} with $\text{Adv}_{\Pi[E]}^{\text{auth}}(\mathcal{A}) \leq 2^{-\tau} + \ell(\text{Cost}(\Pi) + 2)/2^n + \text{Adv}_E^{\pm \text{PRP}}(\mathcal{B})$, where ℓ is the number of (full) blocks in the forgery output by \mathcal{A} . Adversary \mathcal{B} has the same running time as \mathcal{A} and makes at most $(\text{Cost}(\Pi) + 1) \cdot \sigma/2$ queries, where σ is the total number of (full) message blocks in the queries of \mathcal{A} .

Proof. Adversary \mathcal{B} runs \mathcal{A} . For each of \mathcal{A} 's encryption queries, \mathcal{B} runs the encryption scheme of $\Pi[E]$ but with each call to E_K replaced by a query to \mathcal{B} 's first oracle, and returns the ciphertext to \mathcal{A} . When \mathcal{A}

<pre> proc ENCRYPT[Δ, π](N, A, M) // Experiments $H_1, \boxed{H_2}$ $M_1 \cdots M_{2m} M_{2m+1} := M; X := 0^{2n}; v := 1$ // $M_i = n$ for $i \leq 2n$, and $M_{2m+1} < 2n$ for $i = 1$ to m do $T := (N, A, v); \boxed{\pi \leftarrow \text{Perm}(\mathcal{T}, n)}$ $(Y, C_{2i-1} C_{2i}) := \text{Enc}^\pi(T, X, M_{2i-1} M_{2i})$ $v := v + \text{Cost}(\Pi); X := Y$ $(X, C_{2m+1}) := \text{Enc}^\pi((N, A, v), X, M_{2m+1})$ $v := v + \lceil C_{2m+1} /n \rceil; \boxed{\pi \leftarrow \text{Perm}(\mathcal{T}, n)}$ $T := (N, A, 1 - v); V := \overline{\text{Tag}}^\pi(T, X)$ return $C_1 \cdots C_{2m} C_{2m+1} \parallel V[1, \tau]$ </pre>
<pre> proc ENCRYPT[Δ, π](N, A, M) // Experiment H_3 $M_1 \cdots M_{2m} M_{2m+1} := M; X := 0^{2n}; v := 1$ // $M_i = n$ for $i \leq 2n$, and $M_{2m+1} < 2n$ for $i = 1$ to m do $T := (N, A, v); \pi \leftarrow \text{Perm}(\mathcal{T}, n)$ $(Y, C_{2i-1} C_{2i}) := \text{Enc}^\pi(T, X, M_{2i-1} M_{2i})$ $v := v + \text{Cost}(\Pi); X := Y$ $T := (N, A, v); L := (N, A, v + 1)$ if $M_{2m+1} = \varepsilon$ then $C_{2m+1} := \varepsilon$ elseif $M_{2m+1} > n$ then $M'_1 M'_2 := M_{2m+1} \quad // M'_1 = n$ $V := \pi^T(M'_1); C'_2 := V[1, M'_2] \oplus M'_2$ $C'_1 := \pi^L(C'_2 10^*) \oplus M'_1; C_{2m+1} := C'_1 C'_2$ elseif $M_{2m+1} = n$ then $C_{2m+1} := \pi^T(0^n) \oplus M_{2m+1}$ else $C_{2m+1} := \pi^T(0^n)[1, M_{2m+1}] \oplus M_{2m+1}$ $V \leftarrow \{0, 1\}^n$ return $C_1 \cdots C_{2m} C_{2m+1} \parallel V[1, \tau]$ </pre>
<pre> proc ENCRYPT[Δ, π](N, A, M) // Experiments $H_4, \boxed{H_5}$ $M_1 \cdots M_{2m} M_{2m+1} := M; X := 0^{2n}; v := 1$ // $M_i = n$ for $i \leq 2n$, and $M_{2m+1} < 2n$ for $i = 1$ to m do $T := (N, A, v); \pi \leftarrow \text{Perm}(\mathcal{T}, n)$ $(Y, C_{2i-1} C_{2i}) := \text{Enc}^\pi(T, X, M_{2i-1} M_{2i})$ $\boxed{C_{2i-1} C_{2i} \leftarrow \{0, 1\}^{2n}}$ $v := v + \text{Cost}(\Pi); X := Y$ $T := (N, A, v); L := (N, A, v + 1)$ $C_{2m+1} \leftarrow \{0, 1\}^{ M_{2m+1} }; V \leftarrow \{0, 1\}^n$ return $C_1 \cdots C_{2m} C_{2m+1} \parallel V[1, \tau]$ </pre>

Figure B.4: Experiments H_1 – H_5 in the proof of Theorem B.1. Experiments H_2 and H_5 include the corresponding boxed statements, but H_1 and H_4 do not.

outputs a forgery (N, A, C) , adversary \mathcal{B} runs the decryption scheme of $\Pi[E]$ on (N, A, C) , but with each call to E_K/E_K^{-1} replaced by a query to \mathcal{B} 's oracles. Adversary \mathcal{B} returns 1 if \mathcal{A} output a valid forgery, and returns 0 otherwise. Let $\Pi[\pi]$ be the ideal variant of $\Pi[E]$, where calls to E_K/E_K^{-1} are replaced by queries to π/π^{-1} , respectively, with $\pi \leftarrow \text{Perm}(\mathcal{T}, n)$. It suffices to show that $\text{Adv}_{\Pi[\pi]}^{\text{auth}}(\mathcal{A}) \leq 2^{-\tau} + \ell(\text{Cost}(\Pi) + 2)/2^n$.

Consider experiments H_1 – H_3 in Figure B.5. In H_1 , the adversary has oracle access to the encryption and decryption schemes of $\Pi[\pi]$. Experiment H_2 is identical to H_1 , except that when running the decryption

<pre> proc DECRYPT[π](N, A, C) // Experiments $H_1, \boxed{H_2}$ if $C \not\equiv \tau \pmod{2n}$ then return \perp $C_1 \cdots C_{2m} C_{2m+1} \parallel \text{tag} := C$ // $C_i = n$ for $i \leq 2m$, $C_{2m+1} < 2n$, and $\text{tag} = \tau$ $X := 0^{2n}$; $v := 1$ for $i = 1$ to ℓ do $T := (N, A, v)$; $(Y, M_{2i-1}M_{2i}) := \text{Dec}^{\pi, \pi^{-1}}(T, X, C_{2i-1}C_{2i})$ $v := v + \mathbf{Cost}(\Pi)$; $X := Y$ $(X, M_{2m+1}) := \# \mathbf{Enc}^{\pi}((N, A, v), X, C_{2m+1})$ $\pi \leftarrow \text{s Perm}(\mathcal{T}, n)$; $v := v + \lceil M_{2m+1} /n \rceil - 1$ $T := (N, A, 1 - v)$; $V := \mathbf{Tag}^{\pi}(T, X)$ if $\text{tag} \neq V[1, \tau]$ then return \perp return $M_1 \cdots M_{2m} M_{2m+1}$ </pre>
<pre> proc DECRYPT[π](N, A, C) // Experiment H_3 if $C \not\equiv \tau \pmod{2n}$ then return \perp $C_1 \cdots C_{2m} C_{2m+1} \parallel \text{tag} := C$ // $C_i = n$ for $i \leq 2m$, $C_{2m+1} < 2n$, and $\text{tag} = \tau$ $X := 0^{2n}$; $v := 1$ for $i = 1$ to ℓ do $T := (N, A, v)$; $(Y, M_{2i-1}M_{2i}) := \text{Dec}^{\pi, \pi^{-1}}(T, X, C_{2i-1}C_{2i})$ $v := v + \mathbf{Cost}(\Pi)$; $X := Y$ $(X, M_{2m+1}) := \# \mathbf{Enc}^{\pi}((N, A, v), X, C_{2m+1})$; $V \leftarrow \text{s } \{0, 1\}^n$ if $\text{tag} \neq V[1, \tau]$ then return \perp return $M_1 \cdots M_{2m} M_{2m+1}$ </pre>

Figure B.5: Experiments H_1 – H_3 in the proof of Theorem B.2. Experiment H_2 includes the corresponding boxed statement, but experiment H_1 does not. Each experiment also has a procedure $\mathbf{ENCRYPT}[\pi]$, implementing the encryption algorithm of $\Pi[\pi]$, that is not shown for simplicity.

algorithm, we re-sample $\pi \leftarrow \text{s Perm}(\mathcal{T}, n)$ before using it in \mathbf{Tag} . Experiment H_3 is identical to H_2 , except that instead of using \mathbf{Tag} to generate the tag, we sample the tag uniformly.

Let (N, A, C) be the forgery output by \mathcal{A} . Suppose that there is no encryption query (N, A, M') with $\lceil |M'|/n \rceil = \lceil (|C| - \tau)/n \rceil$. Since decryption of the forgery query involves calling \mathbf{Tag} with a tweak that has never been used before, we have $\Pr[\mathcal{A} \text{ forges in } H_1] = \Pr[\mathcal{A} \text{ forges in } H_2]$. An application of Lemma 3.1 to lines 12–13 of \mathbf{Auth} , shows that the string $V := \mathbf{Tag}^{\pi}(T, X)$ is uniform. Thus $\Pr[\mathcal{A} \text{ forges in } H_2]$ is equal to $\Pr[\mathcal{A} \text{ forges in } H_3]$, which is in turn at most $2^{-\tau}$. Hence $\mathbf{Adv}_{\Pi[\pi]}^{\text{auth}}(\mathcal{A}) \leq 2^{-\tau}$.

Now, suppose that there is an encryption query (N, A, M') such that $\lceil |M'|/n \rceil = \lceil (|C| - \tau)/n \rceil$. (Note that there can be at most one such query, since the attacker is not allowed to re-use a nonce value in two encryption queries.) Let C' be the corresponding ciphertext output by this encryption query. We say that a ciphertext is *complete* if its length (excluding the tag) is a multiple of n . First consider the case that one of C and C' is complete and the other is incomplete. For example, suppose that C is complete and C' is incomplete. Let S be the subset of $\text{Perm}(\mathcal{T}, n)$ such that for any $f \in S$ and for any query (T, X) that encrypting (N, A, M') via $\Pi[\pi]$ makes to π , we have $f(T, X) = \pi(T, X)$. From the proof of Theorem B.1, we have (i) the outputs of the encryption oracle are independent of the key Δ , and (ii) $\pi(T, \cdot)$, for any $T \in \mathbb{N} \times \mathcal{A} \times \mathbb{Z}^+$, is independent of Δ . Let X be the state that \mathbf{Tag} receives on querying (N, A, C) , and let X' be the corresponding state on querying (N, A, C') . Then X is independent of Δ , but $X'[1, n]$ is the xor of Δ and another string that is independent of Δ . Hence $X[1, n] = X'[1, n]$ with probability at most 2^{-n} . If $X[1, n] \neq X'[1, n]$, by applying Lemma 3.2 to lines 20–22 of \mathbf{Auth} , the tag of C is within statistical distance at most $2/2^n$ from a uniform string, independent of the prior ciphertexts. Hence the chance of forgery is at most $3/2^n + 2^{-\tau}$.

<pre> proc DECRYPT[Δ, π](N, A, C) // Experiments P_1, P_2 if $C \not\equiv \tau \pmod{2n}$ then return \perp $C_1 \cdots C_{2m} C_{2m+1} \parallel \text{tag} := C; X := 0^{2n}; v := 1$ // $C_i = n$ and $\text{tag} = \tau$ for $i = 1$ to m do $T := (N, A, v); \pi \leftarrow S$ $(Y, M_{2i-1} M_{2i}) := \text{Dec}^{\pi, \pi^{-1}}(T, X, C_{2i-1} C_{2i})$ $v := v + \text{Cost}(\Pi); X := Y$ $M_{2m+1} := M'_{2m+1} \oplus C'_{2m+1} \oplus C_{2m+1}; X := X \oplus (M_{2m+1} \parallel 0^n)$ $T := (N, A, 1 - v); \pi \leftarrow S; V := \text{Tag}^\pi(T, X)$ if $\text{tag} \neq V[1, \tau]$ then return false return true </pre>
<pre> // Experiments P_{3+j}, for $0 \leq j \leq m - r + 1$ proc DECRYPT[Δ, π](N, A, C) if $C \not\equiv \tau \pmod{2n}$ then return \perp $C_1 \cdots C_{2m} C_{2m+1} \parallel \text{tag} := C; X := 0^{2n}; v := 1$ // $C_i = n$, and $\text{tag} = \tau$ for $i = 1$ to m do $T := (N, A, v); \pi \leftarrow S$ $(Y, M_{2i-1} M_{2i}) := \text{Dec}^{\pi, \pi^{-1}}(T, X, C_{2i-1} C_{2i})$ if $r \leq i \leq r + j$ then $Y \leftarrow \{0, 1\}^n$ $v := v + \text{Cost}(\Pi); X := Y$ $M_{2m+1} := M'_{2m+1} \oplus C'_{2m+1} \oplus C_{2m+1}; X := X \oplus (M_{2m+1} \parallel 0^n)$ $\pi \leftarrow S; V := \text{Tag}^\pi(T, X)$ if $j = m - r + 1$ then $V \leftarrow \{0, 1\}^n$ if $\text{tag} \neq V[1, \tau]$ then return false return true </pre>

Figure B.6: Experiments P_1, \dots, P_{4+m-r} in the proof of Theorem B.2. Experiment P_2 includes the corresponding boxed statement, but P_1 does not. Each experiment also has a procedure $\text{ENCRYPT}[\Delta, \pi]$, implementing the encryption algorithm of $\Pi[\pi]$, that is not shown for simplicity. Here S is the set of $f \in \text{Perm}(\mathcal{T}, n)$ such that for any query (T, X) that $\text{ENCRYPT}[\Delta, \pi](N, A, M')$ makes to π , $f(T, X) = \pi(T, X)$.

Suppose that either (i) both C and C' are complete, or (ii) both C and C' are incomplete. We consider case (i); case (ii) is similar. Let $C = C_1 \cdots C_{2m} C_{2m+1} \parallel \text{tag}$ and $C' = C'_1 \cdots C'_{2m} C'_{2m+1} \parallel \text{tag}'$. If $|C_{2m+1}| = |C'_{2m+1}| = 0$ then the situation is like the basic case, which is already proven by Theorem 3.4. What remains is the case $|C_{2m+1}| = |C'_{2m+1}| = n$. We consider the following cases.

Case 1: $C_j = C'_j$ for every $j \leq 2m + 1$. Then tag and tag' must be different and thus, since Tag is deterministic, the forgery is invalid.

Case 2: $C_j = C'_j$ for every $j \leq 2m$, but $C_{2m+1} \neq C'_{2m+1}$. Let X be the state that Tag receives on querying (N, A, C) , and let X' be the corresponding state on querying (N, A, C') . Then $X[1, n] = X'[1, n] \oplus C_{2m+1} \oplus C'_{2m+1} \neq X'[1, n]$. By applying Lemma 3.2 to lines 20–22 of procedure Auth , the tag of C is within statistical distance at most $2/2^n$ from a uniform string, independent of the prior ciphertexts. Hence the chance of forgery is at most $2/2^n + 2^{-\tau}$.

Case 3: There is $r \leq m$ such that $C_{2r-1} C_{2r} \neq C'_{2r-1} C'_{2r}$, and $C_j = C'_j$ for every $j < 2r - 1$. Consider experiments P_1, \dots, P_{m-r+4} in Figure B.6. In P_1 , the adversary has two oracles: ENCRYPT and DECRYPT . The first implements the encryption scheme of $\Pi[\pi]$, and the second implements the decryption scheme of $\Pi[\pi]$ but returns false if the decrypted value is \perp and returns true otherwise. Experiment P_2 is identical to P_1 , except that in procedure DECRYPT , each time we call Dec or Tag we resample $\pi \leftarrow S$. Since in the

forgery query we do not repeat the tweak of any encryption query other than (N, A, M') , and π and π^{-1} are called with distinct tweaks, we have $\Pr[\mathcal{A} \text{ forges in } P_1] = \Pr[\mathcal{A} \text{ forges in } P_2]$. In experiment P_3 we sample Y uniformly instead of computing $Y := \text{Dec}^{\pi, \pi^{-1}}(T, X, C_{2r-1}C_{2r})$. Applying Lemma 3.2 to lines 14–16 of procedure `Auth`, we have $\Pr[\mathcal{A} \text{ forges in } P_2] - \Pr[\mathcal{A} \text{ forges in } P_3] \leq \frac{2\text{Cost}(\Pi)+2}{2^n}$.

For $j = 1, \dots, m-r$, experiment P_{3+j} is identical to P_{2+j} , except that we sample Y uniformly instead of computing $Y := \text{Dec}^{\pi, \pi^{-1}}(T, X, C_{2r+2j-1}C_{2r+2j})$. Applying Lemma 3.2 to lines 17–19 of procedure `Auth`, we can conclude that $\Pr[\mathcal{A} \text{ forges in } P_{2+j}] - \Pr[\mathcal{A} \text{ forges in } P_{3+j}] \leq \frac{2\text{Cost}(\Pi)+2}{2^n}$.

Experiment P_{m-r+4} is identical to P_{m-r+3} except that we sample V uniformly when checking validity of the forgery instead of computing $V := \text{Tag}^\pi(T, X)$. Let X be the incoming state for `Dec` on querying (N, A, C) and let Y be the outgoing state. Let X' and Y' be the corresponding states on querying (N, A, C') respectively. Then X is a random string, independent of prior ciphertexts and C, X', Y' . Moreover, $Y[1, n] = X[1, n] \oplus M'_{2m+1} \oplus C'_{2m+1} \oplus C_{2m+1}$, which is independent of $Y'[1, n]$. Hence $Y[1, n] \neq Y'[1, n]$ with probability at least $1 - 2^{-n}$. Note that Y and Y' are the state given to `Tag` in `ENCRYPT` $[\Delta, \pi](N, A, C)$ and `ENCRYPT` $[\Delta, \pi](N, A, C')$ respectively. If $Y[1, n] \neq Y'[1, n]$, by applying Lemma 3.2 to lines 20–22 of procedure `Auth`, we have $\Pr[\mathcal{A} \text{ forges in } P_{m-r+3}] - \Pr[\mathcal{A} \text{ forges in } P_{m-r+4}] \leq \frac{2}{2^n}$. Finally, $\Pr[\mathcal{A} \text{ forges in } P_{m-r+4}] \leq 2^{-\tau}$. Summing up,

$$\text{Adv}_{\Pi[\pi]}^{\text{auth}}(\mathcal{A}) \leq 2^{-\tau} + \frac{2(m-r+1)(\text{Cost}(\Pi) + 1) + 3}{2^n} \leq 2^{-\tau} + \frac{\ell(\text{Cost}(\Pi) + 2)}{2^n}. \quad \square$$

C Generating Attacks Automatically

In this section we show how to automatically generate attacks on schemes that do not pass the `Priv` and `Auth` tests. The attacks are sound, but it does not mean that schemes that generate no attack are secure. We consider AE schemes that use the extension in Appendix B to handle fragmentary strings. Below, we categorize the attacks based on the graphs (`Tag`, `Enc`, and `Dec`) in which the schemes fail to pass the required tests. We begin by describing a simple property on graphs, which is a cornerstone to our attacks.

A simple linear relation. Let $G = (d, r, F, P, L)$ be a graph and let E be a tweakable blockcipher with key space \mathcal{K} . Let $S = \{v_1, \dots, v_s\} \subseteq \{1, \dots, r\}$ be the set of all `IN`, `INI`, and `TBC` nodes. Without loss of generality, assume that $v_j = j$ for every $j \leq d$. Then for any node i , there are constants $a_1, \dots, a_s \in \{0, 1\}$ such that the n -bit value of node i is the linear combination of the values at nodes v_1, \dots, v_s with coefficients a_1, \dots, a_s respectively. That is, for any $Z_1, \dots, Z_d \in \{0, 1\}^n$, any key $K \in \mathcal{K}$, and any tweak T , if $(Z_1, \dots, Z_r) := \text{Eval}^{E_K, E_K^{-1}}(G, T, Z_1, \dots, Z_d)$ then $Z_i = a_1 \cdot Z_{v_1} \oplus \dots \oplus a_s \cdot Z_{v_s}$, where $1 \cdot X = X$ and $0 \cdot X = 0^n$.

To justify this claim, we give an induction proof on i . If $i \leq d$ then let $a_j = 1$ if $v_j = i$, and let $a_j = 0$ otherwise. The claim trivially holds for this base case. Suppose that the claim holds for $1, \dots, i-1$. We now show that it holds for i as well.

- If i is a `TBC` node then let $a_j = 1$ if $v_j = i$ and let $a_j = 0$ otherwise.
- If i is a `FIN/DUP/OUT` node then let p be the parent of i . Due to the induction hypothesis, there are constants b_1, \dots, b_s such that the claim holds for p . Let $a_j = b_j$ for every $j \leq s$.
- If i is an `XOR` node then let p_1 and p_2 be the parents of i . Due to the induction hypothesis, there are constants b_1, \dots, b_s such that the claim holds for p_1 , and constants c_1, \dots, c_s such that the claim holds for p_2 . Let $a_j = b_j \oplus c_j$ for every $j \leq s$.

This completes the proof.

We can thus conclude that node i can be represented by a corresponding vector (a_1, \dots, a_s) . Note that the above inductive proof also gives an algorithm to identify the representative vectors of all nodes.

Attacks based on a `Tag` graph. There are two tests that `Priv` and `Auth` perform on a `Tag` graph. First, in lines 02–03 (and also lines 12–13) of Figure 3.7, we check if the output value of the graph is random for

a fresh TBC tweak. If this test fails, then the representative vector (a_1, a_2, a_3) of the **OUT** node must have $a_3 = 0$, since this coordinate corresponds to the TBC node (here we assume that **Tag** has only one TBC instruction). We consider the following cases.

Case 1: $(a_1, a_2) = (0, 0)$, meaning that the output value of the **Tag** graph is always 0^n , regardless of the input. This immediately leads the following privacy attack: query 0^{2n} and return 1 if the tag of the answer is 0^τ , and return 0 otherwise. The adversary wins with advantage at least $1 - 2^{-\tau}$.

Case 2: (a_1, a_2) is $(1, 0)$ or $(0, 1)$. Without loss of generality, assume that $(a_1, a_2) = (1, 0)$, meaning that the output value of G_1 is always the same as the value of the first **INI** node. This is a vulnerability if the first input value of the **Tag** graph is a linear combination of the plaintext blocks, as in OCB or OTR. To determine this linearity, we check if in the **Enc** graph, the representative vector (b_1, \dots, b_s) of the first **FIN** node satisfies $b_j = 0$ for every $j > 4$. If this happens then we can also launch the privacy attack in Case 1. In the **Enc** graph, the values of the **IN** and **INI** nodes are 0^n , and thus the value at the first **FIN** node must be 0^n as well. Hence the output value of the **Tag** graph is also 0^n . The adversary wins with advantage at least $1 - 2^{-\tau}$.

Case 3: $(a_1, a_2) = (1, 1)$, meaning that the output value of the **Tag** graph is always the same as the XOR of the values at the two **INI** nodes. This is a vulnerability if this XOR value is a linear combination of the plaintext blocks. We check if in the **Enc** graph, the representative vectors (b_1, \dots, b_s) and (c_1, \dots, c_s) of the two **FIN** nodes satisfy $b_j = c_j$ for every $j > 4$. If this happens then we can also launch the privacy attack in Case 1. The adversary wins with advantage at least $1 - 2^{-\tau}$.

In the second test, in lines 20–22 of Figure 3.7, we check if the output values are independent when we evaluate the **Tag** graph *twice* on the same tweak but different inputs. (Actually, in the test, the two inputs must differ only in the first half.) Without loss of generality, assume that the **Tag** graph passes the test at line 02–03. Then the representative vector (a_1, a_2, a_3) of the **OUT** node must have $a_3 = 1$. Let p be the parent of the TBC node, and let (d_1, d_2, d_3) be the representative vector of p . Then $d_3 = 0$ because the **Tag** graph has only a single TBC node, and this node is not an ancestor of p . We consider the following cases.

Case 1: $d_1 = 0$, meaning that if we evaluate the **Tag** graph twice, on the same tweak, such that the inputs agree on the second half, then the two results agree on the value at node p , and thus they also agree on the value at the **OUT** node. We now give an authenticity attack. First query 0^n to get answer $C \parallel T$, with $|C| = n$. (Note that this query is a fragmentary string.) Next produce $C^* \parallel T$ as a forgery, for any arbitrary $C^* \neq C$. In both the encryption query and the forgery query, the values at the second **INI** node of the **Tag** graph are 0^n . Hence tag T is also a valid tag for C^* , and thus the adversary wins with advantage 1.

Case 2: $d_1 = d_2 = 1$. We first check if a_1 and a_2 are also 1. If this happens then the output value of the **Tag** graph depends solely on the XOR of the values at the two **INI** nodes. This will be a vulnerability if the XOR value can be linearly determined from the ciphertext blocks. We check if in the **Dec** graph, the representative vectors (x_1, \dots, x_s) and (y_1, \dots, y_s) of the two **FIN** nodes satisfy $x_j = y_j$ for every $j > 4$. If this happens then we launch the following authenticity attack. First query 0^{2n} to get ciphertext $C_1 C_2 \parallel T$, with $|C_1| = |C_2| = n$. Recall that (x_3, x_4) and (y_3, y_4) are coordinates corresponding to the ciphertext blocks in the **FIN** nodes of the **Dec** graph. Let $b = x_3 \oplus x_4 \oplus y_3 \oplus y_4$. If $b = 0$ then output $(C_1 \oplus 1^n) \parallel (C_2 \oplus 1^n) \parallel T$ as a forgery. Note that in the encryption query, the XOR of the values at the two **INI** nodes of the **Tag** graph is $x_3 \cdot C_1 \oplus x_4 \cdot C_2 \oplus y_3 \cdot C_1 \oplus y_4 \cdot C_2$, whereas its counterpart in the forgery query is

$$x_3 \cdot (C_1 \oplus 1^n) \oplus x_4 \cdot (C_2 \oplus 1^n) \oplus y_3 \cdot (C_1 \oplus 1^n) \oplus y_4 \cdot (C_2 \oplus 1^n) = x_3 \cdot C_1 \oplus x_4 \cdot C_2 \oplus y_3 \cdot C_1 \oplus y_4 \cdot C_2.$$

If $b = 1$ then one of $x_3 \oplus y_3$ and $x_4 \oplus y_4$ must be 0 while the other is 1. Without loss of generality, assume that $x_3 \oplus y_3 = 0$. Then output $(C_1 \oplus 1^n) \parallel C_2 \parallel T$ as a forgery. Now in the forgery query, the XOR of the values at the two **INI** nodes of the **Tag** graph is

$$x_3 \cdot (C_1 \oplus 1^n) \oplus x_4 \cdot C_2 \oplus y_3 \cdot (C_1 \oplus 1^n) \oplus y_4 \cdot C_2 = x_3 \cdot C_1 \oplus x_4 \cdot C_2 \oplus y_3 \cdot C_1 \oplus y_4 \cdot C_2.$$

Hence, regardless of the value of b , the encryption query and the forgery query always agree on the XOR of the values at the two **INI** nodes of the **Tag** graph. The adversary thus wins with advantage 1.

Note that the case $(d_1, d_2) = (1, 0)$ cannot happen, because it passes the test at line 20–22: if the two inputs of the **Tag** graph differ in the first half then their values at the TBC node are independent random strings and so are the values at the **OUT** nodes.

Attacks based on an Enc graph. There is a single test that **Priv** performs on an **Enc** graph. At lines 04–06 of Figure 3.7, we check if the output values of the graph are random and independent, for a fresh TBC tweak. If the **Enc** graph fails this test, let (b_1, \dots, b_s) and (c_1, \dots, c_s) be the representative vectors of the two **OUT** nodes of the **Enc** graph. We consider the following cases.

Case 1: $b_j = 0$ for every $j > 4$, meaning that the value of the first **OUT** node can be linearly determined from the values at the **IN** and **INI** nodes. We launch the following privacy attack: query 0^{2n} to get the answer C , and then output 1 if the first n -bit block of C is 0^n , and output 0 otherwise. The adversary wins with advantage $1 - 2^{-\tau}$.

Case 2: $c_j = 0$ for every $j > 4$. This case is similar to Case 1.

Case 3: $b_j = c_j$ for every $j > 4$. Then the XOR of the values at two **OUT** nodes can be linearly determined from the values at the **IN** and **INI** nodes. We launch the following privacy attack. First query 0^{2n} to get the answer $C_1 C_2 T$, with $|C_1| = |C_2| = n$. Next, output 1 if $C_1 = C_2$, and output 0 otherwise. The adversary wins with advantage $1 - 2^{-\tau}$.

Attacks based on a Dec graph. There are two tests that **Auth** performs on a **Dec** graph. First, in lines 14–16 of Figure 3.7, we evaluate the graph *twice* on the same tweak with distinct inputs that agree on the values of the **INI** nodes to check if the values at the first **FIN** node are random and independent.

A possible way to fail this test is that the value in the first **FIN** node depends only on at most a single **IN** node. We thus check if both **IN** nodes are ancestors of the first **FIN** node. If this does not happen then without loss of generality suppose that the first **IN** node is *not* an ancestor of the first **FIN** node. We can launch the following authenticity attack. First query 0^{2n} to get $C_1 C_2 T$, with $|C_1| = |C_2| = n$. Next, output $(C_1 \oplus 1^n) \parallel C_2 \parallel T$ as the forgery attempt. The adversary wins with advantage 1.

Now suppose that both **IN** nodes are ancestors of the first **FIN** node. An example of such a scheme that still fails the test is the insecure variant of **OTR** as illustrated in Figure 3.9. We can exploit this if the **Dec** graph satisfies the following additional constraints:

- (i) There are $b_1, b_2, c_1, c_2 \in \{0, 1\}$ such that $(b_1, b_2) \neq (0, 0)$ and XORing $b_1 \cdot \Delta$ and $b_2 \cdot \Delta$ to the values of the two **IN** nodes, respectively, while keeping the same values of the **INI** nodes, results in $(c_1 \cdot \Delta)$ - and $(c_2 \cdot \Delta)$ -change in the values of the first and second **FIN** nodes, respectively. For example, in the insecure variant of **OTR**, (i) holds for $b_1 = c_1 = 1$ and $b_2 = c_2 = 0$.
- (ii) There exist $a_1, a_2 \in \{0, 1\}$ such that for any $\Delta \in \{0, 1\}^n$, XORing $c_1 \cdot \Delta, c_2 \cdot \Delta, a_1 \cdot \Delta, a_2 \cdot \Delta$ to the values of the **INI** nodes and **IN** nodes, respectively, will not change the value of the first **FIN** node. For example, in the insecure variant of **OTR**, (ii) holds for $a_1 = 1$ and $a_2 = 0$.

To check for these properties, we do the following:

- Consider all $(r_1, r_2) \in \{(1, 0), (0, 1), (1, 1)\}$. If all TBC ancestors (x_1, \dots, x_s) of at least one **FIN** node satisfy $x_3 r_1 = x_4 r_2$ then let $b_1 = r_1$ and $b_2 = r_2$.
- Let (u_1, \dots, u_s) and (w_1, \dots, w_s) be the representative vectors of the first and second **FIN** nodes, respectively. Let $c_1 = u_3 b_1 \oplus u_4 b_2$ and $c_2 = w_3 b_1 \oplus w_4 b_2$.
- Consider all $t_1, t_2 \in \{0, 1\}$. If $c_1 u_1 \oplus c_2 u_2 = t_1 u_3 \oplus t_2 u_4$, and all TBC ancestors (x_1, \dots, x_s) of the first **FIN** node satisfy $c_1 x_1 \oplus c_2 x_2 = t_1 x_3 \oplus t_2 x_4$, then let $a_1 = t_1$ and $a_2 = t_2$.

If in the first check, there are multiple pairs (r_1, r_2) then we find corresponding (t_1, t_2) for each such pair, but only choose a tuple (r_1, r_2, t_1, t_2) to determine $(a_1, a_2, b_1, b_2, c_1, c_2)$. If those checks pass, we launch the following authenticity attack. Query 0^{4n} to get answer $C_1 C_2 C_3 C_4 T$, where $|C_1| = |C_2| = |C_3| = |C_4| = n$.

Then output $(C_1 \oplus b_1 \cdot 1^n) \parallel (C_2 \oplus b_2 \cdot 1^n) \parallel (C_3 \oplus a_1 \cdot 1^n) \parallel (C_4 \oplus a_2 \cdot 1^n) \parallel T$ as the forgery attempt. The adversary wins with advantage 1.

In the second test, in lines 20–22 of Figure 3.7, we evaluate the graph *twice*, on the same tweak, and check if the values at the first FIN node are random and independent, where in one input, the value at the first INI is chosen at random. If the Dec graph fails this test, let (y_1, \dots, y_s) be the representative vector of the first FIN node. We determine if the value at the first FIN node is independent of those at the INI nodes, by checking if y_1 and y_2 are both 0. If this happens then we launch the following authenticity attack. First query 0^{4n} to get answer $C_1 C_2 C_3 C_4 T$, where $|C_1| = |C_2| = |C_3| = |C_4| = n$. Next, query $C_1^* C_2 C_3 C_4 T$, for any $C_1^* \neq C_1$, as a forgery attempt. The adversary wins with advantage 1.

Changelog

- Version 2.0 (August 12, 2015): Major revision, with additional illustrations and an approach for generating concrete attacks on schemes not found secure by our analysis. The running time of OCB and the three schemes in Figure 4.3 is updated: TurboBoost has to be disabled for accurate timing.
- Version 1.0 (June 23, 2015): First release.