

This paper was previously titled “The CWC authenticated encryption (associated data) mode.”

High-speed encryption and authentication: A patent-free solution for 10 Gbps network devices

Tadayoshi Kohno
UC San Diego
9500 Gilman Drive, MC 0114
La Jolla, CA 92093
tkohno@cs.ucsd.edu

John Viega
Virginia Tech
6066 Leesburg Pike, Suite 500
Falls Church, VA 22041
viega@securesoftware.com

Doug Whiting
Hifn, Inc.
5973 Avenida Encinas, Suite 110
Carlsbad, CA 92009
dwhiting@hifn.com

September 1, 2003

Abstract

We introduce CWC, the first patent-free and parallelizable dedicated block cipher mode of operation capable of encrypting and authenticating data at 10 Gbps in hardware using conventional ASIC technology. In addition to being designed for use with future 10 Gbps IPsec network devices, CWC was also designed to be efficient in software on modern CPUs. CWC is also provably secure under the standard “authenticated encryption with associated data” notion assuming that the underlying block cipher is a secure pseudorandom permutation, which is a reasonable assumption if the underlying block cipher is AES. All other “authenticated encryption” block cipher modes are either patent-encumbered (e.g., OCB) or are not parallelizable and therefore not capable of processing data beyond about 2 Gbps in hardware with a single processing unit (e.g., CCM and EAX). Although CWC requires more chip area than OCB, our calculations suggest that the extra silicon costs less than the intellectual property fees for the patented modes. Furthermore, we remark that at least one standardization body (IEEE 802.11) has rejected patented-encumbered modes in favor of patent-free modes, suggesting that the demand for patent-free modes is very high.

1 Introduction

There has recently been a thrust toward producing dedicated block cipher modes of operation capable of simultaneously encrypting and authenticating data. Such modes of operation are often called *authenticated encryption (AE) modes* or, if the modes are capable of authenticating more data than they encrypt, *authenticated encryption with associated data (AEAD) modes*. Despite the previous work in this area, however, there remains at least one area of deficiency: of the previously-existing modes, none of the patent-free ones are capable of encrypting and authenticating data faster than about 2 Gbps in hardware. Yet future high-speed IPsec network devices will be expected to process data at a rate of 10 Gbps.

We address this deficiency in this paper by presenting a patent-free and parallelizable AEAD mode of operation (CWC) capable of encrypting and authenticating data at 10 Gbps using conventional ASIC technology. We do not, however, sacrifice performance in software. In fact, in addition to requiring high-performance in hardware, high-performance in software was an explicit design criterion. We also required that our mode be provably-secure, under the standard AEAD notion [16], assuming that the underlying block cipher is a secure pseudorandom permutation, which is a reasonable assumption if the underlying block cipher is AES. Our resulting construction has other desirable properties as well. For example, it is clean and simple (in our opinion), can process the data online (in the algorithmic sense), uses a single key (thereby avoiding expensive memory accesses in hardware), and allows for pre-processing of associated data and other header fields. Finding a patent-free solution that simultaneously satisfied our hardware, software, and provable-security goals proved to be one of our main challenges; we believe that we have met that challenge.

Let us continue by elaborating on some of the motivations for CWC.

WHY DO WE WANT DEDICATED AUTHENTICATED ENCRYPTION MODES? The traditional approach to achieving authenticated encryption is to combine some standard encryption mode (e.g., CBC mode) with some standard message authentication scheme (e.g., HMAC). This is known as the *generic-composition* approach and was first formally explored in [1] and [10]. Unfortunately, such generic-composition constructions are often *ad hoc* and, as illustrated in [1] and [10], it is very easy to accidentally combine secure encryption modes with secure MACs and still get insecure authenticated encryption modes.

One of the biggest advantages of dedicated AEAD modes over generic-composition AEAD modes is that dedicated AEAD modes are not prone to such accidental errors. That is, since dedicated AEAD modes clearly specify how to achieve both privacy and authenticity, there is no longer the risk of someone accidentally combining a privacy/encryption component with an authenticity/MAC component in an insecure fashion. Furthermore, since most applications that require privacy also require integrity, it is logical to focus on tools capable of providing both services simultaneously. There is thus great value in developing and standardizing dedicated AEAD modes, as evidenced by a wealth of papers in this area [8, 5, 7, 17, 21, 16, 2].

PATENTS. Pragmatically, patents are a major impediment to the standardization and wide-spread deployment of some of the modes presented in the above-mentioned papers. In particular, three independent parties have applied for patents on single-pass authenticated encryption modes. It is not our purpose to describe the specifics of these patent applications (and, indeed, the specifics are not completely known to the public). Rather, we point out that the existence of these patent applications makes many existing authenticated encryption modes less attractive, and therefore less amenable to standardization and deployment. To exemplify this point, we note that although Rogaway, Bellare, Black, and Krovetz's OCB mode [17] is very efficient and elegant, it was appar-

ently rejected from the IEEE 802.11 wireless working group largely because of the fact that it was covered by patent applications from multiple parties.

WHAT IS NEEDED? Noting the need for patent-free dedicated AEAD modes, Whiting, Ferguson, and Housley proposed a patent-free AEAD mode called CCM [21] which, apparently because of its patent-free nature, has been adopted by the IEEE 802.11 working group. CCM was recently followed by another construction, called EAX, by Bellare, Rogaway, and Wagner [2]. Since CCM and EAX are based on the generic-composition approach (they both essentially combine standard counter (CTR) mode encryption with variants of CBC-MAC message authentication), CCM and EAX do not fall under the aforementioned patent applications.

There is, however, one significant disadvantage with both CCM and EAX: the CCM and EAX encryption and decryption operations are not parallelizable. That is, although the CTR mode portions of CCM and EAX are clearly parallelizable, their CBC-MAC portions are not. Parallelizability is, however, very important. For example, without the ability to parallelize the encryption process, using current technology it does not seem possible to build a single hardware engine for CCM or EAX capable of encrypting beyond approximately 2 Gbps.¹ Although 2 Gbps might be adequate for today's applications, such speeds will not be adequate for the coming 10 Gbps network devices.

Therefore, there is a need for a patent-free dedicated mode of operation capable of encrypting and authenticating data at 10 Gbps in hardware. One major motivating example is future IPsec network devices, which may soon have to process data at 10 Gbps.

THE CWC SOLUTION. We propose a general AEAD paradigm, called CWC, that addresses all the aforementioned issues. Our preferred instantiation of CWC for 128-bit block ciphers is un-patented, provably-secure, parallelizable, and efficient in both hardware and software. The parallelizability enables high-speed hardware implementations to encrypt at 10 Gbps when using AES.

The general CWC paradigm is based on what is called the "Encrypt-then-Authenticate generic composition paradigm." In particular, CWC essentially combines a Carter-Wegman message authentication scheme [20] with CTR mode encryption in an Encrypt-then-Authenticate manner. The general idea is as follows: given a pair of strings (A, M) and a nonce N as input, the CWC encapsulation algorithm encrypts M with CTR mode to get some intermediate ciphertext σ . It then uses a Carter-Wegman MAC and the nonce N to MAC the pair (A, σ) . If we let τ denote the resulting MAC tag, then the output of the CWC encapsulation algorithm is the concatenation of σ and τ . CWC is designed to protect the privacy of M and the integrity of both A and M . We defer the intricacies of our specific construction to the body of this paper.

Although based on the Encrypt-then-Authenticate generic composition paradigm, CWC is not a generic composition construction; for example, for efficiency reasons the CWC encryption and MAC components share the same block cipher key. This means, among other things, that we had to prove the security of CWC directly, rather than invoke previous results about the generic composition paradigm. Additionally, because of our performance goals, we developed a new, parallelizable Carter-Wegman MAC for use with our specific CWC instantiation. We again stress that our design was influenced by both our hardware and software goals and our provable-security goals (as well as our patent-free requirement). For example, we rejected designs that performed favorably in software but not in hardware, and we rejected designs that were slightly more efficient but that had weaker provable-security bounds than we desired.

THE CWC INSTANTIATION FOR 128-BIT BLOCK CIPHERS. Throughout the body of this paper

¹It is always possible to build two totally independent units and process two packets at a time, but this is dramatically more complex, requiring twice the area, plus a load balancer.

we will focus on our instantiation of the CWC paradigm for 128-bit block ciphers.² In particular, we focus on CWC-AES, a CWC instantiation with AES as the underlying block cipher. When our results apply to AES with specific key lengths, we shall state so explicitly. Instead of writing CWC-AES, we shall write CWC-BC or simply CWC when we mean the general CWC paradigm instantiated like CWC-AES but with any 128-bit block cipher BC in place of AES.

Note the difference in font between CWC, the general paradigm, and CWC, our specific proposal.

ACHIEVING PARALLELISM. Clearly the CTR mode portion of CWC is parallelizable. Furthermore, the core of the Carter-Wegman MAC portion of CWC (a.k.a. the universal hashing portion of CWC) can be made parallelizable. In the case of CWC, the universal hashing step essentially works by computing

$$Y_1x^n + Y_2x^{n-1} + Y_3x^{n-2} + Y_4x^{n-3} + \dots + Y_nx + Y_{n+1} \bmod 2^{127} - 1 .$$

where Y_1, \dots, Y_n are 96-bit integers and Y_{n+1} is a 127-bit integer corresponding to the pair (A, σ) and x is an integer modulo the prime $2^{127} - 1$. It is well-known that the computation of this polynomial is parallelizable. For example, if we have two engines available, we can rewrite the above polynomial as

$$\left(Y_1y^m + Y_3y^{m-1} + \dots + Y_n \right) x + \left(Y_2y^m + Y_4y^{m-1} + \dots + Y_{n+1} \right) \bmod 2^{127} - 1 ,$$

where $y = x^2 \bmod 2^{127} - 1$, $m = (n - 1)/2$, and we assume for illustrative purposes that n is odd. We can then compute both the left and the right portions of the above in parallel. Additional parallelism can be achieved by further splitting the original polynomial into j polynomials in $y^j = x^j \bmod 2^{127} - 1$.

PERFORMANCE. Let (A, M) be some input to the CWC encapsulation algorithm (recall that A is the associated data and M is the message to encrypt). Assuming that the universal hashing subkey is maintained across invocations, encapsulating (A, M) takes $\lceil |M|/128 \rceil + 2$ block cipher invocations. The polynomial used in CWC's universal hashing step will have degree $d = \lceil |A|/96 \rceil + \lceil |M|/96 \rceil$. There are several ways to evaluate this polynomial (details in Section 4). As noted above, we could evaluate it in parallel. Serially, assuming no precomputation, we could evaluate this polynomial using d 127x127-bit multiplies. As another example, assuming n precomputed powers of the hash subkey, which are cheap to maintain in software for reasonable n , we could evaluate the polynomial using $d - m$ 96x127-bit multiplies and m 127x127-bit multiplies, where $m = \lceil (d + 1)/n \rceil - 1$.

As noted before, it is possible to implement CWC-AES in hardware at 10 Gbps using conventional ASIC technology. Specifically, at 0.13 micron, it takes approximately 300 Kgates to reach 10 Gbps throughput. Table 1 relates the software performance, on a Pentium III, of CWC-AES to the two other patent-free AEAD modes CCM and EAX. The implementations used to compute Table 1 were written in C by Brian Gladman [4] and all use 128-bit AES keys; the current CWC-AES implementation does not use the above-mentioned precomputation approach for evaluating the polynomial. Table 1 shows that the current implementations of the three modes have comparable performance in software, the relative "best" depending on the OS/compiler and the length of the message. Using the above-mentioned precomputation approach and switching to assembly, we anticipate reducing the cost of CWC's universal hashing step to around 8 cpb, thereby significantly improving the performance of CWC-AES in software compared to CCM-AES and EAX-AES (since the authentication portions of CCM-AES and EAX-AES are limited by the speed of AES). For comparison, Bernstein's related hash127, which also evaluates a polynomial modulo $2^{127} - 1$ but

²If desired, it is possible to instantiate the general CWC paradigm with 64-bit block ciphers, although certain limitations (e.g., nonce size) apply to such variants. We do not present a 64-bit CWC variant here since we are primarily concerned with new, high-speed systems using AES, not legacy applications.

Mode	Linux/gcc-3.2.2					Windows 2000/Visual Studio 6.0				
	Payload message lengths (bytes)					Payload message lengths (bytes)				
	128	256	512	2048	8192	128	256	512	2048	8192
CWC-AES	105.5	88.4	78.9	72.2	70.5	84.7	70.2	62.2	56.5	55.0
CCM-AES	97.9	87.1	82.0	78.0	77.1	64.8	56.7	52.5	49.5	48.7
EAX-AES	114.1	94.9	86.1	79.1	77.5	75.2	61.8	55.3	50.4	49.1

Table 1: Software performance (in clocks per byte) for the three patent-free dedicated AEAD modes on a Pentium III. All implementations were in C and written by Brian Gladman [4] and use 128-bit AES keys. Values are averaged over 50 000 samples. Please see the text for additional information and discussion.

whose specific structure makes it less attractive in hardware, runs around 4 cpb on a Pentium III when written in assembly and using the precomputation approach.

We do not claim that CWC-AES will be particularly efficient on low-end CPUs such as 8-bit smartcards. However, our goal was not to develop an efficient AEAD mode for such low-end processors. Rather, our goal was to develop a parallelizable and efficient AEAD mode for 10 Gbps hardware and for modern CPUs.

1.1 Background and related work

The notion of an *authenticated encryption (AE) mode* was formalized by Katz and Yung [8] and by Bellare and Namprempre [1] and the notion of an *authenticated encryption with associated data (AEAD) mode* was formalized by Rogaway [16]. In [1, 10], Bellare–Namprempre and Krawczyk explored ways to combine standard encryption modes with MACs to achieve authenticated encryption. A number of dedicated AE and AEAD modes also exist, including RPC [8], XCBC [5], IACBC [7], OCB [17], CCM [21], and EAX [2]. Within the scope of dedicated block cipher-based AEAD modes, CWC’s closest relatives are CCM and EAX, which also use two passes and are unpatented. From a broader perspective, CWC is similar to the combination of McGrew’s UST [14] and TMMH [13], where one of the main advantages of CWC over UST+TMMH is CWC’s small key size, which can be a bottleneck for UST+TMMH in hardware at high speeds.

Rogaway and Wagner recently released a critique of CCM [18]. For each issue raised in [18], we find that we have already addressed the issue (e.g., we designed CWC to be on-line) or we disagree with the issue (e.g., we feel that it is sufficient for new modes of operation to handle arbitrary octet-length, as opposed to arbitrary bit-length, messages³).

The integrity portion of CWC builds on top of the Carter-Wegman universal hashing approach to message authentication [20]. Like Bernstein’s hash127 [3], CWC’s universal hash function evaluates a polynomial over the integers modulo the prime $2^{127} - 1$. One of the main difference between hash127 and CWC’s universal hash function is that hash127 uses signed 32-bit coefficients and CWC uses unsigned 96-bit coefficients. See Remark 3.2 and Section 4 for discussions on why we chose to use 96-bit coefficients.

In April 2003 we introduced an Internet-Draft, within the IRTF Crypto Forum Research Group, specifying the CWC-AES mode of operation. The latest version of the Internet-Draft can be found at <http://www.zork.org/cwc> or on the IETF website <http://www.ietf.org>.

³Although we stress that, if desired, it is easy to modify CWC to handle arbitrary bit-length messages. See Remark 3.9.

1.2 Outline

We begin in Section 2 with some preliminaries and then describe the CWC mode of operation in Section 3. In Section 4 we discuss the performance of CWC and in Section 5 we present our provable-security results for CWC. Appendix A contains our intellectual property statement. Appendix B presents a summary of CWC’s properties. Appendix C contains the formal proofs of security for CWC, as well as a description of our general CWC paradigm. Appendix D contains test vectors.

2 Preliminaries

NOTATION. If x is a string then $|x|$ denotes its length in bits (not octets). Let ε denote the empty string. If x and y are two equal-length strings, then $x \oplus y$ denotes the XOR of x and y . If x and y are strings, then $x\|y$ denotes their concatenation. If N is a non-negative integer and l is an integer such that $0 \leq N < 2^l$, then $\text{tostr}(N, l)$ denotes the encoding of N as an l -bit string in big-endian format. If x is a string, then $\text{toint}(x)$ denotes the integer corresponding to string x in big-endian format (the most significant bit is *not* interpreted as a sign bit). For example, $\text{toint}(10000010) = 2^7 + 2 = 130$. Let $x \leftarrow y$ denote the assignment of y to x . If X is a set, let $x \xleftarrow{\$} X$ denote the process of uniformly selecting at random an element from X and assigning it to x . If f is a randomized algorithm, let $x \xleftarrow{\$} f(y)$ denote the process of running f with input y and a uniformly selected random tape. When we refer to the time of an algorithm or experiment in the provable security section of this paper, we include the size of the code (in some fixed encoding). There is also an implicit big- \mathcal{O} surrounding all such time references.

AUTHENTICATED ENCRYPTION MODES WITH ASSOCIATED DATA. We use Rogaway’s notion of an *authenticated encryption with associated data (AEAD) mode* [16]. An AEAD mode $\mathcal{SE} = (\mathcal{K}_e, \mathcal{E}, \mathcal{D})$ consists of three algorithms and is defined over some key space $\text{KeySp}_{\mathcal{SE}}$, some nonce space $\text{NonceSp}_{\mathcal{SE}} = \{0, 1\}^n$, n a positive integer, some associated data (header) space $\text{AdSp}_{\mathcal{SE}} \subseteq \{0, 1\}^*$, and some payload message space $\text{MsgSp}_{\mathcal{SE}} \subseteq \{0, 1\}^*$. We require that membership in $\text{MsgSp}_{\mathcal{SE}}$ and $\text{AdSp}_{\mathcal{SE}}$ can be efficiently tested and that if M, M' are two strings such that $M \in \text{MsgSp}_{\mathcal{SE}}$ and $|M'| = |M|$, then $M' \in \text{MsgSp}_{\mathcal{SE}}$.

The randomized key generation algorithm \mathcal{K}_e returns a key $K \in \text{KeySp}_{\mathcal{SE}}$; we denote this process as $K \xleftarrow{\$} \mathcal{K}_e$. The deterministic encryption algorithm \mathcal{E} takes as input a key $K \in \text{KeySp}_{\mathcal{SE}}$, a nonce $N \in \text{NonceSp}_{\mathcal{SE}}$, a header (or associated data) $A \in \text{AdSp}_{\mathcal{SE}}$, and a payload message $M \in \text{MsgSp}_{\mathcal{SE}}$, and returns a ciphertext $C \in \{0, 1\}^*$; we denote this process as $C \leftarrow \mathcal{E}_K^{N,A}(M)$ or $C \leftarrow \mathcal{E}_K(N, A, M)$. The deterministic decryption algorithm \mathcal{D} takes as input a key $K \in \text{KeySp}_{\mathcal{SE}}$, a nonce $N \in \text{NonceSp}_{\mathcal{SE}}$, a header $A \in \text{AdSp}_{\mathcal{SE}}$, and a string $C \in \{0, 1\}^*$ and outputs a message $M \in \text{MsgSp}_{\mathcal{SE}}$ or the special symbol INVALID on error; we denote this process as $M \leftarrow \mathcal{D}_K^{N,A}(C)$. We require that $\mathcal{D}_K^{N,A}(\mathcal{E}_K^{N,A}(M)) = M$ for all $K \in \text{KeySp}_{\mathcal{SE}}$, $N \in \text{NonceSp}_{\mathcal{SE}}$, $A \in \text{AdSp}_{\mathcal{SE}}$, and $M \in \text{MsgSp}_{\mathcal{SE}}$. Let $l(\cdot)$ denote the *length function* of \mathcal{SE} ; i.e., for all keys K , nonces N , headers A , and messages M , $|\mathcal{E}_K^{N,A}(M)| = l(|M|)$.

Under the correct usage of an AEAD mode, after a random key is selected, the application should never invoke the encryption algorithm twice with the same nonce value until a new key is randomly selected. In order to ensure that a nonce does not repeat, implementations typically use nonces that contain counters. We use the notion of a nonce, rather than simply a counter, because the notion of a nonce is more general and allows the developer the freedom to structure the nonce as he or she desires.

BLOCK CIPHERS. A block cipher $E : \{0, 1\}^k \times \{0, 1\}^L \rightarrow \{0, 1\}^L$ is a function from k -bit keys and L -bit blocks to L -bit blocks. We use $E_K(\cdot)$, $K \in \{0, 1\}^k$, to denote the function $E(K, \cdot)$ and we

use $f \stackrel{\$}{\leftarrow} E$ as short hand for $K \stackrel{\$}{\leftarrow} \{0,1\}^k$; $f \leftarrow E_K$. Block ciphers are families of permutations; namely, for each key $K \in \{0,1\}^k$, E_K is a permutation on $\{0,1\}^L$. We call k the key length of E and we call L the block length.

3 The CWC mode of operation

We now describe the CWC mode of operation for 128-bit block ciphers. See Appendix C for a description of the general CWC construction.

If BC denotes a block cipher with 128-bit blocks and kl -bit keys, and if $tl \leq 128$ is the desired tag length for CWC in bits, then let CWC-BC- tl denote the CWC mode of operation instantiated with BC using tag length tl . Throughout the remainder of this section, fix BC and tl and let CWC-BC- $tl = (\mathcal{K}, \text{CWC-ENC}, \text{CWC-DEC})$.

We associate to CWC-BC- tl the following sets:

$$\begin{aligned} \text{MsgSp}_{\text{CWC-BC-}tl} &= \{ x \in (\{0,1\}^8)^* : |x| \leq \text{MaxMsgLen} \} \\ \text{AdSp}_{\text{CWC-BC-}tl} &= \{ x \in (\{0,1\}^8)^* : |x| \leq \text{MaxAdLen} \} \\ \text{KeySp}_{\text{CWC-BC-}tl} &= \{0,1\}^{kl} \\ \text{NonceSp}_{\text{CWC-BC-}tl} &= \{0,1\}^{88} \end{aligned}$$

where MaxMsgLen and MaxAdLen are both $128 \cdot (2^{32} - 1)$. That is, the payload and associated data spaces for CWC-BC- tl consist of all strings of octets that are at most $2^{32} - 1$ blocks long.

3.1 The CWC core

The key generation algorithm \mathcal{K} returns a randomly selected key from $\text{KeySp}_{\text{CWC-BC-}tl}$ (i.e., the key generation returns a random kl -bit string). The encryption algorithm CWC-ENC works as follows:

```
Algorithm CWC-ENCK(N, A, M) // CWC encryption
  σ ← CWC-CTRK(N, M)
  τ ← CWC-MACK(N, A, σ)
  Return σ||τ
```

where CWC-CTR and CWC-MAC are described in Section 3.2. The decryption algorithm CWC-DEC works as follows:

```
Algorithm CWC-DECK(N, A, C) // CWC decryption
  If |C| < tl then return INVALID
  Parse C as σ||τ where |τ| = tl
  If A ∉ AdSpCWC-BC-}tl or σ ∉ MsgSpCWC-BC-}tl then return INVALID
  If τ ≠ CWC-MACK(N, A, σ) then return INVALID
  Return CWC-CTRK(N, σ)
```

3.2 The CWC subroutines

The remaining CWC algorithms are defined as follows:

```
Algorithm CWC-CTRK(N, M) // CWC counter mode module
  α ← ⌈|M|/128⌉
  For i = 1 to α do
```

$ks_i \leftarrow \text{BC}_K(10^7 \| N \| \text{tostr}(i, 32))$ // Note that 10^7 means a one bit followed by 7 zeros
 $\sigma \leftarrow (\text{first } |M| \text{ bits of } ks_1 \| ks_2 \| \dots \| ks_\alpha) \oplus M$
 Return σ

Algorithm $\text{CWC-MAC}_K(N, A, \sigma)$ // CWC authentication module
 $R \leftarrow \text{BC}_K(\text{CWC-HASH}_K(A, \sigma))$
 $\tau \leftarrow \text{BC}_K(10^7 \| N \| 0^{32}) \oplus R$
 Return first tl bits of τ

Algorithm $\text{CWC-HASH}_K(A, \sigma)$ // CWC universal hashing module
 $Z \leftarrow \text{last 127 bits of } \text{BC}_K(110^{126})$
 $K_h \leftarrow \text{toint}(Z)$ // The same K_h value is used in every invocation of CWC-HASH_K .
 $l \leftarrow \text{minimum integer such that } 96 \text{ divides } |A| \| 0^l$
 $l' \leftarrow \text{minimum integer such that } 96 \text{ divides } |\sigma| \| 0^{l'}$
 $X \leftarrow A \| 0^l \| \sigma \| 0^{l'}$; $\beta \leftarrow |X|/96$; $l_\sigma \leftarrow |\sigma|/8$; $l_A \leftarrow |A|/8$
 Break X into chunks X_1, X_2, \dots, X_β // $|X_1| = |X_2| = \dots = |X_\beta| = 96$
 For $i = 1$ to β do
 $Y_i \leftarrow \text{toint}(X_i)$
 $Y_{\beta+1} \leftarrow 2^{64} \cdot l_A + l_\sigma$ // Include the lengths of A and σ in the polynomial.
 $R \leftarrow Y_1 K_h^\beta + \dots + Y_\beta K_h + Y_{\beta+1} \text{ mod } 2^{127} - 1$
 Return $\text{tostr}(R, 128)$ // Note: first bit of result will always be 0

3.3 Remarks

We now highlight some features of CWC, explain some of our design decisions, and discuss some of the alternatives we explored. We have additional remarks in Section 5.4.

Remark 3.1 [Computing the CWC-HASH polynomial.] The polynomial

$$Y_1 K_h^\beta + \dots + Y_\beta K_h + Y_{\beta+1} \text{ mod } 2^{127} - 1$$

in CWC-HASH can be computed using Horner's Rule as

$$((((Y_1)K_h + Y_2)K_h + \dots)K_h + Y_\beta)K_h + Y_{\beta+1} \text{ mod } 2^{127} - 1.$$

Alternatively, if the values K_h^i are precomputed, the polynomial can be computed directly.

Furthermore, as discussed in the introduction, computation of the polynomial in CWC-HASH can be parallelized by splitting the polynomial into multiple polynomials in K_h^i for some i .

As we will see in Section 4, different implementations will want to evaluate the polynomial in different ways. For example, in software it is advantageous to precompute the powers of the K_h and evaluate the polynomial directly. To avoid unnecessary memory accesses, hardware implementations will likely evaluate the polynomial using Horner's rule (perhaps by first splitting the polynomial in order to exploit CWC-HASH's parallelism).

Remark 3.2 [On the size of the CWC-HASH coefficients.] All the coefficients Y_1, \dots, Y_β in CWC-HASH are 96-bit integers. When evaluating the polynomial using precomputed powers of K_h , the cost for each coefficient includes the cost of a 96x127-bit multiply. When evaluating the polynomial using Horner's rule, the cost for each coefficient includes the cost of a 127x127-bit multiply (since the accumulated value will be 127 bits long). (See Remark 5.4 for why we chose not to use 96-bit hash subkeys and, particularly relevant here, the fact that when we split the polynomial and evaluate two polynomials in K_h^i , $i \geq 2$, K_h^i will likely be 127 bits long even if K_h is

96 bits). Since we are already performing 127x127-bit multiplies, to increase the performance when using Horner’s rule it would easily be possible to define CWC to use coefficients up to 126-bits in size. Such an approach would speed up the Horner’s rule computation by a ratio of 126/96 (nearly 4:3), but would require considerable additional complexity to perform bit and byte shifting within the coefficients. Note that Bernstein’s related hash127 [3] uses smaller 32-bit coefficients, which makes it more costly in hardware when evaluating the polynomial using Horner’s rule, but cheaper in software when using precomputed powers of the hash subkey. We use 96-bit coefficients because it provides for fast hardware implementations (using Horner’s rule) and fast (although not as fast as hash127) software implementation when using precomputation. See Section 4 for additional discussion. (Finally, the final $Y_{\beta+1}$ may be larger than 96-bits since $Y_{\beta+1}$ does not have to be multiplied with anything.)

Remark 3.3 [Why a single key.] It would be perfectly acceptable from a security perspective to make the block cipher key K and hash key K_h independent. The main motivation for using a single key, and deriving the hash key K_h from the block cipher key K , was simplicity of key management. From a performance perspective, we note that fetching key material can be a bottleneck in high-speed hardware.

Remark 3.4 [Separating block cipher inputs.] The input to the block cipher when generating the hash key K_h begins with the bits 11. All the inputs when generating CTR mode keystream begin with the bits 10. The input to the keystream generator in CWC-MAC has the last 32 bits all zero and the input to the block cipher in CWC-CTR never has the last 32 bits zero. All the outputs of CWC-HASH begin with a 0 bit. These properties ensure that there is never an overlap in the inputs between the different uses of the underlying block cipher. For example, the output of the universal hash function (which is enciphered with the block cipher) will never collide with one of the inputs to the block cipher in CWC-CTR. Essentially, separating the block cipher inputs in this way is what allows us to use a single block cipher key in all applications of the block cipher.

Remark 3.5 [Why not derive multiple keys from a single key?] It would be possible to define a mode of operation that takes a single master key and that derives “independent” encryption and MAC block cipher keys from the master key. Doing so would eliminate the need to be careful about separating inputs to the block cipher (Remark 3.4), but would require additional computations (most likely block cipher invocations) to derive the encryption and MAC keys if implementations only store the master in memory. Furthermore, unless implementations store the expanded keys in memory, there would be the additional cost of expanding the key schedules for the derived encryption and MAC keys. Since we can provably use the same block cipher key for all applications of the underlying block cipher (Remark 3.4 and Section 5), since our solution avoids unnecessary precomputation steps, and since we believe our solution is still clean and simple, we chose not to derive “independent” encryption and MAC block cipher keys from a single master key.

Remark 3.6 [Computing the universal hash subkey.] Although CWC-HASH shows the hash subkey K_h being computed upon every invocation, it is possible to compute K_h in the key generation step of CWC. Doing so would save one block cipher application per message but would require maintaining an additional 128 bits across invocations. We anticipate that in hardware, where fetching key material can be expensive, the hash subkey will be re-computed on every invocation of the encryption and decryption algorithms. In software, however, we anticipate that the subkey K_h will be computed once and maintained across invocations.

Remark 3.7 [On the choice of parameters.] The parameters (e.g., the nonce length and the way the nonce is encoded in the input to the block cipher) are fixed for CWC. This is in order to promote interoperability. In CWC the block counter length is set to 32 bits in order to allow CWC to be used with IPsec jumbograms and other large packets up to $2^{32} - 1$ blocks long. The nonce length is set to 88 bits in order to handle future IPsec sequence numbers.

Remark 3.8 [Byte ordering.] CWC uses big-endian byte ordering. We do so for consistency purposes and in order to maintain compatibility with McGrew’s ICM Internet-Draft [12] and the IETF, which strongly favors the big-endian byte-ordering.

Remark 3.9 [Handling arbitrary bit-length messages.] Although we could have specified CWC to take arbitrary bit-length messages as input (just change the definitions of the message spaces and compute $l_A \leftarrow |A|$ and $l_\sigma \leftarrow |\sigma|$ in CWC-HASH), we do not specify CWC this way simply because there does not appear to be a significant need to handle arbitrary bit-length messages and we do not consider it a good trade-off to define a mode for arbitrary bit-length messages at the expense of octet-oriented systems.

If, in the future, such a need arises, it will still be possible to modify the current CWC construction to take arbitrary bit-length messages as input without affecting interoperability with existing CWC implementations when octet-strings are communicated. Although other possibilities exist, one method would be to augment the computation of $Y_{\beta+1}$ in CWC-HASH as follows:

$$r_A \leftarrow |A| \bmod 8; r_\sigma \leftarrow |\sigma| \bmod 8; Y_{\beta+1} \leftarrow 2^{120} \cdot r_A + 2^{112} \cdot r_\sigma + 2^{64} \cdot l_A + l_\sigma .$$

Remark 3.10 [64-bit block ciphers.] It is possible to instantiate the general CWC paradigm (Appendix C) with 64-bit block ciphers like DES and 3DES. We do not do so in this paper since we are targeting future high-speed cryptographic applications.

Remark 3.11 [Initial counter for CTR-mode.] Motivated by EAX2 [2], one possible alternative to CWC might be to use $\text{BC}_K(1110^5\|N)$ both as the value to encrypt R in CWC-MAC and as the initial counter to CTR mode-encrypt M (with the first two bits of the counter always set to 10). Other EAX2-motivated constructions also exist. For example, the tag might be set to $\text{BC}_K(h(X_0\|N)) \oplus \text{BC}_K(h(X_1\|A)) \oplus \text{BC}_K(h(X_2\|\sigma))$, where X_0, X_1, X_2 are strings, none of which is a prefix of the other, and h is a parallelizable universal hash function, like CWC-HASH but hashing only single strings (as opposed to pairs of strings). Compared to CWC, these alternatives have the ability to take longer nonces as input, and, from a functional perspective, can be applied to strings up to 2^{126} blocks long. But we do not view this as a reason to prefer these alternatives over CWC. From a practical perspective, we do not foresee applications needing nonces longer than 11 octets, or needing to encrypt messages longer than $2^{32} - 1$ blocks. Moreover, from a security perspective, applications should not encrypt too many packets between rekeyings, implying that even 11 octet nonces are more than sufficient.

4 Performance

4.1 Hardware

Since one of our main goals is to achieve 10 Gbps in hardware, and in particular for future high-speed IPsec network devices, let us focus first on hardware costs. As noted in the introduction, using 0.13 micron CMOS ASIC technology, it should take approximately 300 Kgates to achieve 10 Gbps throughput for CWC-AES. This estimate, which is applicable to AES with all key lengths, includes

four AES counter-mode encryption engines, each running at 200 MHz and requiring about 25K gates each. In addition, there are two 32x128-bit multiply/accumulate engines, each running at 200 MHz with a latency of four clocks, one each for the even and odd polynomial coefficients. Of course, simply keeping these engines “fed” may be quite a feat in itself, but that is generally true of any 10 Gbps path. Also, there may well be better methods to structure an implementation, depending on the particular ASIC vendor library and technology, but, regardless of the implementation strategy, 10 Gbps is quite achievable because of the inherent parallelism of CWC.

Since OCB is CWC’s main competitor for high-speed environments, it is worth comparing CWC with OCB instantiated with AES (we do not compare CWC with CCM and EAX here since the latter two are not parallelizable). We first note that CWC-AES saves some gates because we only have to implement AES encryption in hardware. However, at 10 Gbps, OCB still probably requires only about half the silicon area of CWC-AES. The main question for many hardware designers is thus whether the extra silicon area for CWC-AES costs more than three royalty payments, as well as negotiation costs and overhead. Our estimates indicate that, given today’s silicon costs, the extra silicon for CWC-AES is probably cheaper than the IP fees for OCB.

4.2 Software

CWC-AES can also be implemented efficiently in software. Table 1 shows timing information for CWC-AES, as well as CCM-AES and EAX-AES, on a 1.133GHz mobile Pentium III dual-booting RedHat Linux 9 (kernel 2.4.20-8) and Windows 2000 SP2. The numbers in the table are the clocks per byte for different message lengths averaged over 50 000 runs and include the entire time for setting up (e.g., expanding the AES key-schedule) and encrypting. All implementations were in C and written by Brian Gladman [4] and use 128-bit AES keys. The Linux compiler was gcc version 3.2.2; the Windows compiler was Visual Studio 6.0.

From Table 1 we conclude that the three patent-free modes, as currently implemented by Gladman, share similar software performances. The “best” performing one appears to depend on OS/compiler and the length of the message being processed. On Linux, it appears that CWC-AES performs slightly better than EAX-AES for all message lengths that we tested, and better than CCM-AES for the longer messages, whereas Gladman’s CCM-AES and EAX-AES implementations slightly outperform his CWC-AES implementation on Windows for all the message lengths that we tested.

Note, however, that all the implementations used to compute Table 1 were written in C. Furthermore, the current CWC-AES code does not make use of all of the optimization techniques (and in particular precomputation) that we describe below. By switching to assembly and using the additional optimization techniques, we anticipate the speed for CWC-HASH to drop to better than 8 clocks per byte, whereas the speed for the CBC-MAC portion of CCM-AES and EAX-AES will be limited by the speed of AES (the best reported speed for AES on a Pentium III is 14.1 cpb, due to a proprietary library by Helger Lipmaa; Gladman’s free hand-optimized Windows assembly implementation runs at 17.5 cpb [11]). Returning to the speed of CWC-HASH, for reference we note that Bernstein’s related hash127 [3] runs around 4 cpb on a Pentium III when written in assembly and using the precomputation approach. Bernstein’s hash127 also works by evaluating a polynomial modulo $2^{127} - 1$; the main difference is that the coefficients for hash127 are 32 bits long, whereas the coefficients for CWC-HASH are 96 bits long (recall Remark 3.2, which discusses why we use 96-bit coefficients).

4.2.1 Implementing CWC-HASH in software

Since the implementation of CWC-HASH is more complicated than the implementation of the CWC-CTR portion of CWC, we devote the rest of this section to discussing CWC-HASH.

PRECOMPUTATION. As noted in Remark 3.1, there are two general approaches to implementing CWC-HASH in software. The first is to use Horner’s rule. The second is to evaluate the polynomial directly, which can be faster if one precomputes powers of the hash key K_h at setup time (here the powers of K_h can be viewed as an expanded key-schedule). In particular, as noted in Remark 3.2, evaluating the polynomial using Horner’s rule requires a 127x127-bit multiply for each coefficient, whereas evaluating the polynomial directly using precomputed powers of K_h requires a 96x127-bit multiply for each coefficient.⁴ The advantage with precomputation was first observed by Bernstein in the context of hash127 [3].

The above description of the precomputation approach assumed that if the polynomial is $Y_1K_h^{\gamma-1} + \dots + Y_{\gamma-1}K_h + Y_\gamma$ (i.e., the polynomial has γ coefficients), then we had precomputed the powers of K_h^i for all $i \in \{1, \dots, \gamma - 1\}$. The precomputation approach extends naturally to the case where we have precomputed the powers K_h^j , $j \in \{1, \dots, n\}$, for some $n \leq \gamma - 1$. For simplicity, first assume that we know the polynomial has a multiple of n coefficients. For such a polynomial, one processes the first n coefficients (to get $Y_1K_h^{n-1} + \dots + Y_{n-1}K_h + Y_n$), then multiplies the intermediate result by K_h^n (to get $Y_1K_h^{2n-1} + \dots + Y_{n-1}K_h^{n+1} + Y_nK_h^n$). After that, one can continue processing data with the same precomputed values (to get $Y_1K_h^{2n-1} + \dots + Y_{2n-1}K_h + Y_{2n}$), and so on. Note that each chunk of n coefficients takes $(n - 1)$ 96x127-bit multiplies, and all but the last chunk takes an additional 127x127-bit multiply. Now assume that the number of coefficients m in the polynomial is not necessarily a multiple of n . If m is known in advance, one could first process $m \bmod n$ coefficients, multiply by K_h^n , then process in n -coefficient chunks as before. Alternately, as long as the end of the message is known n coefficients in advance, one could process n -coefficient chunks, and then finish off the final $m \bmod n$ coefficients using Horner’s rule. Or, if the number of coefficients in the polynomial is not known until the final coefficient is reached, one could process the message in n -coefficient chunks and then multiply by a precomputed power of K_h^{-1} once the end of the message hash been reached.

Naturally, precomputation requires extra memory, but that is usually cheap and plentiful in a software-based environment. Using 32-bit multiplies, the precomputation approach requires 12 32-bit multiplies per 96-bit coefficient, as well as 17 adds, all of which may carry. In assembly, most of these carry operations can be implemented for free, or close to it by using a special variant of the add instruction that adds in the operand as well as the value of the carry from the previous add operation. But when implemented in C, they will generally compile to code that requires a conditional branch and an extra addition. An implementation using Horner’s rule requires an additional four multiplies and three additions with carry per coefficient, adding about 33% overhead, since the multiplies dominate the additions. A 64-bit platform only requires four multiplies and four adds (which may all carry), no matter the implementation strategy taken. The multiply being far more expensive than other operations, we would thus expect a 64-bit integer implementation to run in one third the time of a 32-bit implementation, assuming that the cost of primitive operations does not increase.

EXPLOITING THE PARALLELISM OF SOME INSTRUCTION SETS. On most platforms, it turns out that the integer execution unit is not the fastest way to implement CWC-HASH. Many platforms have multimedia instructions that can be used to speed up the implementation. As another alternative,

⁴As an aside, see Remark 5.4 for why we did not make the hash subkey 96-bits, which could have sped up a serial Horner’s rule implementation.

Bernstein demonstrated that, on most platforms, the floating point unit can be used to implement this class of universal hash functions far more efficiently than can be done in the integer unit. This is particularly true on the x86 platform where, in contrast to using the standard registers, two floating point multiplies can be started in close proximity without introducing a pipeline stall. That is, the x86 can effectively perform two floating-point operations in parallel. The disadvantage of using floating-point registers is that the operands for the individual multiplies need to be small, so that the operations can be done without loss of precision. On the x86, Bernstein multiplies 24-bit values, allowing the sums of product terms to fit into double precision values with 53 bits of precision without loss of information. Bernstein details many ways to optimize this sort of calculation in [3].

As noted before, there are only two main differences between the structure of the polynomials of Bernstein’s hash127 and CWC-HASH. The first is that Bernstein uses signed coefficients, whereas CWC-HASH uses unsigned coefficients; this should not have an impact on efficiency. The other difference is that Bernstein uses 32-bit coefficients, whereas CWC-HASH uses 96-bit coefficients. While both solutions average one multiplication per byte when using integer math, Bernstein’s solution requires only .75 additions per byte, whereas CWC-HASH requires 1.42 additions per byte, nearly twice as many. Using 32-bit multiplies to build a 96x127 multiplier (assuming precomputation), CWC-HASH should therefore perform no worse than at half the speed of hash127. When using 24-bit floating point coefficients to build a multiply (without applying any non-obvious optimizations), hash127 requires 12 multiplies and 16 adds per 32-bit word. CWC can get by with 8 multiples per word and 12.67 additions per word. This is because a 96-bit coefficient fits exactly into four 24-bit values, meaning we can use a 6x4 multiply for every three words. With 32-bit coefficients, we need to use two 24-bit values to represent each coefficient, resulting in a single 6x2 multiply that needs to be performed for each word.

Gladman’s implementation of CWC-HASH uses floating point arithmetic, but uses Horner’s rule instead of performing precomputation to achieve extra speed. Nothing about the CWC hash indicates that it should run any worse than half the speed of hash127, if implemented in a similar manner, in assembly, and using the floating point registers and precomputation. This upper-bound paints an encouraging picture for CWC performance, because hash127 on a Pentium III runs around 4 cpb when implemented in assembly and using the floating point registers and precomputation. This indicates that a well-optimized software version of CWC-HASH should run no slower than 8 cycles per byte.

Finally, it may be possible to further improve the performance of CWC-HASH. For example, literature from the gaming community [6] indicates that one can use both integer and floating point registers in parallel. Although we have not tested this approach, it seems reasonable to conclude that one might be able to interleave integer operations, and thereby obtain additional speedups.

5 Theorem statements

In addition to parallelizability and performance, provable-security was one of our major design requirements (we rejected several constructions that had weaker provable-security results than we desired). Consequently, the CWC mode is a provably secure AEAD mode assuming that the underlying block cipher (e.g., AES) is a secure pseudorandom permutation. This is a quite reasonable assumption since most modern block ciphers (including AES) are believed to be pseudorandom. Furthermore, all provably-secure block cipher modes of operation that we are aware of make the same assumptions we make (and some modes, e.g. OCB [17], make even stronger, albeit still reasonable, assumptions).

The specific results for CWC appear as Theorem 5.1 and Theorem 5.2 below. In Appendix C we also present results for the general CWC paradigm, from which Theorems 5.1 and 5.2 follow.

5.1 Preliminaries

Before presenting our provable security results, we must first formally describe what we mean by privacy and integrity/authenticity. Our privacy and integrity/authenticity notions for AEAD modes come from [17]. We must also describe the notion of a pseudorandom permutation.

PRIVACY OF AEAD MODES. Let $\mathcal{SE} = (\mathcal{K}_e, \mathcal{E}, \mathcal{D})$ be an AEAD mode with length function $l(\cdot)$. Let $\mathcal{E}(\cdot, \cdot, \cdot)$ be an oracle that, on input $(N, A, M) \in \text{NonceSp}_{\mathcal{SE}} \times \text{AdSp}_{\mathcal{SE}} \times \text{MsgSp}_{\mathcal{SE}}$, returns a random string of length $l(|M|)$. Let B be an adversary with access to an oracle and that returns a bit. Then

$$\mathbf{Adv}_{\mathcal{SE}}^{\text{priv}}(B) = \Pr \left[K \xleftarrow{\$} \mathcal{K}_e : B^{\mathcal{E}_K(\cdot, \cdot, \cdot)} = 1 \right] - \Pr \left[B^{\mathcal{E}(\cdot, \cdot, \cdot)} = 1 \right]$$

is the *IND \mathcal{E} -CPA-advantage* of B in breaking the privacy of \mathcal{SE} under chosen-plaintext attacks; i.e., $\mathbf{Adv}_{\mathcal{SE}}^{\text{priv}}(B)$ is the advantage of B in distinguishing between ciphertexts from $\mathcal{E}_K(\cdot, \cdot, \cdot)$ and random strings. An adversary B is *nonce-respecting* if it never queries its oracle with the same nonce twice. Intuitively, a mode \mathcal{SE} preserves privacy under chosen plaintext attacks if the IND \mathcal{E} -CPA-advantage of all nonce-respecting adversaries using reasonable resources is small.

INTEGRITY/AUTHENTICITY OF AEAD MODES. Let $\mathcal{SE} = (\mathcal{K}_e, \mathcal{E}, \mathcal{D})$ be an AEAD mode. Let F be a forging adversary and consider an experiment in which we first pick a random key $K \xleftarrow{\$} \mathcal{K}_e$ and then run F with oracle access to $\mathcal{E}_K(\cdot, \cdot, \cdot)$. We say that F *forges* if F returns a pair (N, A, C) such that $\mathcal{D}_K^{N, A}(C) \neq \text{INVALID}$ but F did not make a query (N, A, M) to $\mathcal{E}_K(\cdot, \cdot, \cdot)$ that resulted in a response C . Then

$$\mathbf{Adv}_{\mathcal{SE}}^{\text{auth}}(F) = \Pr \left[K \xleftarrow{\$} \mathcal{K}_e : F^{\mathcal{E}_K(\cdot, \cdot, \cdot)} \text{ forges} \right]$$

is the *AUTH-advantage* of F in breaking the integrity/authenticity of \mathcal{SE} . Intuitively, the mode \mathcal{SE} preserves integrity/authenticity if the AUTH-advantage of all nonce-respecting adversaries using reasonable resources is small.

PSEUDORANDOM PERMUTATIONS. If X is a set, then $\text{Perm}(X)$ denotes the set of all permutations on X . If L is a positive integer, then $\text{Perm}(L)$ denotes the set of all permutations on $\{0, 1\}^L$. Let F be a family of functions from set D to D . Let A be an adversary with access to an oracle and that returns a bit. Then

$$\mathbf{Adv}_F^{\text{PRP}}(A) = \Pr \left[f \xleftarrow{\$} F : A^{f(\cdot)} = 1 \right] - \Pr \left[g \xleftarrow{\$} \text{Perm}(D) : A^{g(\cdot)} = 1 \right]$$

denotes the PRP-advantage of A in distinguishing a random instance of F from a random permutation. Intuitively, we say that F is a secure PRP if the PRP-advantages of all adversaries using reasonable resources is small. Modern block ciphers, such as AES, are believed to be secure PRPs.

5.2 Integrity/authenticity

Theorem 5.1 [Integrity/authenticity of CWC.] Let CWC-BC-tl be as in Section 3. (Recall that BC is a 128-bit block cipher and that the tag length tl is ≤ 128 .) Consider a nonce-respecting AUTH adversary A against CWC-BC-tl. Assume the execution environment allows A to query its oracle with associated data that are at most $n \leq \text{MaxAdLen}$ bits long and with messages that are at most $m \leq \text{MaxMsgLen}$ bits long. Assume A makes at most $q - 1$ oracle queries and the total length of all the payload data (both in these $q - 1$ oracle queries and the forgery attempt) is at

most μ . Then given A we can construct a PRP adversary C_A against BC such that

$$\mathbf{Adv}_{\text{CWC-BC-tl}}^{\text{auth}}(A) \leq \mathbf{Adv}_{\text{BC}}^{\text{prp}}(C_A) + \frac{(\mu/128 + 3q + 1)^2}{2^{129}} + \frac{n + m}{2^{133}} + \frac{1}{2^{125}} + \frac{1}{2^{\text{tl}}}. \quad (1)$$

Furthermore, the experiment for C_A takes the same time as the experiment for A and C_A makes at most $\mu/128 + 3q + 1$ oracle queries. ■

The above theorem means that if the underlying block cipher is a secure pseudorandom permutation, then CWC-BC will preserve authenticity. If the underlying block cipher is something like AES, then this initial assumption seems quite reasonable and, therefore, CWC-AES will preserve authenticity.

Let us elaborate on why Theorem 5.1 implies that CWC-BC will preserve authenticity. Assume BC is a secure block cipher. This means that $\mathbf{Adv}_{\text{BC}}^{\text{prp}}(C)$ must be small for all adversaries C using reasonable resources and, in particular, this means that, for C_A as described in the theorem statement, $\mathbf{Adv}_{\text{BC}}^{\text{prp}}(C_A)$ must be small assuming that A uses reasonable resources. And if $\mathbf{Adv}_{\text{BC}}^{\text{prp}}(C_A)$ is small and μ, q, m and n are small, then, because of the above equations, $\mathbf{Adv}_{\text{CWC-BC-tl}}^{\text{auth}}(A)$ must also be small as well. I.e., any adversary A using reasonable resources will only be able to break the authenticity of CWC-BC-tl with some small probability.

Let us consider some concrete examples. Let $n = \text{MaxAdLen}$ and $m = \text{MaxMsgLen}$, which is the maximum possible allowed by the CWC-BC construction. Then Equation 1 becomes

$$\mathbf{Adv}_{\text{CWC-BC-tl}}^{\text{auth}}(A) \leq \mathbf{Adv}_{\text{BC}}^{\text{prp}}(C_A) + \frac{(\mu/128 + 3q + 1)^2}{2^{129}} + \frac{1}{2^{93}} + \frac{1}{2^{\text{tl}}}.$$

If we limit the number of applications of CWC-BC between rekeyings to some reasonable value such as $q = 2^{32}$, if we limit the total number of payload bits between rekeyings to $\mu = 2^{50}$, and if we take $\text{tl} \geq 43$, then the above equation becomes

$$\mathbf{Adv}_{\text{CWC-BC-tl}}^{\text{auth}}(A) \leq \mathbf{Adv}_{\text{BC}}^{\text{prp}}(C_A) + \frac{1}{2^{41}}$$

which means that, assuming that the underlying block cipher is a secure PRP, an attacker will not be able to break the unforgeability of CWC-BC-tl with probability much greater than 2^{-41} .

5.3 Privacy

Theorem 5.2 [Privacy of CWC.] Let CWC-BC-tl be as in Section 3. Then given a nonce-respecting IND \mathcal{S} -CPA adversary A against CWC-BC-tl one can construct a PRP adversary C_A against BC such that if A makes at most q oracle queries totaling at most μ bits of payload message data, then

$$\mathbf{Adv}_{\text{CWC-BC-tl}}^{\text{priv}}(A) \leq \mathbf{Adv}_{\text{BC}}^{\text{prp}}(C_A) + \frac{(\mu/128 + 3q + 1)^2}{2^{129}}. \quad (2)$$

Furthermore, the experiment for C_A takes the same time as the experiment for A and C_A makes at most $\mu/128 + 3q + 1$ oracle queries. ■

We interpret Theorem 5.2 in the same way we interpreted Theorem 5.1. In particular, this theorem shows that if BC is a secure pseudorandom permutation, then CWC-BC-tl preserves privacy under chosen-plaintext attacks.

As a concrete example of why Theorem 5.2 implies that CWC-BC preserves privacy under chosen-plaintext attacks, let us again consider the case where $q = 2^{32}$ and $\mu = 2^{50}$. Then Equation 2 becomes

$$\mathbf{Adv}_{\text{CWC-BC-tl}}^{\text{priv}}(A) \leq \mathbf{Adv}_{\text{BC}}^{\text{prp}}(C_A) + \frac{1}{2^{42}}$$

which means that, assuming that the underlying block cipher is a secure PRP, an attacker will not be able to break the privacy of CWC-BC-tl with advantage much greater than 2^{-42} .

Remark 5.3 [Chosen-ciphertext privacy.] Since CWC-BC-tl preserves privacy under chosen-plaintext attacks (Theorem 5.2) *and* provides integrity (Theorem 5.1) assuming that BC is a secure pseudorandom permutation, it also provides privacy under chosen-ciphertext attacks under the same assumption about BC. See [1, 16] for a discussion of the relationship between chosen-plaintext privacy, integrity, and chosen-ciphertext privacy; this relationship was also used, for example, by the designers of OCB [17].

5.4 Remarks

We close this section with some additional remarks on the design of CWC and several additional variants that we considered.

Remark 5.4 [On the length of the hash subkey.] It is possible to use smaller subkeys K_h in CWC-HASH (simply truncate $\text{BC}_K(110^{126})$ appropriately). Recall that we have fixed the block length of BC to 128 bits. Let hkl denote the length of the hash subkey in an altered construction. If $\text{hkl} < 127$, then the upper-bound in Equation 1 becomes

$$\text{Adv}_{\text{BC}}^{\text{prp}}(C_A) + \frac{(\mu/128 + 3q + 1)^2}{2^{129}} + \frac{(n + m)/96 + 2}{2^{\text{hkl}}} + \frac{1}{2^{\text{tl}}}.$$

Consider an application that sets hkl to 96. If we replace m and n by their maximum possible values, the upper-bound becomes

$$\text{Adv}_{\text{BC}}^{\text{prp}}(C_A) + \frac{(\mu/128 + 3q + 1)^2}{2^{129}} + \frac{1}{2^{62}} + \frac{1}{2^{\text{tl}}}.$$

Since 2^{-62} is already very small (and, in fact, dominated by the $(\mu/128 + 3q + 1)^2 \cdot 2^{-129}$ term for some reasonable values of q and μ), from a provable-security perspective, developers would be justified in using 96-bit hash subkeys.

Rather than use shorter hash subkeys, however, our current CWC instantiation in Section 3 uses 127-bit hash subkeys. We do so for several reasons. First, in hardware, to obtain maximum speed, one would parallelize the CWC hash function by evaluating, for example, two polynomials in K_h^2 in parallel. Since K_h^2 would generally not be 96-bits long, there is no performance advantage with using 96-bit subkeys K_h in this situation. In software, the use of 96-bit hash subkeys could lead to improved performance when evaluating the polynomial using Horner’s rule. However, the performance of such a construction is essentially equivalent to the performance of the current construct when not using Horner’s rule but using pre-computed powers of K_h . Since we believe that high-performance implementations will not benefit from the use of 96-bit hash subkeys (i.e., the additional 31 key bits come with no or negligible additional cost), we have chosen to fix the length of our hash subkeys to 127 bits.

Developers of CWC derivatives may, however, wish to use shorter hash subkeys, and we do not prevent that (although we do suggest referring to such modes in such a way as to avoid confusion with CWC-BC). We also suggest that developer’s understand the impact of using shorter hash subkeys. For example, using a 64-bit hash subkey would increase the upper-bound on the probability of an adversary forging to around 2^{-30} , which may be too large for some applications.

Remark 5.5 [On computing the tag.] In CWC the MAC consisted of hashing (A, σ) , enciphering the hash with the block cipher, and then XORing the result with some keystream (i.e., in the current proposal the tag is $\text{BC}_K(10^7 \| N \| 0^{32}) \oplus \text{BC}_K(\text{CWC-HASH}_K(A, \sigma))$). One question the reader

might have is whether two block cipher invocations are necessary. We first comment that the cost of two block cipher operations per MAC is not particularly significant compared to the total cost of CWC. CWC-AES as currently specified already achieves its design goal of encrypting 10 Gbps in hardware. And, in software, the extra cost of one block cipher operation is quite minor for average packets, and less than approximately 15% for 64-byte packets. Nevertheless, the use of two block cipher applications for the tag might seem aesthetically unappealing to some.

Instead of the two block cipher applications, one could use $\text{BC}_K(h'_K(N, A, \sigma))$ as the tag, where h' is a modified version of CWC-HASH designed to hash 3-tuples instead of pairs of strings (this is important because the nonce must also be authenticated). The main disadvantage of this approach is that it would change the upper-bound in Equation 1 to

$$\text{Adv}_{\text{BC}}^{\text{PRP}}(C_A) + \frac{(\mu/128 + 3q + 1)^2}{2^{129}} + q^2 \cdot \left(\frac{n + m}{2^{133}} + \frac{1}{2^{125}} \right) + \frac{1}{2^{\text{tl}}}$$

(note the new q^2 term). If we set $n = \text{MaxAdLen}$, $m = \text{MaxMsgLen}$, $q = 2^{32}$, and $\mu = 2^{50}$, then for any $\text{tl} \geq 29$, we get that the advantage of an adversary in breaking the unforgeability of this modified CWC variant is upper-bounded by 2^{-27} , which, although not extremely large, is worse than the upper-bound of 2^{-41} we get using Equation 1. Even if n and m are at most one million blocks long, we see that the integrity upper-bound for the altered CWC construction is worse than the upper-bound for the CWC construction we present in Section 3. More generally, this means that for reasonable values of n, m, q, μ , the insecurity upper-bounds of this alternative will be worse than the insecurity upper-bounds of the CWC mode described in Section 3. Furthermore, the upper-bound would be even worse if one keys the hash function with shorter keys, which some developers might choose to do (recall Remark 5.4).

Another possible way to reduce the number of block cipher invocations necessary to compute the MAC would be to take the output of the current hash function and run it through another hash function that is almost-XOR-universal (see Appendix C for a description of this property). However, this approach is not attractive because it requires additional key material. In particular, while this approach may save one block cipher operation, in hardware the block cipher operation is actually smaller and simpler than managing the extra key material, given that the hardware already has a block cipher encryptor running at high speed.

Another possibility would be to use something like $\text{BC}_K(N) + Y_1 K_h^{\beta+2} + \dots + Y_\beta K_h^3 + l_A K_h^2 + l_\sigma K_h \bmod 2^{127} - 1$, encoded as a 127-bit string and truncated to tl bits, as the MAC (here $\text{BC}_K(N)$ is interpreted as an integer). Doing so would, however, result in a new integrity upper-bound

$$\text{Adv}_{\text{BC}}^{\text{PRP}}(C_A) + \frac{(\mu/128 + 2q + 1)^2 + 4q + 4}{2^{129}} + \frac{(n + m)/96 + 5}{2^{\text{tl}}}.$$

If we take n and m to be MaxAdLen and MaxMsgLen , respectively, then the upper-bound becomes

$$\text{Adv}_{\text{BC}}^{\text{PRP}}(C_A) + \frac{(\mu/128 + 2q + 1)^2 + 4q + 4}{2^{129}} + \frac{2^{34}}{2^{\text{tl}}}.$$

Compared to Equation 1, we see the presence of a $2^{34-\text{tl}}$ term. This means that, in some situations, when using the above upper-bound as a guide for parameter selection, tag lengths must be longer than one might expect. For example, if $\text{tl} = 32$, then the above equation would upper-bound the advantage of an adversary against this modified construction as 1. This means that 32-bit tags should not be used with this modified construction when authenticating long messages. While one might consider this more of a “certificational” problem than a real problem, we view this property as undesirable. Hence our decision to specify CWC as in Section 3.

6 Conclusions

In this work we present CWC, the first patent-free, parallelizable, and provably-secure dedicated block cipher mode of operation. Because of its inherent parallelism, CWC-AES is capable of processing data at 10 Gbps in hardware, making it ideal for use with coming 10 Gbps IPsec network devices. CWC-AES is also efficient in software, with the current implementation comparable to current implementations of the other patent-free (albeit not parallelizable) modes of operations CCM-AES and EAX-AES. In software, we anticipate significant speedups after switching to assembly and using the precomputation approach for CWC-HASH discussed in Section 4. Finally, CWC-AES is provably secure assuming that AES is a secure pseudorandom permutation, which is a reasonable assumption and, in fact, was one of the AES design criteria.

Acknowledgments

We thank Peter Gutmann, David McGrew, Fabian Monrose, Avi Rubin, Adam Stubblefield, and David Wagner for their comments. Additionally, we thank Brian Gladman for helping to validate our test vectors and for working with us to obtain timing information. T. Kohno was supported by a National Defense Science and Engineering Fellowship.

References

- [1] M. Bellare and C. Namprempre. Authenticated encryption: Relations among notions and analysis of the generic composition paradigm. In T. Okamoto, editor, *Advances in Cryptology – ASIACRYPT 2000*, volume 1976 of *Lecture Notes in Computer Science*, pages 531–545. Springer-Verlag, Berlin Germany, Dec. 2000.
- [2] M. Bellare, P. Rogaway, and D. Wagner. A conventional authenticated-encryption mode, 2003. Available at <http://eprint.iacr.org/2003/069/>.
- [3] D. Bernstein. Floating-point arithmetic and message authentication, 2000. Available at <http://cr.yp.to/papers.html#hash127>.
- [4] B. Gladman. AES and combined encryption/authentication modes, 2003. Available at <http://fp.gladman.plus.com/AES/index.htm>.
- [5] V. Gligor and P. Donescu. Fast encryption and authentication: XCBC encryption and XECB authentication modes. In *Fast Software Encryption 2001*, Lecture Notes in Computer Science. Springer-Verlag, Berlin Germany, 2001.
- [6] C. Hecker. Perspective texture mapping, part V: It’s about time. *Game Developer*, Apr. 1996. Available at <http://www.d6.com/users/checker/pdfs/gdmtex5.pdf>.
- [7] C. Jutla. Encryption modes with almost free message integrity. In B. Pfitzmann, editor, *Advances in Cryptology – EUROCRYPT 2001*, volume 2045 of *Lecture Notes in Computer Science*, pages 529–544. Springer-Verlag, Berlin Germany, May 2001.
- [8] J. Katz and M. Yung. Unforgeable encryption and chosen ciphertext secure modes of operation. In B. Schneier, editor, *Fast Software Encryption 2000*, volume 1978 of *Lecture Notes in Computer Science*, pages 284–299. Springer-Verlag, Berlin Germany, Apr. 2000.

- [9] H. Krawczyk. LFSR-based hashing and authentication. In Y. Desmedt, editor, *Advances in Cryptology – CRYPTO ’94*, Lecture Notes in Computer Science. Springer-Verlag, Berlin Germany, Aug. 1994.
- [10] H. Krawczyk. The order of encryption and authentication for protecting communications (or: How secure is SSL?). In J. Kilian, editor, *Advances in Cryptology – CRYPTO 2001*, volume 2139 of *Lecture Notes in Computer Science*, pages 310–331. Springer-Verlag, Berlin Germany, Aug. 2001.
- [11] H. Lipmaa. AES/Rijndael: speed, 2003. Available at <http://www.tcs.hut.fi/~helger/aes/rijndael.html>.
- [12] D. McGrew. Integer counter mode, Oct. 2002. Available at <http://www.ietf.org/internet-drafts/draft-irtf-cfrg-icm-00.txt>.
- [13] D. McGrew. The truncated multi-modular hash function (TMMH), version two, Oct. 2002. Available at <http://www.ietf.org/internet-drafts/draft-irtf-cfrg-tmmh-00.txt>.
- [14] D. McGrew. The universal security transform, Oct. 2002. Available at <http://www.ietf.org/internet-drafts/draft-irtf-cfrg-ust-00.txt>.
- [15] P. Rogaway. Bucket hashing and its applications to fast message authentication. *Journal of Cryptology*, 12:91–115, 1999.
- [16] P. Rogaway. Authenticated encryption with associated data. In *Proceedings of the 9th Conference on Computer and Communications Security*, Nov. 2002.
- [17] P. Rogaway, M. Bellare, J. Black, and T. Krovetz. OCB: A block-cipher mode of operation for efficient authenticated encryption. In *Proceedings of the 8th Conference on Computer and Communications Security*, pages 196–205. ACM Press, 2001.
- [18] P. Rogaway and D. Wagner. A critique of CCM, Apr. 2003. Available at <http://eprint.iacr.org/2003/070/>.
- [19] D. Stinson. Universal hashing and authentication codes. *Designs, Codes and Cryptography*, 4:369–380, 1994.
- [20] M. Wegman and L. Carter. New hash functions and their use in authentication and set equality. *Journal of Computer and System Sciences*, 22:265–279, 1981.
- [21] D. Whiting, N. Ferguson, and R. Housley. Counter with CBC-MAC (CCM). Submission to NIST. Available at <http://csrc.nist.gov/CryptoToolkit/modes/proposedmodes/>, 2002.

A Intellectual property statement

The authors hereby explicitly release any intellectual property rights to the CWC mode into the public domain. The authors are not aware of any patent or patent application anywhere in the world that cover this mode.

B Summary of properties

In this appendix we summarize some of the properties of CWC. We include all of the properties listed in the submission guidelines on the NIST Modes of Operation website. We also discuss some additional properties that we feel are important.

SECURITY FUNCTION. CWC is a provably secure *authenticated encryption* with associated data (AEAD) mode. Informally, this means that the encapsulation algorithm, on input a pair of messages (A, M) and some nonce N , encapsulates (A, M) in a way that protects the privacy of M and the integrity of both A and M . Our formal security statements appear in Section 5 and the proofs appear in Appendix C.

ERROR PROPAGATION. Assuming that the underlying block cipher is a secure pseudorandom function or permutation, any attempt, by an adversary using reasonable resources, to forge a new ciphertext will, with very high probability, be detected. This follows from the fact that CWC is a provably-secure AEAD mode.

SYNCHRONIZATION. Synchronization is based on the nonce. As with other nonce-based AEAD modes, the nonce must either be sent with the ciphertext or the receiver must know how to derive the nonce on its own.

PARALLELIZABILITY. CWC is *parallelizable*. The amount of parallelism for the hashing portion can be determined by the implementor without affecting interoperability.

KEYING MATERIAL REQUIRED. CWC is defined to be a *single-key* AEAD mode. However, CWC does internally use two keys (the main block cipher key and a hash key which is derived using the block cipher key). Implementors can decide whether to store the derived hash key in memory or whether to re-derive it as needed.

COUNTER/IV/NONCE REQUIREMENTS. CWC uses a 11-octet *nonce*. CWC is provably secure as long as one does not query the encryption algorithm twice with the same nonce. Although it is possible to instantiate the generic CWC paradigm with other nonce lengths, for CWC the nonce size is fixed at 11-octets in order to minimize interoperability issues.

MEMORY REQUIREMENTS. The software memory requirements are basically those of the underlying block cipher. For example, fast AES in software requires 4K bytes of table, and about 200 bytes of expanded key material. In some situations, software implementations may precompute powers of the hash subkey.

PRE-PROCESSING CAPABILITY. The underlying CTR mode keystream can be *precomputed*. The only block cipher input that cannot be precomputed is the output of CWC-HASH.

CWC can preprocess its associated data, thereby reducing computation time if the associated data remains static or changes only infrequently.

MESSAGE LENGTH REQUIREMENTS. The associated data and message can both be any string of octets with length at most $128 \cdot (2^{32} - 1)$ bits. Because there does not appear to be a need to handle strings of arbitrary bit-length, CWC as currently specified cannot encapsulate arbitrary bit-length messages. (As discussed in Section 3, it is easy to modify CWC to handle arbitrary bit-length messages, if desired.)

CIPHERTEXT EXPANSION. The ciphertext expansion is the minimum possible while still providing a tl -bit tag. That is, on input a pair (A, M) , a nonce N , and a key K , $\text{CWC-ENC}_K(N, A, M)$ outputs a ciphertext C with length $|C| = |M| + tl$.

BLOCK CIPHER INVOCATIONS. If the hash subkey K_h is computed as part of the key generation

process and not during each invocation of the CWC encapsulation routine, then CWC makes one block cipher invocation during key setup and $\lceil |M|/128 \rceil + 2$ block cipher invocations during encapsulation and decapsulation. If the hash subkey K_h is not computed as part of the key generation process, then CWC makes no block cipher invocations during key setup and $\lceil |M|/128 \rceil + 3$ block cipher invocations during encapsulation and decapsulation.

PROVABLE SECURITY. CWC is a *provably-secure* AEAD mode assuming that the underlying block cipher (e.g., AES) is a secure pseudorandom function or permutation. The proofs of security do not require the block cipher to satisfy the strong notion of super-pseudorandomness required by some other block cipher modes of operation.

NUMBER OF OPTIONS AND INTEROPERABILITY. CWC uses a minimal number of options. The only options are the choice of the underlying block cipher (and key length) and the tag length. Having fewer options makes interoperability easier.

ON-LINE. The CWC encryption algorithm is on-line. This means that CWC can process data as it arrives, rather than waiting for the entire message to be buffered before beginning the encryption processes. This may be advantageous when encrypting streaming data sources. (Note, however, that, like any other AEAD mode, the decryptor should still buffer the entire message and check the tag τ before revealing the plaintext and associated data.)

PATENT STATUS. To the best of our knowledge CWC is *not* covered by any patents.

PERFORMANCE. CWC is efficient in both hardware and software. In hardware, CWC can process data at 10 Gbps.

SIMPLICITY. Although simplicity is a matter of perspective, we believe that CWC is a very simple construction. It combines standard CTR mode encryption with the evaluation of a polynomial modulo $2^{127}-1$. Because of its simplicity, we believe that CWC is easy to implement and understand.

C Proofs of Theorem 5.1 and Theorem 5.2

Before proving Theorem 5.1 and Theorem 5.2, we first state results about the general CWC paradigm (see Lemma C.5 and Lemma C.6 below). We then show how Theorems 5.1 and 5.2 follow from Lemmas C.5 and C.6. We then prove these two lemmas.

C.1 More definitions

We begin with a few additional definitions.

UNIVERSAL HASH FUNCTIONS. A hash function $\mathcal{HF} = (\mathcal{K}_h, \mathcal{H})$ consists of two algorithms and is defined over some key space $\text{KeySp}_{\mathcal{HF}}$, some message space $\text{MsgSp}_{\mathcal{HF}}$, and some hash space $\text{HashSp}_{\mathcal{HF}}$. The randomized key generation algorithm returns a random key $K \in \text{KeySp}_{\mathcal{HF}}$; we denote this as $K \stackrel{\$}{\leftarrow} \mathcal{K}_h$. The deterministic hash algorithm takes a key $K \in \text{KeySp}_{\mathcal{HF}}$ and a message $M \in \text{MsgSp}_{\mathcal{HF}}$ and returns a hash value $h \in \text{HashSp}_{\mathcal{HF}}$; we denote this as $h \leftarrow \mathcal{H}_K(M)$. Let $H \stackrel{\$}{\leftarrow} \mathcal{HF}$ be shorthand for $K \stackrel{\$}{\leftarrow} \mathcal{K}_h; H \leftarrow \mathcal{H}_K$.

The hash function \mathcal{HF} is said to be ϵ -almost universal (ϵ -AU) if for all distinct $m, m' \in \text{MsgSp}_{\mathcal{HF}}$,

$$\Pr \left[H \stackrel{\$}{\leftarrow} \mathcal{HF} : H(m) = H(m') \right] \leq \epsilon .$$

The hash function \mathcal{HF} is said to be ϵ -almost XOR universal (ϵ -AXU) if $\text{HashSp}_{\mathcal{HF}} = \{0, 1\}^n$ for some positive integer n and for all distinct $m, m' \in \text{MsgSp}_{\mathcal{HF}}$ and $c \in \{0, 1\}^n$,

$$\Pr \left[H \stackrel{\$}{\leftarrow} \mathcal{HF} : H(m) \oplus H(m') = c \right] \leq \epsilon .$$

PSEUDORANDOM FUNCTIONS. If X and Y are sets, then $\text{Func}(X, Y)$ denotes the set of all functions from X to Y . If l and L are positive integers, then $\text{Func}(l, L)$ denotes the set of all functions from $\{0, 1\}^l$ to $\{0, 1\}^L$.

Let F be a family of functions from D to R . Let A be an adversary with access to an oracle and that returns a bit. Then

$$\text{Adv}_F^{\text{prf}}(A) = \Pr \left[f \stackrel{\$}{\leftarrow} F : A^{f(\cdot)} = 1 \right] - \Pr \left[g \stackrel{\$}{\leftarrow} \text{Func}(D, R) : A^{g(\cdot)} = 1 \right]$$

denotes the PRF-advantage of A in distinguishing a random instance of F from a random function. Intuitively, we say that F is a secure PRF if the PRF-advantages of all adversaries using reasonable resources is small.

MESSAGE AUTHENTICATION. A nonced message authentication scheme $\mathcal{MA} = (\mathcal{K}_m, \mathcal{T}, \mathcal{V})$ consists of three algorithms and is defined over some key space $\text{KeySp}_{\mathcal{MA}}$, some nonce space $\text{NonceSp}_{\mathcal{MA}}$, some message space $\text{MsgSp}_{\mathcal{MA}}$, and some tag space $\text{TagSp}_{\mathcal{MA}}$. The randomized key generation algorithm returns a key $K \in \text{KeySp}_{\mathcal{MA}}$; we denote this as $K \stackrel{\$}{\leftarrow} \mathcal{K}_m$. The deterministic tagging algorithm \mathcal{T} takes a key $K \in \text{KeySp}_{\mathcal{MA}}$, a nonce $N \in \text{NonceSp}_{\mathcal{MA}}$, and a message $M \in \text{MsgSp}_{\mathcal{MA}}$ and returns a tag $\tau \in \text{TagSp}_{\mathcal{MA}}$; we denote this process as $\tau \leftarrow \mathcal{T}_K^N(M)$ or $\tau \leftarrow \mathcal{T}_K(N, M)$. The deterministic verification algorithm \mathcal{V} takes as input a key $K \in \text{KeySp}_{\mathcal{MA}}$, a nonce $N \in \text{NonceSp}_{\mathcal{MA}}$, a message $M \in \text{MsgSp}_{\mathcal{MA}}$, and a candidate tag $\tau \in \{0, 1\}^*$, computes $\tau' = \mathcal{T}_K^N(M)$, and returns `accept` if $\tau' = \tau$ and returns `reject` otherwise.

Let F be a forging adversary and consider an experiment in which we first pick a random key $K \stackrel{\$}{\leftarrow} \mathcal{K}_m$ and then run F with oracle access to $\mathcal{T}_K(\cdot, \cdot)$. We say that F *forges* if F returns a triple (N, M, τ) such that $\mathcal{V}_K^N(M, \tau) = \text{accept}$ but F did not make a query (N, M) to $\mathcal{T}_K(\cdot, \cdot)$ that resulted in a response τ . Then

$$\text{Adv}_{\mathcal{MA}}^{\text{uf}}(F) = \Pr \left[K \stackrel{\$}{\leftarrow} \mathcal{K}_m : F^{\mathcal{T}_K(\cdot, \cdot)} \text{ forges} \right]$$

denotes the UF-advantage of F in breaking the *unforgeability* of \mathcal{MA} . An adversary is *nonce-respecting* if it never queries its tagging oracle with the same nonce twice. Intuitively, \mathcal{MA} is unforgeable if the UF-advantage of all nonce-respecting adversaries with reasonable resources is small.

C.2 The general CWC construction

We now describe our generalization of the CWC construction.

Construction C.1 [General CWC.] Let l, L, n, o, t, k be positive integers such that $t \leq L$. (Further restrictions will be placed shortly.) Essentially, l is the length of the input to a PRF (e.g., 128), L is the length of the output from the PRF (e.g., 128), n is the length of the nonce (e.g., 88), o is the length of the offset (e.g., 32), t is the length of the desired tag (e.g., 64 or 128), k is the length of the hash function's keysize (e.g., 127).

Let F be a family of functions from $\{0, 1\}^l$ to $\{0, 1\}^L$. Let $\mathcal{HF} = (\mathcal{K}_h, \mathcal{H})$ be a family of hash functions with $\text{HashSp}_{\mathcal{HF}} = \{0, 1\}^l$ and $\text{KeySp}_{\mathcal{HF}} = \{0, 1\}^k$ (and \mathcal{K}_h works by randomly selecting and returning an element from $\{0, 1\}^k$ with uniform probability). Let $\text{ctr0} : \mathbb{Z}_{\lceil k/L \rceil} \rightarrow \{0, 1\}^l$, $\text{ctr1} : \{0, 1\}^n \times (\mathbb{Z}_{2^o} - \{0\}) \rightarrow \{0, 1\}^l$ and $\text{ctr2} : \{0, 1\}^n \rightarrow \{0, 1\}^l$ be efficiently-computable injective functions. If $W = \{ \text{ctr0}(O) : O \in \mathbb{Z}_{\lceil k/L \rceil} \}$, $X = \{ \text{ctr1}(N, O) : N \in \{0, 1\}^n, O \in (\mathbb{Z}_{2^o} - \{0\}) \}$, $Y = \{ \text{ctr2}(N) : N \in \{0, 1\}^n \}$, and $Z = \{ \mathcal{H}_K(M) : K \in \text{KeySp}_{\mathcal{HF}}, M \in \text{MsgSp}_{\mathcal{HF}} \}$, we require that W, X, Y , and Z be pairwise mutually exclusive.

Let $\text{extract} : \{0, 1\}^{\lceil k/L \rceil \cdot L} \rightarrow \{0, 1\}^k$ be a function that takes as input a $\lceil k/L \rceil \cdot L$ -bit string and that outputs a k -bit string. We require that extract always pick the same k bits from the input

string and always outputs those bits in the exact same order (e.g., `extract` returns the first k bits of its input).

Let $\mathcal{SE}[F, \mathcal{HF}] = (\mathcal{K}_e, \mathcal{E}, \mathcal{D})$ be an AEAD mode built from function family F and hash function \mathcal{HF} and using the above functions `extract`, `ctr0`, `ctr1`, `ctr2`. We assume that $\text{AdSp}_{\mathcal{SE}[F, \mathcal{HF}]} \times \text{MsgSp}_{\mathcal{SE}[F, \mathcal{HF}]} \subseteq \text{MsgSp}_{\mathcal{HF}}$ and that all messages in $\text{MsgSp}_{\mathcal{SE}[F, \mathcal{HF}]}$ have length at most $L \cdot (2^o - 1)$. Note that the former means that the message space of \mathcal{HF} actually consists of pairs of strings. Let $\text{NonceSp}_{\mathcal{SE}[F, \mathcal{HF}]} = \{0, 1\}^n$. Let $\mathcal{SE}[F, \mathcal{HF}]$'s component algorithms be defined as follows:

Algorithm \mathcal{K}_e

```

 $f \xleftarrow{\$} F$ 

$$K_h \leftarrow \text{extract}(f(\text{ctr0}(0)) \| f(\text{ctr0}(1)) \| \dots \| f(\text{ctr0}(\lceil k/L \rceil - 1))) ; H \leftarrow \mathcal{H}_{K_h}$$

Return  $\langle f, H \rangle$ 

```

Algorithm $\mathcal{E}_{\langle f, H \rangle}^{N, A}(M)$

```

 $\sigma \leftarrow \text{CTR-MODE}_f^N(M)$ 

$$\tau \leftarrow \text{first } t \text{ bits of } (f(\text{ctr2}(N)) \oplus f(H(A, \sigma)))$$

Return  $\sigma \| \tau$ 

```

Algorithm $\mathcal{D}_{\langle f, H \rangle}^{N, A}(C)$

```

If  $|C| < t$  then return INVALID
Parse  $C$  as  $\sigma \| \tau$  //  $|\tau| = t$ 
If  $A \notin \text{AdSp}_{\mathcal{SE}[F, \mathcal{HF}]}$  or  $\sigma \notin \text{MsgSp}_{\mathcal{SE}[F, \mathcal{HF}]}$  then return INVALID

$$\tau' \leftarrow \text{first } t \text{ bits of } (f(\text{ctr2}(N)) \oplus f(H(A, \sigma)))$$

If  $\tau \neq \tau'$  return INVALID
 $M \leftarrow \text{CTR-MODE}_f^N(\sigma)$ 
Return  $M$ 

```

Algorithm $\text{CTR-MODE}_f^N(X)$

```

 $\alpha \leftarrow \lceil |X|/L \rceil$ 
For  $i = 1$  to  $\alpha$  do
     $Z_i \leftarrow f(\text{ctr1}(N, i))$ 
 $Y \leftarrow (\text{first } |X| \text{ bits of } Z_1 \| Z_2 \| \dots \| Z_\alpha) \oplus X$ 
Return  $Y$  ■

```

Remark C.2 Recall that one requirement on the message space for any AEAD mode is that if it contains any string M , then it contains all strings of length $|M|$. This means that the membership test $\sigma \notin \text{MsgSp}_{\mathcal{SE}[F, \mathcal{HF}]}$ and the application of H to (A, σ) makes sense.

Remark C.3 As specified in the definition, $\text{AdSp}_{\mathcal{SE}[F, \mathcal{HF}]} \times \text{MsgSp}_{\mathcal{SE}[F, \mathcal{HF}]} \subseteq \text{MsgSp}_{\mathcal{HF}}$. This means that we \mathcal{HF} is used to hash *pairs* of strings, not just string. This is not a serious restriction since given any hash function that hashes strings, it is trivial to construct a hash function that hashes pairs of strings (by encoding the pair of strings as a single string in some appropriate manner).

Remark C.4 It is also worth commenting on the purpose of `ctr0`, `ctr1`, and `ctr2`. As shown in Construction C.1, these functions are used to derive the inputs to the construction's underlying function f . By requiring that none of the outputs collide (i.e., that the sets W, X, Y, Z in the definition are pairwise mutually exclusive), we ensure that the inputs to f for different purposes never collide. For example, the inputs to f used for counter mode encryption will always be different than the inputs to f when enciphering the output of H .

C.3 The security of the general CWC construction

We now state the following results for all Construction C.1-style AEAD modes. We shall prove Lemmas C.5 and C.6 in Appendices C.5 and C.6, respectively.

Lemma C.5 [Integrity of Construction C.1.] Let $\mathcal{SE}[F, \mathcal{HF}]$ be as in Construction C.1 and let \mathcal{HF} be an ϵ -AU hash function. Then given any nonce-respecting AUTH adversary A against $\mathcal{SE}[F, \mathcal{HF}]$, we can construct a PRF adversary B_A against F such that

$$\mathbf{Adv}_{\mathcal{SE}[F, \mathcal{HF}]}^{\text{auth}}(A) \leq \mathbf{Adv}_F^{\text{prf}}(B_A) + \epsilon + 2^{-t}.$$

Furthermore, the experiment for B_A takes the same time as the experiment for A and, if A makes at most $q - 1$ oracle queries and a total of at most μ bits of payload data (for both these $q - 1$ oracle queries and the forgery attempt), then B_A makes at most $\mu/L + 3q + \lceil k/L \rceil$ oracle queries. ■

Lemma C.6 [Privacy of Construction C.1.] Let $\mathcal{SE}[F, \mathcal{HF}]$ be as in Construction C.1. Then given a nonce-respecting IND $\$$ -CPA adversary A against $\mathcal{SE}[F, \mathcal{HF}]$ one can construct a PRF adversary B_A against F such that

$$\mathbf{Adv}_{\mathcal{SE}[F, \mathcal{HF}]}^{\text{priv}}(A) \leq \mathbf{Adv}_F^{\text{prf}}(B_A).$$

Furthermore, the experiment for B_A takes the same time as the experiment for A and, if A makes at most q oracle queries totaling at most μ bits of payload data, then B_A makes at most $\mu/L + 3q + \lceil k/L \rceil$ oracle queries. ■

We interpret these lemmas as follows. Intuitively, the first lemma states that if F is a secure PRF, if \mathcal{HF} is ϵ -AU where ϵ is not too large, and if t is not too small, then $\mathcal{SE}[F, \mathcal{HF}]$ preserves integrity. We comment that most modern block ciphers (e.g., AES) are considered to be secure PRPs (and therefore also secure PRFs up to a birthday term). We also comment that we can construct hash functions \mathcal{HF} with provably small ϵ .

Intuitively, the second lemma states that if F is a secure PRF, then $\mathcal{SE}[F, \mathcal{HF}]$ will preserve privacy. We discuss the meaning of these types of proofs in more detail in Section 5.

C.4 Proof of Theorem 5.1 and Theorem 5.2

The security of the CWC construction from Section 3 follows from Lemmas C.5 and C.6 assuming that (1) CWC as described in Section 3 is really an instantiation of Construction C.1 and (2) that the hash function used in Section 3 is ϵ -AU for some small ϵ . We begin by justifying the second bullet.

Lemma C.7 [CWC-HASH (Section 3) is ϵ -almost universal.] Consider the CWC-BC-tl construction from Section 3. Let $\mathcal{HF} = (\mathcal{K}_h, \mathcal{H})$ be the hash function whose key generation algorithm selects a random key K from $\{0, 1\}^{127}$ and let \mathcal{H}_K be the CWC-HASH function except that we replace

$$Z \leftarrow \text{last 127 bits of } \text{BC}_K(110^{126})$$

with

$$Z \leftarrow K.$$

Note that $\text{AdSp}_{\text{CWC-BC-tl}} \times \text{MsgSp}_{\text{CWC-BC-tl}} \subseteq \text{MsgSp}_{\mathcal{HF}}$; that is, \mathcal{H}_K takes two strings as input. Assume \mathcal{HF} hashes pairs of strings where the first string is always at most $n \leq \text{MaxAdLen}$ bits long

and the second string is always at most $m \leq \text{MaxMsgLen}$ bits long. Then \mathcal{HF} is ϵ -almost universal where

$$\epsilon \leq \frac{n+m}{2^{133}} + \frac{1}{2^{125}} \cdot \blacksquare$$

Proof of Lemma C.7: Let (A, σ) and (A', σ') be two distinct inputs to \mathcal{H}_K and let $X = (B_1, \dots, B_{\beta+1})$ and $Y = (C_1, \dots, C_{\gamma+1})$ respectively denote their encodings as vectors of 96-bit integers (with $B_{\beta+1}$ and $C_{\gamma+1}$ possibly longer than 96-bits long). Without loss of generality, assume $\beta \leq \gamma$ and let $X' = (B'_1, \dots, B'_{\gamma+1})$ where $B'_j = 0$ for $j \in \{1, \dots, \gamma - \beta\}$ and $B'_j = B_{j-\gamma+\beta}$ for $j \in \{\gamma - \beta + 1, \dots, \gamma + 1\}$ (i.e., prepend $\gamma - \beta$ zero elements to the X vector).

If $(A, \sigma) \neq (A', \sigma')$ then $X' \neq Y$. This follows from the fact that $B'_{\gamma+1}$ and $C_{\gamma+1}$ respectively encode the lengths of A and σ and of A' and σ' and that if $X' = Y$, then the $B'_{\gamma+1} = C_{\gamma+1}$ and $(A, \sigma) = (A', \sigma')$.

Note that $\mathcal{H}_K(A, \sigma) = \mathcal{H}_K(A', \sigma')$ when

$$\left(B'_1 \cdot K_h^\gamma + \dots + B'_\gamma \cdot K_h + B'_{\gamma+1} \right) - \left(C_1 \cdot K_h^\gamma + \dots + C_\gamma \cdot K_h + C_{\gamma+1} \right) = 0 \pmod{2^{127} - 1} \quad (3)$$

where K_h is the hash key derived from K as specified in CWC-HASH. Since the vectors X' and Y are not equal, $\left(B'_1 \cdot K_h^\gamma + \dots + B'_\gamma \cdot K_h + B'_{\gamma+1} \right) - \left(C_1 \cdot K_h^\gamma + \dots + C_\gamma \cdot K_h + C_{\gamma+1} \right)$ is a non-zero polynomial of degree at most γ . Therefore, by the Fundamental Theorem of Algebra, Equation 3 has at most γ solution modulo $2^{127} - 1$.

Since we are interested in the probability, over the 127-bit keys K , that Equation 3 is true, we note that all keys K_h modulo $2^{127} - 1$ (except 0) have exactly one ways of occurring and that the 0 key can occur in one additional way (i.e., the all 0 string and the all 1 string). This means that of the 2^{127} possible keys K , at most $\gamma + 1$ can lead to keys K_h such that Equation 3 is true.

Finally, note that γ is at most $2 + (n+m)/96$ (the +2 comes from the fact that we append 0 bits to A and σ). Consequently

$$\epsilon \leq \frac{\frac{n+m}{96} + 3}{2^{127}} \leq \frac{n+m}{2^{133}} + \frac{1}{2^{125}}$$

as desired. \blacksquare

We now prove Theorem 5.1 and Theorem 5.2, which are corollaries of Lemmas C.5, C.6, and C.7.

Proof of Theorem 5.1 and Theorem 5.2: To prove these theorems we must show that the CWC-BC-tl constructions from Section 3 are instantiations of Construction C.1. We begin by noting that the block cipher BC in CWC-BC-tl plays the role of F in Construction C.1 and that the hash function CWC-HASH (with the simplified key generation algorithm from Lemma C.7) plays the role of \mathcal{HF} in Construction C.1.

Since BC plays the role of F , we have that $l = L = 128$. Furthermore, as described in Section 3, $n = 88$, $o = 32$, $t = \text{tl}$, and $k = 127$. We note that the output the hash function is a 128-bit string whose first bit is always 0. This property, as well as the encodings for the nonce/offsets when encrypting the message and the Carter-Wegman MAC and when generating the hash key, ensure that requisite properties for the interactions between the hash function, ctr0 , ctr1 , and ctr2 .

A direct comparison of the Construction C.1 algorithms and the algorithms from Section 3 shows that they are equivalent. CWC-BC-tl is therefore an instantiation of Construction C.1 and the provable security of CWC-BC-tl follows.

Finally, we apply the standard PRF-PRP switching technique in order to model the underlying block cipher as a PRP rather than a PRF in Theorem 5.1 and Theorem 5.2. \blacksquare

C.5 Proof of Lemma C.5

We begin by sketching the proof of Lemma C.5. We first show that applying a random function to the output of an ϵ -AU hash function yields an ϵ' -AXU hash function (Proposition C.9). We then recall the result of Krawczyk [9] that XORing the output of an AXU hash function with a one-time pad yields a secure MAC (Proposition C.11). Such a MAC essentially corresponds to the second and third boxed steps in Construction C.1. (We do not need this final block cipher application if the input to the hash includes the nonce and if we accept a birthday term of the form $q^2\epsilon$.)

We then observe that if we consider a construction like Construction C.1 but with the latter two boxed steps replaced with calls to a secure MAC that tags pairs of strings (A, σ) with nonces N , then that construction would be unforgeable (Proposition C.13). In Proposition C.16 we use the above results to show that $\mathcal{SE}[\text{Func}(l, L), \mathcal{HF}]$ preserves integrity (where $\mathcal{SE}[\text{Func}(l, L), \mathcal{HF}]$ is as in Construction C.1). Lemma C.5 follows.

FROM AU TO AXU. Let us begin with the following construction.

Construction C.8 [Building AXU hash functions from AU hash functions.] Let $\mathcal{HF} = (\mathcal{K}_h, \mathcal{H})$ be a hash function and let $\overline{\mathcal{HF}[t]} = (\overline{\mathcal{K}_h}, \overline{\mathcal{H}})$, t a positive integer, be the hash function defined as follows:

$$\begin{array}{l} \overline{\mathcal{K}_h} \\ H \stackrel{\$}{\leftarrow} \mathcal{HF} \\ e \stackrel{\$}{\leftarrow} \text{Func}(\text{HashSp}_{\mathcal{HF}}, \{0, 1\}^t) \\ \text{Return } \langle H, e \rangle \end{array} \left| \begin{array}{l} \overline{\mathcal{H}}_{\langle H, e \rangle}(M) \\ \text{Return } e(H(M)) \end{array} \right.$$

Note that $\text{MsgSp}_{\overline{\mathcal{HF}[t]}} = \text{MsgSp}_{\mathcal{HF}}$ and $\text{HashSp}_{\overline{\mathcal{HF}[t]}} = \{0, 1\}^t$. \blacksquare

Proposition C.9 Let \mathcal{HF} , t , and $\overline{\mathcal{HF}[t]}$ be as in Construction C.8. If \mathcal{HF} is ϵ -AU, then $\overline{\mathcal{HF}[t]}$ is $(\epsilon + 2^{-t})$ -AXU.

This result follows from a result in [19, 15] which states that the composition of an ϵ' -AXU hash function, with domain B and range C , with an ϵ -AU hash function, with domain A and range B , is an $(\epsilon + \epsilon')$ -AXU hash function with domain A and range C , and the fact that the hash function whose key generation algorithm returns a random function from $\text{Func}(\text{HashSp}_{\mathcal{HF}}, \{0, 1\}^t)$ is 2^{-t} -AXU.

CARTER-WEGMAN MACS. Consider now the following construction.

Construction C.10 [Building MACs from AXU hash functions.] Let $\mathcal{HF} = (\mathcal{K}_h, \mathcal{H})$ be a hash function with hash space $\{0, 1\}^t$, t a positive integer. We can construct a nonced message authentication scheme $\mathcal{MA} = (\mathcal{K}_m, \mathcal{T}, \mathcal{V})$ as follows:

$$\begin{array}{l} \mathcal{K}_m \\ H \stackrel{\$}{\leftarrow} \mathcal{HF} \\ g \stackrel{\$}{\leftarrow} \text{Func}(\text{NonceSp}_{\mathcal{MA}}, \{0, 1\}^t) \\ \text{Return } \langle H, g \rangle \end{array} \left| \begin{array}{l} \mathcal{T}_{\langle H, g \rangle}(N, M) \\ \text{Return } g(N) \oplus H(M) \end{array} \right| \begin{array}{l} \mathcal{V}_{\langle H, g \rangle}(N, M, \tau) \\ \text{If } g(N) \oplus H(M) = \tau \text{ then} \\ \quad \text{return accept} \\ \text{Else return reject} \end{array}$$

Note that $\text{MsgSp}_{\mathcal{MA}} = \text{MsgSp}_{\mathcal{HF}}$, $\text{TagSp}_{\mathcal{MA}} = \{0, 1\}^t$, and that $\text{NonceSp}_{\mathcal{MA}}$ is arbitrary. \blacksquare

We now state the following result, due to Krawczyk [9].

Proposition C.11 Let \mathcal{HF} and \mathcal{MA} be as in Construction C.10. If \mathcal{HF} is ϵ -AXU, then for all nonce-respecting UF adversaries F attacking \mathcal{MA} , $\mathbf{Adv}_{\mathcal{MA}}^{\text{uf}}(F) \leq \epsilon$. **■**

As noted in [9], this proposition follows from the facts that XORing the output of the hash function with $g(N)$ prevents any loss of information (assuming that the adversary is nonce-respecting), that a forgery attempt with a previous nonce is upper-bounded by ϵ , and that a forgery attempt with a new nonce is upper-bounded by $2^{-t} \leq \epsilon$.

ENCRYPT-THEN-AUTHENTICATE. Consider the following Encrypt-then-Authenticate [1, 10] construction.

Construction C.12 [Encrypt-then-Authenticate.] Let l, L, n, o, t be positive integers. (Further restrictions will be placed shortly.) Essentially, l is the length of the input to a PRF (e.g., 128), L is the length of the output from the PRF (e.g., 128), n is the length of the nonce (e.g., 88), o is the length of the offset (e.g., 32).

Let F be a family of functions from $\{0, 1\}^l$ to $\{0, 1\}^L$. Let $\mathcal{MA} = (\mathcal{K}_m, \mathcal{T}, \mathcal{V})$ be a message authentication scheme with $\text{NonceSp}_{\mathcal{MA}} = \{0, 1\}^n$ and $\text{TagSp}_{\mathcal{MA}} = \{0, 1\}^t$. Let $\text{ctr1} : \{0, 1\}^n \times (\mathbb{Z}_{2^o} - \{0\}) \rightarrow \{0, 1\}^l$ be an efficiently-computable injective function.

Let $\mathcal{SE}[F, \mathcal{MA}] = (\mathcal{K}_e, \mathcal{E}, \mathcal{D})$ be an AEAD mode built from function family F and message authentication scheme \mathcal{MA} and using the above function ctr1 . We assume that $\text{AdSp}_{\mathcal{SE}[F, \mathcal{MA}]} \times \text{MsgSp}_{\mathcal{SE}[F, \mathcal{MA}]} \subseteq \text{MsgSp}_{\mathcal{MA}}$ and that all messages in $\text{MsgSp}_{\mathcal{SE}[F, \mathcal{MA}]}$ have length at most $L \cdot (2^o - 1)$. Note that the former means that the message space of \mathcal{MA} actually consists of pairs of strings. Let $\text{NonceSp}_{\mathcal{SE}[F, \mathcal{MA}]} = \text{NonceSp}_{\mathcal{MA}}$. Let $\mathcal{SE}[F, \mathcal{MA}]$'s component algorithms be defined as follows:

Algorithm \mathcal{K}_e

```

 $f \xleftarrow{\$} F$ 
 $K \xleftarrow{\$} \mathcal{K}_m$ 
Return  $\langle f, K \rangle$ 

```

Algorithm $\mathcal{E}_{\langle f, K \rangle}^{N, A}(M)$

```

 $\sigma \leftarrow \text{CTR-MODE}_f^N(M)$ 
 $\tau \leftarrow \mathcal{T}_K^N(A, \sigma)$ 
Return  $\sigma \parallel \tau$ 

```

Algorithm $\mathcal{D}_{\langle f, K \rangle}^{N, A}(C)$

```

If  $|C| < t$  then return INVALID
Parse  $C$  as  $\sigma \parallel \tau$  //  $|\tau| = t$ 
If  $A \notin \text{AdSp}_{\mathcal{SE}[F, \mathcal{MA}]}$  or  $\sigma \notin \text{MsgSp}_{\mathcal{SE}[F, \mathcal{MA}]}$  then return INVALID
 $\tau' \leftarrow \mathcal{T}_K^N(A, \sigma)$ 
If  $\tau \neq \tau'$  return INVALID
 $M \leftarrow \text{CTR-MODE}_f^N(\sigma)$ 
Return  $M$ 

```

Algorithm $\text{CTR-MODE}_f^N(X)$

```

 $\alpha \leftarrow \lceil |X|/L \rceil$ 
For  $i = 1$  to  $\alpha$  do

```

$Z_i \leftarrow f(\text{ctr1}(N, i))$
 $Y \leftarrow (\text{first } |X| \text{ bits of } Z_1 \| Z_2 \| \dots \| Z_\alpha) \oplus X$
 Return Y ■

Proposition C.13 Let $\mathcal{SE}[F, \mathcal{MA}]$ be as in Construction C.12. Then given a nonce-respecting AUTH adversary B against $\mathcal{SE}[F, \mathcal{MA}]$, we can construct a nonce-respecting forgery adversary D_B against \mathcal{MA} such that

$$\mathbf{Adv}_{\mathcal{SE}[F, \mathcal{MA}]}^{\text{auth}}(B) \leq \mathbf{Adv}_{\mathcal{MA}}^{\text{uf}}(D_B).$$

Furthermore the experiment for D_B uses the same time as the experiment for B and if B makes q encryption oracle queries, then D_B makes q tagging oracle queries. ■

The approach used in [1] when analyzing Encrypt-then-Authenticate constructions can be used to prove Proposition C.13. The only difference is that we consider MACs that also take nonces as input.

COMBINING THESE CONSTRUCTIONS. Let us now combine these constructions.

Construction C.14 [Combined CWC.] Let l, L, n, o, t, k be positive integers such that $t \leq L$. (Further restrictions will be placed shortly.) Essentially, l is the length of the input to a PRF (e.g., 128), L is the length of the output from the PRF (e.g., 128), n is the length of the nonce (e.g., 88), o is the length of the offset (e.g., 32), t is the length of the desired tag (e.g., 64 or 128), k is the length of the hash function's keysize (e.g., 128).

Let F be a family of functions from $\{0, 1\}^l$ to $\{0, 1\}^L$. Let $\mathcal{HF} = (\mathcal{K}_h, \mathcal{H})$ be a family of hash functions with $\text{HashSp}_{\mathcal{HF}} = \{0, 1\}^l$ and $\text{KeySp}_{\mathcal{HF}} = \{0, 1\}^k$ (and \mathcal{K}_h works by randomly selecting and returning an element from $\{0, 1\}^k$ with uniform probability). Let $\text{ctr1} : \{0, 1\}^n \times (\mathbb{Z}_{2^o} - \{0\}) \rightarrow \{0, 1\}^l$ be an efficiently-computable injective function. Let $\text{extract} : \{0, 1\}^{\lceil k/L \rceil \cdot L} \rightarrow \{0, 1\}^k$ be a function that takes as input a $\lceil k/L \rceil \cdot L$ -bit string and that outputs a k -bit string. We require that extract always pick the same k bits from the input string and always outputs those bits in the exact same order (e.g., extract returns the first k bits of its input).

Let $\mathcal{SE}[F, \mathcal{HF}] = (\mathcal{K}_e, \mathcal{E}, \mathcal{D})$ be an AEAD mode built from function family F and hash function \mathcal{HF} and using the above functions extract and ctr1 . We assume that $\text{AdSp}_{\mathcal{SE}[F, \mathcal{HF}]} \times \text{MsgSp}_{\mathcal{SE}[F, \mathcal{HF}]} \subseteq \text{MsgSp}_{\mathcal{HF}}$ and that all messages in $\text{MsgSp}_{\mathcal{SE}[F, \mathcal{HF}]}$ have length at most $L \cdot (2^o - 1)$. Note that the former means that the message space of \mathcal{HF} actually consists of pairs of strings. Let $\text{NonceSp}_{\mathcal{SE}[F, \mathcal{HF}]} = \{0, 1\}^n$. Let $\mathcal{SE}[F, \mathcal{HF}]$'s component algorithms be defined as follows:

Algorithm \mathcal{K}_e

$f \xleftarrow{\$} F$
 $d \xleftarrow{\$} \text{Func}(\mathbb{Z}_{\lceil k/L \rceil}, \{0, 1\}^L)$; $e \xleftarrow{\$} \text{Func}(\text{HashSp}_{\mathcal{HF}}, \{0, 1\}^t)$; $g \xleftarrow{\$} \text{Func}(\text{NonceSp}_{\mathcal{SE}[F, \mathcal{HF}]}, \{0, 1\}^t)$
 $K_h \leftarrow \text{extract}(d(0) \| d(1) \| \dots \| d(\lceil k/L \rceil - 1))$; $H \leftarrow \mathcal{H}_{K_h}$
 Return $\langle f, H, e, g \rangle$

Algorithm $\mathcal{E}_{\langle f, H, e, g \rangle}^{N, A}(M)$

$\sigma \leftarrow \text{CTR-MODE}_f^N(M)$
 $\tau \leftarrow g(N) \oplus e(H(A, \sigma))$
 Return $\sigma \| \tau$

Algorithm $\mathcal{D}_{\langle f, H, e, g \rangle}^{N, A}(C)$
 If $|C| < t$ then return INVALID
 Parse C as $\sigma \parallel \tau$ // $|\tau| = t$
 If $A \notin \text{AdSp}_{\mathcal{SE}[F, \mathcal{HF}]}$ or $\sigma \notin \text{MsgSp}_{\mathcal{SE}[F, \mathcal{HF}]}$ then return INVALID
 $\tau' \leftarrow g(N) \oplus e(H(A, \sigma))$
 If $\tau \neq \tau'$ return INVALID
 $M \leftarrow \text{CTR-MODE}_f^N(\sigma)$
 Return M

Algorithm $\text{CTR-MODE}_f^N(X)$
 $\alpha \leftarrow \lceil |X|/L \rceil$
 For $i = 1$ to α do
 $Z_i \leftarrow f(\text{ctr1}(N, i))$
 $Y \leftarrow (\text{first } |X| \text{ bits of } Z_1 \parallel Z_2 \parallel \dots \parallel Z_\alpha) \oplus X$
 Return Y ■

Proposition C.15 Let $\mathcal{SE}[F, \mathcal{HF}]$ be as in Construction C.14 and let \mathcal{HF} be an ϵ -AU hash function. Then the advantage of any nonce-respecting AUTH adversary A in breaking the authenticity of $\mathcal{SE}[F, \mathcal{HF}]$ is upper bounded by

$$\text{Adv}_{\mathcal{SE}[F, \mathcal{HF}]}^{\text{auth}}(A) \leq \epsilon + 2^{-t} . \blacksquare$$

Proof of Proposition C.15: We first note that the steps $d \stackrel{\$}{\leftarrow} \text{Func}(\mathbb{Z}_{\lceil k/L \rceil}, \{0, 1\}^L)$; $K_h \leftarrow \text{extract}(d(0) \parallel d(1) \parallel \dots \parallel d(\lceil k/L \rceil - 1))$; $H \leftarrow \mathcal{H}_{K_h}$ is equivalent to the step $H \stackrel{\$}{\leftarrow} \mathcal{HF}$.

Note that $e(H(A, \sigma))$ can be rewritten as $\overline{\mathcal{H}}_{(H, e)}(A, \sigma)$ where $\overline{\mathcal{H}\mathcal{F}}[t] = (\overline{K}_h, \overline{\mathcal{H}})$ is composed from \mathcal{HF} per Construction C.8.

Also note that $g(N) \oplus \overline{\mathcal{H}}_{(H, e)}(A, \sigma)$ can be replaced with $\mathcal{T}_{(\overline{\mathcal{H}}_{(H, e)}, g)}^N(A, \sigma)$ where $\mathcal{MA} = (\mathcal{K}_m, \mathcal{T}, \mathcal{V})$ is composed from $\overline{\mathcal{H}\mathcal{F}}[t]$ as per Construction C.10.

By Proposition C.13, given A we can construct an adversary B_A against \mathcal{MA} such that

$$\text{Adv}_{\mathcal{SE}[F, \mathcal{HF}]}^{\text{auth}}(A) \leq \text{Adv}_{\mathcal{MA}}^{\text{uf}}(B_A) .$$

By Proposition C.11 we know that

$$\text{Adv}_{\mathcal{MA}}^{\text{uf}}(B_A) \leq \epsilon'$$

where ϵ' is $\epsilon + 2^{-t}$ (the latter by Proposition C.9). ■

INTEGRITY OF $\mathcal{SE}[\text{Func}(l, L), \mathcal{HF}]$. We now consider the integrity of $\mathcal{SE}[\text{Func}(l, L), \mathcal{HF}]$.

Proposition C.16 Let $\mathcal{SE}[\text{Func}(l, L), \mathcal{HF}]$ be a AEAD mode as in Construction C.1. Then for any nonce-respecting AUTH adversary A against $\mathcal{SE}[\text{Func}(l, L), \mathcal{HF}]$, we have that

$$\text{Adv}_{\mathcal{SE}[\text{Func}(l, L), \mathcal{HF}]}^{\text{auth}}(A) \leq \epsilon + 2^{-t} . \blacksquare$$

Proof of Proposition C.16: Let $\mathcal{SE}'[\text{Func}(l, L), \mathcal{HF}]$ be as in Construction C.14. Note that $\mathcal{SE}[\text{Func}(l, L), \mathcal{HF}]$ and $\mathcal{SE}'[\text{Func}(l, L), \mathcal{HF}]$ are identical except that the former uses only one random function f and $\mathcal{SE}'[\text{Func}(l, L), \mathcal{HF}]$ uses four random functions (one to generate the hash key, one to CTR-mode encrypt the message, one to encipher the output of the hash function,

and one to CTR-mode encrypt the output of the hash function). Furthermore, recall that, for $\mathcal{SE}[\text{Func}(l, L), \mathcal{HF}]$, there is never a collision in the input to f between the four different uses of f (this was a requirement imposed on \mathcal{HF} , ctr0 , ctr1 , and ctr2). Consequently, the fact that $\mathcal{SE}'[\text{Func}(l, L), \mathcal{HF}]$ uses four random functions and $\mathcal{SE}[\text{Func}(l, L), \mathcal{HF}]$ uses one is immaterial. Hence the probability that A forges against $\mathcal{SE}[\text{Func}(l, L), \mathcal{HF}]$ is the same as the probability that it forges against $\mathcal{SE}'[\text{Func}(l, L), \mathcal{HF}]$. I.e.,

$$\mathbf{Adv}_{\mathcal{SE}[\text{Func}(l, L), \mathcal{HF}]}^{\text{auth}}(A) = \mathbf{Adv}_{\mathcal{SE}'[\text{Func}(l, L), \mathcal{HF}]}^{\text{auth}}(A).$$

By Proposition C.15, we know the latter probability is upper bounded by $\epsilon + 2^{-t}$. \blacksquare

PROOF OF LEMMA C.5. We now prove Lemma C.5.

Proof of Lemma C.5: Adversary B_A runs A and replies to A 's oracle queries using its oracle f . If A returns a valid forgery, B_A returns 1, otherwise B_A returns 0. This implies that

$$\mathbf{Adv}_{\mathcal{SE}[F, \mathcal{HF}]}^{\text{auth}}(A) = \Pr \left[f \stackrel{\$}{\leftarrow} F : B_A^{f(\cdot)} = 1 \right]$$

and

$$\mathbf{Adv}_{\mathcal{SE}[\text{Func}(l, L), \mathcal{HF}]}^{\text{auth}}(A) = \Pr \left[f \stackrel{\$}{\leftarrow} \text{Func}(l, L) : B_A^{f(\cdot)} = 1 \right].$$

Since

$$\mathbf{Adv}_{\mathcal{SE}[\text{Func}(l, L), \mathcal{HF}]}^{\text{auth}}(A) \leq \epsilon + 2^{-t}$$

by Proposition C.16, we have

$$\begin{aligned} \mathbf{Adv}_{\mathcal{SE}[F, \mathcal{HF}]}^{\text{auth}}(A) &= \mathbf{Adv}_{\mathcal{SE}[F, \mathcal{HF}]}^{\text{auth}}(A) - \mathbf{Adv}_{\mathcal{SE}[\text{Func}(l, L), \mathcal{HF}]}^{\text{auth}}(A) + \mathbf{Adv}_{\mathcal{SE}[\text{Func}(l, L), \mathcal{HF}]}^{\text{auth}}(A) \\ &\leq \Pr \left[f \stackrel{\$}{\leftarrow} F : B_A^{f(\cdot)} = 1 \right] - \Pr \left[f \stackrel{\$}{\leftarrow} \text{Func}(l, L) : B_A^{f(\cdot)} = 1 \right] \\ &\quad + \epsilon + 2^{-t} \\ &= \mathbf{Adv}_F^{\text{prf}}(B_A) + \epsilon + 2^{-t} \end{aligned}$$

as desired. \blacksquare

C.6 Proof of Lemma C.6

Proof of Lemma C.6: Let B_A be a PRF adversary against F that uses adversary A and that has oracle access to a function $g : \{0, 1\}^l \rightarrow \{0, 1\}^L$. Adversary B_A runs A and replies to A 's encryption oracle queries using its own oracle $g(\cdot)$ for the function f in Construction C.1. Adversary B_A returns the same bit that A returns. Then

$$\Pr \left[\langle f, H \rangle \stackrel{\$}{\leftarrow} \mathcal{K}_e : A^{\mathcal{E}(f, h)(\cdot, \cdot)} = 1 \right] = \Pr \left[g \stackrel{\$}{\leftarrow} F : B_A^{g(\cdot)} = 1 \right]$$

since when B_A is given a random instance of F it runs A exactly as if A was given the real encryption oracle. Furthermore

$$\Pr \left[A^{\mathcal{E}(\cdot, \cdot)} = 1 \right] = \Pr \left[g \stackrel{\$}{\leftarrow} \text{Func}(l, L) : B_A^{g(\cdot)} = 1 \right]$$

since B_A replies to all of A 's oracle queries with independently selected random strings. Consequently

$$\mathbf{Adv}_{\mathcal{SE}[F, \mathcal{HF}]}^{\text{priv}}(A) \leq \mathbf{Adv}_F^{\text{prf}}(B_A)$$

as desired. \blacksquare

D Test vectors

Vector #1: CWC-AES-128

AES KEY: 00 01 02 03 04 05 06 07 08 09 0A 0B 0C 0D 0E 0F
PLAINTEXT: 00 01 02 03 04 05 06 07
ASSOC DATA: <None>
NONCE: FF EE DD CC BB AA 99 88 77 66 55

HASH KEY: 34 AE 6A 6F E9 51 78 94 AC CC BB 9E BA E7 20 8C
HASH VALUE: 2B 9E AE BE 67 3F AE 03 6B 16 EA 31 DC A7 AE 6B
AES(HVAL): FC DC 06 4C CD CA FE E3 DE 7A A3 CF 5C 5D B9 7B
MAC CTR PT: 80 FF EE DD CC BB AA 99 88 77 66 55 00 00 00 00
AES(MCPT): AB 89 DD E9 C4 55 C1 FE BE 7E E7 58 82 D4 8A D2
CIPHERTEXT: 88 B8 DF 06 28 FD 51 CC 57 55 DB A5 09 9F 3F 1D
60 04 44 97 DE 89 33 A9

Vector #2: CWC-AES-192

AES KEY: 00 01 02 03 04 05 06 07 08 09 0A 0B 0C 0D 0E 0F
FO E0 D0 C0 B0 A0 90 80
PLAINTEXT: 00 01 02 03 04 05 06 07
ASSOC DATA: <None>
NONCE: FF EE DD CC BB AA 99 88 77 66 55

HASH KEY: 4F A8 88 AF 06 83 60 0C AB 35 75 EF 0A E6 01 A5
HASH VALUE: 40 E6 24 83 4B 27 9A 7B 15 42 C7 FE 29 EB 29 A3
AES(HVAL): 69 CC 0E 3D 96 98 EB 75 1F 06 A5 90 9B C2 4F 5A
MAC CTR PT: 80 FF EE DD CC BB AA 99 88 77 66 55 00 00 00 00
AES(MCPT): C6 B6 F4 33 F9 12 39 4F 6A 8C B9 D3 F2 7B 0C 72
CIPHERTEXT: F0 DB A9 74 12 30 01 B0 AF 7A FA 0E 6F 8A D2 3A
75 8A 1C 43 69 B9 43 28

Vector #3: CWC-AES-256

AES KEY: 00 01 02 03 04 05 06 07 08 09 0A 0B 0C 0D 0E 0F
FO E0 D0 C0 B0 A0 90 80 70 60 50 40 30 20 10 00
PLAINTEXT: 00 01 02 03 04 05 06 07
ASSOC DATA: <None>
NONCE: FF EE DD CC BB AA 99 88 77 66 55

HASH KEY: 35 8F 2B 0C FF E9 84 BE F9 EE EE 55 85 36 BC E5
HASH VALUE: 18 99 E1 A6 1E 6E 37 65 C6 3A 41 99 56 8C D1 BF
AES(HVAL): 1C 56 65 0A 22 BC B5 94 AC F3 CA 24 46 03 B8 5E
MAC CTR PT: 80 FF EE DD CC BB AA 99 88 77 66 55 00 00 00 00
AES(MCPT): 92 0A 3B 46 82 25 16 F1 5A A3 1B AE 8D EB 72 A0
CIPHERTEXT: 7B CF 73 BE 46 9C 46 0B 8E 5C 5E 4C A0 99 A3 65
F6 50 D1 8A CB E8 CA FE

Vector #4: CWC-AES-128

AES KEY: 00 01 02 03 04 05 06 07 08 09 0A 0B 0C 0D 0E 0F

PLAINTEXT: 00 01 02 03 04 05 06 07
ASSOC DATA: 54 68 69 73 20 69 73 20 61 20 70 6C 61 69 6E 74
65 78 74 20 68 65 61 64 65 72 2E 00
NONCE: FF EE DD CC BB AA 99 88 77 66 55

HASH KEY: 34 AE 6A 6F E9 51 78 94 AC CC BB 9E BA E7 20 8C
HASH VALUE: 2E A9 2A A5 28 B1 1C 08 1C C8 2F 24 9B E4 19 8D
AES(HVAL): EA 54 F8 3D 56 7F 53 05 88 B1 EA 96 36 79 CD AC
MAC CTR PT: 80 FF EE DD CC BB AA 99 88 77 66 55 00 00 00 00
AES(MCPT): AB 89 DD E9 C4 55 C1 FE BE 7E E7 58 82 D4 8A D2
CIPHERTEXT: 88 B8 DF 06 28 FD 51 CC 41 DD 25 D4 92 2A 92 FB
36 CF 0D CE B4 AD 47 7E

Vector #5: CWC-AES-192
AES KEY: 00 01 02 03 04 05 06 07 08 09 0A 0B 0C 0D 0E 0F
F0 E0 D0 C0 B0 A0 90 80
PLAINTEXT: 00 01 02 03 04 05 06 07
ASSOC DATA: 54 68 69 73 20 69 73 20 61 20 70 6C 61 69 6E 74
65 78 74 20 68 65 61 64 65 72 2E 00
NONCE: FF EE DD CC BB AA 99 88 77 66 55

HASH KEY: 4F A8 88 AF 06 83 60 0C AB 35 75 EF 0A E6 01 A5
HASH VALUE: 60 3F FC 24 71 64 2E D9 57 E1 B1 EA F2 F8 B0 34
AES(HVAL): D8 39 86 2A 33 5A 54 68 C8 16 DA 47 69 A2 10 EB
MAC CTR PT: 80 FF EE DD CC BB AA 99 88 77 66 55 00 00 00 00
AES(MCPT): C6 B6 F4 33 F9 12 39 4F 6A 8C B9 D3 F2 7B 0C 72
CIPHERTEXT: F0 DB A9 74 12 30 01 B0 1E 8F 72 19 CA 48 6D 27
A2 9A 63 94 9B D9 1C 99

Vector #6: CWC-AES-256
AES KEY: 00 01 02 03 04 05 06 07 08 09 0A 0B 0C 0D 0E 0F
F0 E0 D0 C0 B0 A0 90 80 70 60 50 40 30 20 10 00
PLAINTEXT: 00 01 02 03 04 05 06 07
ASSOC DATA: 54 68 69 73 20 69 73 20 61 20 70 6C 61 69 6E 74
65 78 74 20 68 65 61 64 65 72 2E 00
NONCE: FF EE DD CC BB AA 99 88 77 66 55

HASH KEY: 35 8F 2B 0C FF E9 84 BE F9 EE EE 55 85 36 BC E5
HASH VALUE: 0A C6 B1 39 57 7F 26 DA 94 16 42 E1 6D 73 EC B5
AES(HVAL): 4B A5 AD 1E 74 A2 C5 BE AB D0 DA 4D F4 29 83 0C
MAC CTR PT: 80 FF EE DD CC BB AA 99 88 77 66 55 00 00 00 00
AES(MCPT): 92 0A 3B 46 82 25 16 F1 5A A3 1B AE 8D EB 72 A0
CIPHERTEXT: 7B CF 73 BE 46 9C 46 0B D9 AF 96 58 F6 87 D3 4F
F1 73 C1 E3 79 C2 F1 AC

Vector #7: CWC-AES-128
AES KEY: 00 01 02 03 04 05 06 07 08 09 0A 0B 0C 0D 0E 0F
PLAINTEXT: 00 01 02 03 04 05 06 07 08 09 0A 0B 0C 0D 0E

ASSOC DATA: <None>

NONCE: FF EE DD CC BB AA 99 88 77 66 55

HASH KEY: 34 AE 6A 6F E9 51 78 94 AC CC BB 9E BA E7 20 8C
HASH VALUE: 79 00 74 72 E1 C8 36 96 ED 7A B1 F9 03 6E 94 8B
AES(HVAL): 2B 0F 24 69 B1 2B BE 39 C9 40 67 BA F1 25 E2 5B
MAC CTR PT: 80 FF EE DD CC BB AA 99 88 77 66 55 00 00 00 00
AES(MCPT): AB 89 DD E9 C4 55 C1 FE BE 7E E7 58 82 D4 8A D2
CIPHERTEXT: 88 B8 DF 06 28 FD 51 CC 31 E6 6E 57 0B 0F 77 80
86 F9 80 75 7E 7F C7 77 3E 80 E2 73 F1 68 89

Vector #8: CWC-AES-192

AES KEY: 00 01 02 03 04 05 06 07 08 09 0A 0B 0C 0D 0E 0F
FO E0 D0 C0 B0 A0 90 80

PLAINTEXT: 00 01 02 03 04 05 06 07 08 09 0A 0B 0C 0D 0E

ASSOC DATA: <None>

NONCE: FF EE DD CC BB AA 99 88 77 66 55

HASH KEY: 4F A8 88 AF 06 83 60 0C AB 35 75 EF 0A E6 01 A5
HASH VALUE: 2C 5E 3A A4 37 1C 27 D6 E8 6B 76 DC 3D 93 BC 87
AES(HVAL): 48 6E 9C E5 C3 16 3E A6 9C D4 D7 E2 7C 9D 92 D2
MAC CTR PT: 80 FF EE DD CC BB AA 99 88 77 66 55 00 00 00 00
AES(MCPT): C6 B6 F4 33 F9 12 39 4F 6A 8C B9 D3 F2 7B 0C 72
CIPHERTEXT: F0 DB A9 74 12 30 01 B0 E1 42 B7 58 87 C9 00 8E
D8 68 D6 3A 04 07 E9 F6 58 6E 31 8E E6 9E A0

Vector #9: CWC-AES-256

AES KEY: 00 01 02 03 04 05 06 07 08 09 0A 0B 0C 0D 0E 0F
FO E0 D0 C0 B0 A0 90 80 70 60 50 40 30 20 10 00

PLAINTEXT: 00 01 02 03 04 05 06 07 08 09 0A 0B 0C 0D 0E

ASSOC DATA: <None>

NONCE: FF EE DD CC BB AA 99 88 77 66 55

HASH KEY: 35 8F 2B 0C FF E9 84 BE F9 EE EE 55 85 36 BC E5
HASH VALUE: 4A 70 29 CC 58 25 52 CB 75 AD C9 60 FF B3 F7 55
AES(HVAL): 2B 64 0E 02 CE 51 DE 22 B2 0F 2A 8D C4 23 CD C0
MAC CTR PT: 80 FF EE DD CC BB AA 99 88 77 66 55 00 00 00 00
AES(MCPT): 92 0A 3B 46 82 25 16 F1 5A A3 1B AE 8D EB 72 A0
CIPHERTEXT: 7B CF 73 BE 46 9C 46 0B 9B C6 2D DE 26 DD 47 B9
6E 35 44 4C 74 C8 D3 E8 AC 31 23 49 C8 BF 60

Vector #10: CWC-AES-128

AES KEY: 00 01 02 03 04 05 06 07 08 09 0A 0B 0C 0D 0E 0F

PLAINTEXT: 00 01 02 03 04 05 06 07 08 09 0A 0B 0C 0D 0E

ASSOC DATA: 54 68 69 73 20 69 73 20 61 20 70 6C 61 69 6E 74
65 78 74 20 68 65 61 64 65 72 2E 00

NONCE: FF EE DD CC BB AA 99 88 77 66 55

HASH KEY: 34 AE 6A 6F E9 51 78 94 AC CC BB 9E BA E7 20 8C
HASH VALUE: 51 AE 9D 7E 86 BD E0 B2 AA 18 2C 91 87 0A 9C A5
AES(HVAL): DF 48 30 BD 1D DC E0 59 B1 C2 0B 29 01 4F 80 10
MAC CTR PT: 80 FF EE DD CC BB AA 99 88 77 66 55 00 00 00 00
AES(MCPT): AB 89 DD E9 C4 55 C1 FE BE 7E E7 58 82 D4 8A D2
CIPHERTEXT: 88 B8 DF 06 28 FD 51 CC 31 E6 6E 57 0B 0F 77 74
C1 ED 54 D9 89 21 A7 0F BC EC 71 83 9B 0A C2

Vector #11: CWC-AES-192

AES KEY: 00 01 02 03 04 05 06 07 08 09 0A 0B 0C 0D 0E 0F
FO E0 D0 C0 B0 A0 90 80
PLAINTEXT: 00 01 02 03 04 05 06 07 08 09 0A 0B 0C 0D 0E
ASSOC DATA: 54 68 69 73 20 69 73 20 61 20 70 6C 61 69 6E 74
65 78 74 20 68 65 61 64 65 72 2E 00
NONCE: FF EE DD CC BB AA 99 88 77 66 55

HASH KEY: 4F A8 88 AF 06 83 60 0C AB 35 75 EF 0A E6 01 A5
HASH VALUE: 51 60 E7 81 DC 64 F9 CD 54 BA 02 40 A2 E8 EE 99
AES(HVAL): A0 30 58 13 22 B6 80 53 64 B0 3E 52 41 D2 2D 0A
MAC CTR PT: 80 FF EE DD CC BB AA 99 88 77 66 55 00 00 00 00
AES(MCPT): C6 B6 F4 33 F9 12 39 4F 6A 8C B9 D3 F2 7B 0C 72
CIPHERTEXT: F0 DB A9 74 12 30 01 B0 E1 42 B7 58 87 C9 00 66
86 AC 20 DB A4 B9 1C 0E 3C 87 81 B3 A9 21 78

Vector #12: CWC-AES-256

AES KEY: 00 01 02 03 04 05 06 07 08 09 0A 0B 0C 0D 0E 0F
FO E0 D0 C0 B0 A0 90 80 70 60 50 40 30 20 10 00
PLAINTEXT: 00 01 02 03 04 05 06 07 08 09 0A 0B 0C 0D 0E
ASSOC DATA: 54 68 69 73 20 69 73 20 61 20 70 6C 61 69 6E 74
65 78 74 20 68 65 61 64 65 72 2E 00
NONCE: FF EE DD CC BB AA 99 88 77 66 55

HASH KEY: 35 8F 2B 0C FF E9 84 BE F9 EE EE 55 85 36 BC E5
HASH VALUE: 3F F5 0C 60 E6 01 7A 3C A1 BB B3 54 65 02 85 7C
AES(HVAL): 3E EF A2 E4 97 91 82 86 73 0C F6 E9 46 2C CA 15
MAC CTR PT: 80 FF EE DD CC BB AA 99 88 77 66 55 00 00 00 00
AES(MCPT): 92 0A 3B 46 82 25 16 F1 5A A3 1B AE 8D EB 72 A0
CIPHERTEXT: 7B CF 73 BE 46 9C 46 0B 9B C6 2D DE 26 DD 47 AC
E5 99 A2 15 B4 94 77 29 AF ED 47 CB C7 B8 B5

Vector #13: CWC-AES-128

AES KEY: 00 01 02 03 04 05 06 07 08 09 0A 0B 0C 0D 0E 0F
PLAINTEXT: 00 01 02 03 04 05 06 07 08 09 0A 0B 0C 0D 0E 0F
80 81 82 83 84 85 86 87 88 89 8A 8B 8C 8D 8E 8F
ASSOC DATA: <None>
NONCE: FF EE DD CC BB AA 99 88 77 66 55

HASH KEY: 34 AE 6A 6F E9 51 78 94 AC CC BB 9E BA E7 20 8C

HASH VALUE: 58 D5 28 89 4F 1F 6A 52 A6 44 FA 69 65 C0 73 A6
AES(HVAL): A3 9E F3 6F 67 1F FA F8 71 0C 83 BB 49 A6 6E BC
MAC CTR PT: 80 FF EE DD CC BB AA 99 88 77 66 55 00 00 00 00
AES(MCPT): AB 89 DD E9 C4 55 C1 FE BE 7E E7 58 82 D4 8A D2
CIPHERTEXT: 88 B8 DF 06 28 FD 51 CC 31 E6 6E 57 0B 0F 77 0F
48 5B 82 64 6E CF B9 F9 A0 B0 75 4F D5 94 36 5A
08 17 2E 86 A3 4A 3B 06 CF 72 64 E3 CB 72 E4 6E

Vector #14: CWC-AES-192

AES KEY: 00 01 02 03 04 05 06 07 08 09 0A 0B 0C 0D 0E 0F
F0 E0 D0 C0 B0 A0 90 80
PLAINTEXT: 00 01 02 03 04 05 06 07 08 09 0A 0B 0C 0D 0E 0F
80 81 82 83 84 85 86 87 88 89 8A 8B 8C 8D 8E 8F
ASSOC DATA: <None>
NONCE: FF EE DD CC BB AA 99 88 77 66 55

HASH KEY: 4F A8 88 AF 06 83 60 0C AB 35 75 EF 0A E6 01 A5
HASH VALUE: 0D 0A D2 78 1E 8F E8 47 00 85 31 28 B1 E3 49 3A
AES(HVAL): 5A 05 AA 45 88 06 A9 C1 DC 5A F6 AF 6F 8F EC F6
MAC CTR PT: 80 FF EE DD CC BB AA 99 88 77 66 55 00 00 00 00
AES(MCPT): C6 B6 F4 33 F9 12 39 4F 6A 8C B9 D3 F2 7B 0C 72
CIPHERTEXT: F0 DB A9 74 12 30 01 B0 E1 42 B7 58 87 C9 00 A3
A4 C4 70 6D 40 41 F4 F9 58 E1 3F D0 D7 60 4D 1E
9C B3 5E 76 71 14 90 8E B6 D6 4F 7C 9D F4 E0 84

Vector #15: CWC-AES-256

AES KEY: 00 01 02 03 04 05 06 07 08 09 0A 0B 0C 0D 0E 0F
F0 E0 D0 C0 B0 A0 90 80 70 60 50 40 30 20 10 00
PLAINTEXT: 00 01 02 03 04 05 06 07 08 09 0A 0B 0C 0D 0E 0F
80 81 82 83 84 85 86 87 88 89 8A 8B 8C 8D 8E 8F
ASSOC DATA: <None>
NONCE: FF EE DD CC BB AA 99 88 77 66 55

HASH KEY: 35 8F 2B 0C FF E9 84 BE F9 EE EE 55 85 36 BC E5
HASH VALUE: 02 F2 DA E9 83 72 0E BC DC 77 89 3B 67 CB 3D B7
AES(HVAL): B7 F6 AE DE A3 95 35 FE 03 93 08 DF E0 C7 F1 78
MAC CTR PT: 80 FF EE DD CC BB AA 99 88 77 66 55 00 00 00 00
AES(MCPT): 92 0A 3B 46 82 25 16 F1 5A A3 1B AE 8D EB 72 A0
CIPHERTEXT: 7B CF 73 BE 46 9C 46 0B 9B C6 2D DE 26 DD 47 B5
D2 41 06 CA 5D EB 80 A7 B5 71 0A 38 A4 39 8D BA
25 FC 95 98 21 B0 23 0F 59 30 13 71 6D 2C 83 D8

Vector #16: CWC-AES-128

AES KEY: 00 01 02 03 04 05 06 07 08 09 0A 0B 0C 0D 0E 0F
PLAINTEXT: 00 01 02 03 04 05 06 07 08 09 0A 0B 0C 0D 0E 0F
80 81 82 83 84 85 86 87 88 89 8A 8B 8C 8D 8E 8F
ASSOC DATA: 54 68 69 73 20 69 73 20 61 20 70 6C 61 69 6E 74
65 78 74 20 68 65 61 64 65 72 2E 00

NONCE: FF EE DD CC BB AA 99 88 77 66 55

HASH KEY: 34 AE 6A 6F E9 51 78 94 AC CC BB 9E BA E7 20 8C
HASH VALUE: 05 EE B6 CB DF A6 E5 B8 4C 65 DD F4 8C C8 25 23
AES(HVAL): 62 E5 23 FE 48 8F BC 14 E3 77 15 6C 4D 0F D0 8B
MAC CTR PT: 80 FF EE DD CC BB AA 99 88 77 66 55 00 00 00 00
AES(MCPT): AB 89 DD E9 C4 55 C1 FE BE 7E E7 58 82 D4 8A D2
CIPHERTEXT: 88 B8 DF 06 28 FD 51 CC 31 E6 6E 57 0B 0F 77 0F
48 5B 82 64 6E CF B9 F9 A0 B0 75 4F D5 94 36 5A
C9 6C FE 17 8C DA 7D EA 5D 09 F2 34 CF DB 5A 59

Vector #17: CWC-AES-192

AES KEY: 00 01 02 03 04 05 06 07 08 09 0A 0B 0C 0D 0E 0F
F0 E0 D0 C0 B0 A0 90 80
PLAINTEXT: 00 01 02 03 04 05 06 07 08 09 0A 0B 0C 0D 0E 0F
80 81 82 83 84 85 86 87 88 89 8A 8B 8C 8D 8E 8F
ASSOC DATA: 54 68 69 73 20 69 73 20 61 20 70 6C 61 69 6E 74
65 78 74 20 68 65 61 64 65 72 2E 00
NONCE: FF EE DD CC BB AA 99 88 77 66 55

HASH KEY: 4F A8 88 AF 06 83 60 0C AB 35 75 EF 0A E6 01 A5
HASH VALUE: 10 E1 48 E2 D0 68 39 EC C4 0A 6C A3 D6 8B 47 54
AES(HVAL): 23 0A 37 C3 48 7C 9F 76 05 B9 5D 1A 21 D5 D5 FD
MAC CTR PT: 80 FF EE DD CC BB AA 99 88 77 66 55 00 00 00 00
AES(MCPT): C6 B6 F4 33 F9 12 39 4F 6A 8C B9 D3 F2 7B 0C 72
CIPHERTEXT: F0 DB A9 74 12 30 01 B0 E1 42 B7 58 87 C9 00 A3
A4 C4 70 6D 40 41 F4 F9 58 E1 3F D0 D7 60 4D 1E
E5 BC C3 F0 B1 6E A6 39 6F 35 E4 C9 D3 AE D9 8F

Vector #18: CWC-AES-256

AES KEY: 00 01 02 03 04 05 06 07 08 09 0A 0B 0C 0D 0E 0F
F0 E0 D0 C0 B0 A0 90 80 70 60 50 40 30 20 10 00
PLAINTEXT: 00 01 02 03 04 05 06 07 08 09 0A 0B 0C 0D 0E 0F
80 81 82 83 84 85 86 87 88 89 8A 8B 8C 8D 8E 8F
ASSOC DATA: 54 68 69 73 20 69 73 20 61 20 70 6C 61 69 6E 74
65 78 74 20 68 65 61 64 65 72 2E 00
NONCE: FF EE DD CC BB AA 99 88 77 66 55

HASH KEY: 35 8F 2B 0C FF E9 84 BE F9 EE EE 55 85 36 BC E5
HASH VALUE: 09 4D C5 21 94 79 E0 58 4E E9 C1 2C 29 6A E3 A4
AES(HVAL): E9 69 49 47 09 07 62 3B A9 8D AD 51 9F D5 D1 F7
MAC CTR PT: 80 FF EE DD CC BB AA 99 88 77 66 55 00 00 00 00
AES(MCPT): 92 0A 3B 46 82 25 16 F1 5A A3 1B AE 8D EB 72 A0
CIPHERTEXT: 7B CF 73 BE 46 9C 46 0B 9B C6 2D DE 26 DD 47 B5
D2 41 06 CA 5D EB 80 A7 B5 71 0A 38 A4 39 8D BA
7B 63 72 01 8B 22 74 CA F3 2E B6 FF 12 3E A3 57