

CWC: A high-performance conventional authenticated encryption mode

Tadayoshi Kohno¹, John Viega², and Doug Whiting³

¹ UC San Diego, tkohno@cs.ucsd.edu

² Virginia Tech, viega@securesoftware.com

³ Hifn, Inc., dwhiting@hifn.com

Abstract. We introduce CWC, a new block cipher mode of operation for protecting both the privacy and the authenticity of encapsulated data. CWC is currently the only such mode having all five of the following properties: provable security, parallelizability, high performance in hardware, high performance in software, and no intellectual property concerns. We believe that having all five of these properties makes CWC a powerful tool for use in many performance-critical cryptographic applications. CWC is also the only appropriate solution for some applications; e.g., standardization bodies like the IETF and NIST prefer patent-free modes, and CWC is the only such mode capable of processing data at 10Gbps in hardware, which will be important for future IPsec (and other) network devices. As part of our design, we also introduce a new parallelizable universal hash function optimized for performance in both hardware and software.

Keywords: Authenticated encryption, modes of operation, parallelism, performance, security proofs.

1 Introduction

An *authenticated encryption associated data* (AEAD) scheme is a symmetric encryption scheme designed to protect both the privacy and the authenticity of encapsulated data. There has recently been a strong push toward producing block cipher-based AEAD schemes [13, 10, 12, 23, 27, 22, 5]. Despite this push, among the previous works there does not exist any AEAD scheme simultaneously having all of the following properties: provable security, parallelizability, high performance in hardware, high performance in software, and free from intellectual property concerns. Even though not all applications will require all five of the these properties, almost all applications will require at least one of the them, and may very likely have to be able to interoperate with an application requiring a different property. We thus view finding an appropriate scheme having all five of these properties as a very important research goal.

Finding an appropriate balance between all five of the aforementioned properties is, however, not easy because the most natural approaches to addressing

some of the properties are actually disadvantageous with respect to other properties. We believe we have overcome these challenges and, in doing so, introduce a new mode of operation called CWC, or Carter-Wegman Counter mode.

MOTIVATING EXAMPLE. One of the primary reasons for desiring such a block cipher-based AEAD scheme is IPsec. From a pragmatic perspective, we note that many vendors and standardization bodies prefer patent-free modes over patented modes (the elegant OCB mode was apparently rejected from the IEEE 802.11 working group because of patent concerns). And, from a hardware performance perspective, we note that because none of the existing patent-free AEAD schemes are parallelizable, it is impossible to make existing patent-free AEAD schemes run faster than about 2Gbps using conventional ASIC technology and a single processing unit. Nevertheless, future network devices will be expected to run at 10Gbps. CWC addresses these issues, being both patent-free and capable of processing data at 10Gbps using conventional ASIC technology.

THE CWC SOLUTION. Our new mode of operation, called CWC, has all five of the properties mentioned above. It is provably secure. Moreover, our provable security-based analyses helped guide our research and helped us reject other schemes with similar performance properties but with slightly worse provable security bounds. CWC is also parallelizable, which means that we can make CWC-AES run at 10Gbps using conventional ASIC technology. CWC is also fast in software. Our current implementation of CWC-AES runs at about the same speed as the other patent-free modes on 32-bit architectures (Table 1), and we anticipate significant performance gains on 32-bit CPUs when using more sophisticated implementation techniques (Section 6), and we also see significantly better performance on 64-bit architectures. Of course, we do remark that the patented modes like OCB are capable of running even faster in software, which would make them very attractive were they not also encumbered in intellectual property issues.

Like the other two unpatented block cipher-based AEAD modes, CCM [27] and EAX [5], CWC avoids patents by making two passes on the data: one pass to “encrypt” and one pass to “authenticate.” Adopting the terminology used in [5], it is because of the two-pass nature of CWC that we call it a “conventional” block cipher-based AEAD scheme. By using the conventional approach, CCM, EAX, and CWC are very much like composition-based AEAD scheme [4, 15], or AEAD schemes composed from existing encryption schemes and MACs. Unlike composition-based AEAD schemes, however, by designing CWC directly from a block cipher, we eliminate redundant steps and fine-tune CWC for efficiency, again keeping in mind both our hardware and software goals. For example, we use only one block cipher key, which saves expensive memory access in hardware.

The encryption core of CWC is essentially counter (CTR) mode encryption, which is well-known to be efficient and parallelizable. Finding an appropriate algorithm for the authentication core of CWC proved to be more of a challenge. For authentication, we decided to base our design on the Carter-Wegman [26] universal hash function approach for message authentication. Part of the difficulty in the design came down to choosing the right type of universal hash function,

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 block cipher-based AEAD modes on a Pentium III. Values are averaged over 50 000 samples.

with the right parameters. Since polynomial evaluation can be parallelized (if the polynomial is in x , one can split it into i polynomials in x^i), we chose to use a universal hash function consisting of evaluating a polynomial modulo the prime $2^{127} - 1$. We note that our hash function is similar to Bernstein’s hash127 [6] except that Bernstein’s hash function was optimized for software performance at the expense of hardware performance. To address this issue, we use larger coefficients than Bernstein uses. We believe our hardware- and software-optimized universal hash function to be of independent interest.

NOTATION. As part of our research, we first created a general approach for combining CTR mode encryption with a universal hash function in order to provide authenticated encryption. We shall refer to this general approach as CWC (note no change in font), and shall use CWC-BC to refer to a CWC instantiation with a 128-bit block cipher BC as the underlying block cipher and with the universal hash function described briefly above. We shall use CWC as shorthand for CWC-BC and use CWC-AES to mean CWC-BC with AES [8] as the underlying block cipher. Other instantiations of the general CWC approach are possible, e.g., for legacy 64-bit block ciphers. Since we are primarily targeting new applications, and since a mode using a 128-bit block cipher will never be asked to interoperate with a mode using a 64-bit block cipher, we focus this paper only on our 128-bit CWC instantiation.

When we say that an AEAD scheme’s encryption algorithm takes a pair (A, M) as input and produces a ciphertext as output, we mean that the AEAD scheme is designed to protect the privacy of M and the authenticity of both A and M . This will be made more formal in the body.

PERFORMANCE. Let (A, M) be some input to the CWC encryption algorithm. The CWC encryption algorithm derives a universal hash subkey from the block cipher key. Assuming that the universal hash subkey is maintained across invocations, encrypting (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 6). 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 $m = \lceil (d + 1)/n \rceil - 1$.

In hardware using conventional ASIC technology at 0.13 micron, it takes approximately 300 Kgates to reach 10 Gbps throughput for CWC-AES. This is around twice as much as OCB, but avoids IP negotiation overhead and royalty payments to three parties. 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 patented modes such as OCB are not included in this table, but are about twice as fast as the times given for the patent-free modes. The implementations used to compute Table 1 were written in C by Brian Gladman [9] 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 but the authentication portion of CWC-AES is limited by the speed of the universal hash function). For comparison, Bernstein’s related hash127, which also evaluates a polynomial modulo $2^{127} - 1$ but 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. On 64-bit G5s, our initial implementation of the hash function runs at around 6 cpb, thus showing that CWC-AES is very attractive on 64-bit architectures (when running the G5 in 32-bit mode, our implementation runs at around 15 cpb).

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 AEAD scheme for such low-end processors.

THE PATENT ISSUE. The patent issue is a very peculiar one. While it may initially sound odd to let patents influence research, we note that it is also not uncommon, especially in other sciences. Indeed, we view this line of research as discovering the most appropriate solution given real-world constraints. And, just like performance constraints, intellectual property constraints are very real.

BACKGROUND AND RELATED WORK. The notion of an *authenticated encryption (AE) scheme* was formalized by Katz and Yung [13] and by Bellare and Namprempre [4] and the notion of an *authenticated encryption with associated data (AEAD) scheme* was formalized by Rogaway [22]. Bellare and Namprempre [4] and Krawczyk [15] explored ways to combine standard encryption schemes with MACs to achieve authenticated encryption. A number of dedicated AE and AEAD schemes also exist, including RPC [13], XECB [10], IAPM [12], OCB [23], CCM [27], and EAX [5]. CWC is similar to the combination of McGrew’s UST [20] and TMMH [19], where one of the main advantages of CWC over UST+TMMH is CWC’s small key size, which, as the author of UST and TMMH noted, can be

a bottleneck for UST+TMMH in hardware at high speeds. The integrity portion of CWC builds on top of the Carter-Wegman universal hashing approach to message authentication [26]. The specific hash function CWC uses is similar to Bernstein’s hash127 [6], but is better suited for hardware. Shoup [25] and Nevelsteen and Preneel [21] also worked on software optimizations for universal hash functions. Rogaway and Wagner released a critique of CCM [24]. For each issue raised in [24], we find that we have 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, messages; we stress, however, that, if desired, it is easy to modify CWC to handle arbitrary bit-length messages, see Section 5).

2 Preliminaries

NOTATION. If x is a string then $|x|$ denotes its length in bits. 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, $\text{toint}(10000010) = 2^7 + 2 = 130$. If b is a bit and n a non-negative integer, then b^n denote b concatenated with itself n times; e.g., 10^7 is the string 10000000. 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, 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 SCHEMES WITH ASSOCIATED DATA. We use Rogaway’s notion of an *authenticated encryption with associated data (AEAD) scheme* or *mode* [22]. An AEAD scheme $\mathcal{SE} = (\mathcal{K}_e, \mathcal{E}, \mathcal{D})$ 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 scheme, 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 $f \stackrel{s}{\leftarrow} E$ as short hand for $K \stackrel{s}{\leftarrow} \{0, 1\}^k$; $f \leftarrow E_K$. 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.

We adopt the notion of security for block ciphers introduced in [17] and adopted for the concrete setting in [2]. Let $E : \{0, 1\}^k \times \{0, 1\}^L \rightarrow \{0, 1\}^L$ be a block cipher and let $\text{Perm}(L)$ denote the set of all permutations on $\{0, 1\}^L$. Let A be an adversary with access to an oracle and that returns a bit. Then

$$\text{Adv}_F^{\text{PRP}}(A) = \Pr \left[f \stackrel{s}{\leftarrow} E : A^{f(\cdot)} = 1 \right] - \Pr \left[g \stackrel{s}{\leftarrow} \text{Perm}(L) : A^{g(\cdot)} = 1 \right]$$

denotes the PRP-advantage of A in distinguishing a random instance of E from a random permutation. Intuitively, we say that E is a secure PRP, or a secure block cipher, if the PRP-advantages of all adversaries using reasonable resources is small. Modern block ciphers, such as AES [8], are believed to be secure PRPs.

3 The CWC mode of operation

We now describe our new AEAD scheme. Let $\text{BC} : \{0, 1\}^{kl} \times \{0, 1\}^{128} \rightarrow \{0, 1\}^{128}$ be a 128-bit block cipher. Let $\text{tl} \leq 128$ is the desired tag length in bits. Then the CWC mode of operation using BC with tag length tl , $\text{CWC-BC-tl} = (\mathcal{K}, \text{CWC-ENC}, \text{CWC-DEC})$, is defined as follows. The message spaces are:

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

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.

The CWC-BC-tl key generation, encryption, and decryption algorithms are defined as follows:

<pre> Algorithm \mathcal{K} $K \xleftarrow{\\$} \{0, 1\}^{kl}$ Return K Algorithm $\text{CWC-ENC}_K(N, A, M)$ $\sigma \leftarrow \text{CWC-CTR}_K(N, M)$ $\tau \leftarrow \text{CWC-MAC}_K(N, A, \sigma)$ Return $\sigma \parallel \tau$ </pre>	<pre> Algorithm $\text{CWC-DEC}_K(N, A, C)$ If $C < \text{tl}$ then return INVALID Parse C as $\sigma \parallel \tau$ where $\tau = \text{tl}$ If $A \notin \text{AdSp}_{\text{CWC-BC-tl}}$ or $\sigma \notin \text{MsgSp}_{\text{CWC-BC-tl}}$ then return INVALID If $\tau \neq \text{CWC-MAC}_K(N, A, \sigma)$ then return INVALID $M \leftarrow \text{CWC-CTR}_K(N, \sigma)$ Return M </pre>
---	--

The remaining algorithms (CWC-CTR, CWC-MAC, CWC-HASH) are defined below. The CWC-CTR algorithm handles generating the encryption and decryption keystreams, CWC-MAC handles the generation of an authentication tag, and uses CWC-HASH as the underlying universal hash function.

<pre> Algorithm $\text{CWC-CTR}_K(N, M)$ $\alpha \leftarrow \lceil M /128 \rceil$ For $i = 1$ to α do $s_i \leftarrow \text{BC}_K(10^7 \parallel N \parallel \text{tostr}(i, 32))$ $x \leftarrow \text{first } M \text{ bits of } s_1 \parallel s_2 \parallel \dots \parallel s_\alpha$ $\sigma \leftarrow x \oplus M$ Return σ Algorithm $\text{CWC-MAC}_K(N, A, \sigma)$ $R \leftarrow \text{BC}_K(\text{CWC-HASH}_K(A, \sigma))$ $\tau \leftarrow \text{BC}_K(10^7 \parallel N \parallel 0^{32}) \oplus R$ Return first tl bits of τ </pre>	<pre> Algorithm $\text{CWC-HASH}_K(A, \sigma)$ $Z \leftarrow \text{last 127 bits of } \text{BC}_K(110^{126})$ $K_h \leftarrow \text{toint}(Z)$ $l \leftarrow \text{min int such that 96 divides } A \parallel 0^l$ $l' \leftarrow \text{min int such that 96 divides } \sigma \parallel 0^{l'}$ $X \leftarrow A \parallel 0^l \parallel \sigma \parallel 0^{l'}$; $\beta \leftarrow X /96$ Break X into chunks X_1, X_2, \dots, X_β For $i = 1$ to β do $Y_i \leftarrow \text{toint}(X_i)$ $l_\sigma \leftarrow \sigma /8$; $l_A \leftarrow A /8$ $Y_{\beta+1} \leftarrow 2^{64} \cdot l_A + l_\sigma$ $R \leftarrow Y_1 K_h^\beta + \dots + Y_\beta K_h + Y_{\beta+1}$ mod $2^{127} - 1$ Return $\text{tostr}(R, 128)$ </pre>
---	--

4 Theorem statements

The CWC scheme is a provably secure AEAD scheme 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 at least the same assumptions we make, and some modes, such as OCB [23], require the stronger, albeit still reasonable, assumption of super-pseudorandomness.

The specific results for CWC appear as Theorem 1 and Theorem 2 below, and are proven in the full version of this paper [14]. In [14] we also present results for the general CWC construction, from which Theorems 1 and 2 follow.

4.1 Privacy

We first show that if BC is a secure block cipher, then CWC-BC-tl will preserve privacy under chosen-plaintext attacks. For our notion of privacy for AEAD

schemes, we use the strong definition of indistinguishability from [22]. Let $\mathcal{SE} = (\mathcal{K}_e, \mathcal{E}, \mathcal{D})$ be an AEAD scheme with length function $l(\cdot)$. Let $\mathcal{S}(\cdot, \cdot, \cdot)$ 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

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

is the IND \mathcal{S} -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 scheme \mathcal{SE} preserves privacy under chosen plaintext attacks if the IND \mathcal{S} -CPA-advantage of all nonce-respecting adversaries using reasonable resources is small.

Theorem 1. [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*

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

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. ■

Let us elaborate on why Theorem 1 implies that CWC-BC will preserve privacy under chosen-plaintext attacks. 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 are small, then, because of the above equations, $\mathbf{Adv}_{\text{CWC-BC-tl}}^{\text{priv}}(A)$ must also be small as well. I.e., any adversary A using reasonable resources will only be able to break the privacy of CWC-BC-tl with some small probability.

As a concrete example, let us consider limiting the number of applications of CWC-BC-tl between rekeyings to some reasonable value such as $q = 2^{32}$, and let us limit the total number of payload bits between rekeyings to $\mu = 2^{50}$. Then Equation 1 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} .

4.2 Integrity/authenticity

We now present our results showing that if BC is a secure block cipher, then CWC-BC-tl will protect the authenticity of encapsulated data. We use the strong

notion of authenticity for AEAD schemes from [22]. Let $\mathcal{SE} = (\mathcal{K}_e, \mathcal{E}, \mathcal{D})$ be an AEAD scheme. 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 scheme \mathcal{SE} preserves integrity/authenticity if the AUTH-advantage of all nonce-respecting adversaries using reasonable resources is small.

Theorem 2. [Integrity/authenticity of CWC.] *Let CWC-BC-tl be as specified 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*

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

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. ■

Let us elaborate on why Theorem 2 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 2 becomes

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

If we set $q = 2^{32}$ and $\mu = 2^{50}$ as before, 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} .

Remark 1. [Chosen-ciphertext privacy.] Since CWC-BC-tl preserves privacy under chosen-plaintext attacks (Theorem 1) *and* provides integrity (Theorem 2) assuming that BC is a secure pseudorandom permutation, it also provides privacy under chosen-ciphertext attacks under the same assumption about BC. See [4, 22] 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 [23].

5 Design decisions

Finding an appropriate balance between provable security, hardware efficiency, and software efficiency, while simultaneously avoiding existing intellectual property issues, proved to be one the the biggest challenges of this research project. In this section we discuss how our diverse set of goals affected our design decisions.

THE CWC-HASH UNIVERSAL HASH FUNCTION. We found that the best way to simultaneously achieve our parallelizability, hardware, and software goals was to base the authentication portion of CWC on the Carter-Wegman [26] universal hash function approach to message authentication. This is because universal hash functions, and especially the one we created for CWC, can be implemented in a multitude of ways, thus allowing different platforms and applications to implement CWC-HASH in the way most appropriate for them. For example, hardware implementations will like parallelize the computation of CWC-HASH by splitting it into multiple polynomials in K_h^i for some i . In more detail, if the polynomial is

$$Y_1 K_h^\beta + Y_2 K_h^{\beta-1} + Y_3 K_h^{\beta-2} + Y_4 K_h^{\beta-3} + \dots + Y_\beta K_h + Y_{\beta+1} \pmod{2^{127} - 1}.$$

then, setting $i = 2$, and $y = K_h^2 \pmod{2^{127} - 1}$, and assuming β is odd for illustration purposes, we can rewrite the above polynomial as

$$\left(Y_1 y^m + Y_3 y^{m-1} + \dots + Y_\beta \right) x + \left(Y_2 y^m + Y_4 y^{m-1} + \dots + Y_{\beta+1} \right) \pmod{2^{127} - 1},$$

After splitting the polynomial, hardware implementations will then likely compute each polynomial using Horner's rule (e.g., the polynomial $aK_h^{2i} + bK_h^i + c$ would be evaluated as $((a)K_h^i + b)K_h^i + c$). Software implementations on modern CPUs, for which memory is cheap, will likely precompute a number of powers of K_h and evaluate the CWC-HASH polynomial directly, or almost directly, using a hybrid between a precomputation approach and Horner's rule. We consider a number of possible implementation strategies in more detail in Section 6.

CWC-HASH is an instantiation of the classic polynomial universal hash approach to message authentication [26], and is closely related to Bernstein's hash127 [6], which also evaluates a polynomial modulo $2^{127} - 1$. Although hash127 is very fast in software, its structure makes it less suitable for use on high-speed

hardware. In particular, Bernstein's choice of 32-bit coefficients, while great for software implementations with precomputed powers of K_h , means that hardware implementations using Horner's rule will be "wasting work." Specifically, even with 32-bit coefficients, incorporating each new coefficient using Horner's rule will require a 127x127-bit multiply because the accumulated value will be 127 bits long. By defining the CWC-HASH coefficients to be 96-bits long, we increase the performance of Horner's rule implementations by a factor of three. (Of course, we could have gone even further and made the coefficients 126 bits long, but doing so would have required considerable additional complexity to perform bit and byte shifting within the coefficients.) An alternative approach for increasing the performance of a serial implementation of Horner's rule would be to reduce the size of the CWC-HASH subkey K_h to 96 bits. We discuss why we rejected this option in more detail later, but remark here that there are already more efficient strategies than Horner's rule for implementing CWC-HASH in software, and that in a parallelized approach the values K_h^i , $i \geq 2$, will most often be full 127-bit values even if K_h is only 96-bits long.

ON USING A SINGLE KEY. From a security perspective, it would have been perfectly acceptable, and in fact more traditional, to make the CWC-CTR block cipher key and the two CWC-MAC block cipher keys independent. Like others [27, 5], however, we acknowledge that there are several important reasons for sharing keys between the encryption and authentication portions of modes such as CWC. One of the most important reasons is simplicity of key management. Indeed, fetching key material can be a major bottleneck in high-speed hardware, and minimizing key material is thus important. This fact is also why we derive the hash subkey from the block cipher key rather than use an independent hash subkey. We could, of course, have defined a mode that derived a number of essentially independent block cipher and hash keys from a single block cipher key, but doing so would either have required more memory or more computation and, because we have proofs that our construction works, would have been unnecessary.

Sharing the block cipher key in the way described above and deriving the hash subkey from the block cipher key did, however, mean that we had to be very careful with our proofs of security. To facilitate our proofs, we took extra care in our design to ensure that there would never be a collision in the plaintext inputs to the block cipher between the different usages of the block cipher. For example, by defining CWC-HASH to produce a 127-bit value as output, we know that the first application of BC to $CWC-HASH_K(A, \sigma)$ in CWC-MAC will always have its first bit set to 0. To avoid a collision with the input to the keystream generator, the block cipher inputs in CWC-CTR always have the first two bits set to 10. When using the block cipher to create the hash subkey K_h , the first two bits of the input are set to 11.

ON THE CHOICE OF PARAMETERS. Part of this effort involved specifying the appropriate parameters for the CWC encryption mode. Example parameters include the nonce length and the way the nonce is encoded in the input to the block cipher. We chose to fix these parameters for interoperability purposes, but

note that our general approach in [14] does not have these parameters fixed. We chose to set the nonce length to 88 bits in order to handle future IPsec sequence numbers. And we chose to set the block counter length to 32 bits in order to allow CWC to be used with IPsec jumbograms and other large packets. We also chose to use big-endian byte ordering for consistency purposes and to maintain compatibility with McGrew's ICM Internet-Draft [18] and the IETF, which strongly favors big-endian byte-ordering.

HANDLING ARBITRARY BIT-LENGTH MESSAGES. Since we do not believe that many applications will actually require the ability to encrypt arbitrary bit-length messages, we do not define CWC to take arbitrary bit-length messages as input. That said, we did design CWC in such a way that it will be easy to modify the specification to take arbitrary bit-length messages without affecting interoperability with existing implementations when octet-strings are communicated. For example, one could augment the computation of $Y_{\beta+1}$ in CWC-HASH as follows:

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

Of course, a cleaner approach for handling arbitrary bit-length messages would be to compute $l_A \leftarrow |A|$ and $l_\sigma \leftarrow |\sigma|$ in CWC-HASH. We do not define CWC this way because 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.

64-BIT BLOCK CIPHERS. The reader might naturally ask why we did not define CWC for use with block ciphers with 64-bit blocks. The answer is because we are targeting future high-speed cryptographic applications. Nevertheless, we do remark that the general CWC approach in [14] can be instantiated with 64-bit block ciphers.

SOME POSSIBLE ALTERNATIVES. Here we discuss some other possible alternatives to CWC and why we rejected these alternatives. First, as noted earlier, it is possible to improve the performance in some situations by using shorter hash subkeys K_h , say of length 96 bits. Such an alternative will not increase the performance in high-speed hardware implementations that will parallelize the computation of CWC-HASH by evaluating a polynomial in (at least) K_h^2 . A 96-bit hash subkey would have increased Horner's rule performance in software, but would still be comparable in speed to a software-based approach using pre-computed powers of K_h (see Section 6), so reducing the size of K_h to 96 bits would not provide a significant advantage in software either. In [14] we also consider what happens to our provable security bounds when the length of the hash subkey is reduced to less than 96 bits.

There are a number of possible approaches for reducing the number of block cipher applications in the CWC-MAC algorithm by one. For example, 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. One could also use something like $\text{BC}_K(N) + Y_1 K_h^{\beta+2} + \dots + Y_\beta K_h^3 + l_A K_h^2 + l_\sigma K_h \bmod 2^{127} - 1$ as the tag. In [14] we consider these and other alternatives and discuss why we chose to define CWC the way that we did instead of using an option with one fewer block cipher invocation. In the case of the two alternatives mentioned in this paragraph, we

note that we rejected them because we were able to prove better bounds on the security of CWC as currently defined.

Motivated by EAX2 [5], 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. We do comment, however, that we believe the alternatives discussed in this paragraph are still more attractive than EAX because, like CWC but unlike EAX, these alternatives are parallelizable.

We chose not to base the authentication portion of our new mode on XOR-MAC [3] or PMAC [7] because of patent concerns and our software performance requirements and we chose not to base the authentication portion on software-efficient MACs such as HMAC [1] because of our hardware parallelizability requirement.

6 Performance

HARDWARE. Since one of our main goals was to achieve high performance in hardware and, in particular, to provide a solution for future 10 Gbps IPsec (and other) 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 25Kgates 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. With respect to negotiation costs and royalty payments, we note that despite significant demands, to date the relevant parties have not all offered publicly available IP fee schedules. Given this fact, and given today's silicon costs, we believe that the extra silicon for CWC-AES is probably cheaper overall than the negotiation costs and IP fees required for OCB.

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 [9] and use 128-bit AES keys. The Linux compiler was gcc version 3.2.2; the Windows compiler was Visual Studio 6.0. To be fair, we note that OCB does run at about twice the speeds given in Table 1.

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 [16]). Returning to the speed of CWC-HASH, for reference we note that Bernstein's related hash127 [6] 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 Section 5, which discusses why we use 96-bit coefficients). We also note that the performance of CWC-HASH will increase dramatically on 64-bit architectures with larger multiplies; an initial implementation on a G5

using 64-bit integer operations runs at around 6 cpb (when running the G5 in 32-bit mode, the hash function runs at around 15 cpb).

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 Section 5, 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 Section 5, 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. (We discuss elsewhere why we did not make the hash subkey 96-bits, which could have sped up a serial Horner's rule implementation.) The advantage with precomputation was first observed by Bernstein in the context of hash127 [6].

The above description of the precomputation approach assumed that if the polynomial is $Y_1 K_h^{\gamma-1} + \dots + Y_{\gamma-1} K_h + Y_\gamma$ (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 $Y_1 K_h^{n-1} + \dots + Y_{n-1} K_h + Y_n$), then multiplies the intermediate result by K_h^n (to get $Y_1 K_h^{2n-1} + \dots + Y_{n-1} K_h^{n+1} + Y_n K_h^n$). After that, one can continue processing data with the same precomputed values (to get $Y_1 K_h^{2n-1} + \dots + Y_{2n-1} K_h + Y_{2n}$), 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 implemen-

tation 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, which explains why implementations of CWC-HASH for 64-bit architectures are much faster.

EXPLOITING THE PARALLELISM OF SOME INSTRUCTION SETS. On most 32-bit 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, 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 multiples 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 [6].

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 multipliers 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 C 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 on the same machine.

Finally, it may be possible to further improve the performance of CWC-HASH. For example, literature from the gaming community [11] 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.

7 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.

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, R. Canetti, and H. Krawczyk. Keying hash functions for message authentication. In N. Kobnitz, editor, *CRYPTO '96*, volume 1109 of *LNCS*, pages 1–15. Springer-Verlag, Aug. 1996.
2. M. Bellare, A. Desai, E. Jorjipii, and P. Rogaway. A concrete security treatment of symmetric encryption. In *Proc. of the 38th FOCS*, pages 394–403. IEEE Computer Society Press, 1997.
3. M. Bellare, R. Gu erin, and P. Rogaway. XOR MACs: New methods for message authentication using finite pseudorandom functions. In D. Coppersmith, editor, *CRYPTO '95*, volume 963 of *LNCS*, pages 15–28. Springer-Verlag, Aug. 1995.
4. M. Bellare and C. Namprempre. Authenticated encryption: Relations among notions and analysis of the generic composition paradigm. In T. Okamoto, editor, *ASIACRYPT 2000*, volume 1976 of *LNCS*, pages 531–545. Springer-Verlag, Dec. 2000.
5. M. Bellare, P. Rogaway, and D. Wagner. The EAX mode of operation. In W. Meier and B. Roy, editors, *FSE 2004*, LNCS. Springer-Verlag, 2004.
6. D. Bernstein. Floating-point arithmetic and message authentication, 2000. Available at <http://cr.y.p.to/papers.html#hash127>.
7. J. Black and P. Rogaway. A block-cipher mode of operation for parallelizable message authentication. In L. Knudsen, editor, *EUROCRYPT 2002*, volume 2332 of *LNCS*. Springer-Verlag, 2002.

8. J. Daemen and V. Rijmen. *The Design of Rijndael*. Springer-Verlag, 2002.
9. B. Gladman. AES and combined encryption/authentication modes, 2003. Available at <http://fp.gladman.plus.com/AES/index.htm>.
10. V. Gligor and P. Donescu. Fast encryption and authentication: XCBC encryption and XECB authentication modes. In M. Matsui, editor, *FSE 2001*, LNCS. Springer-Verlag, 2001.
11. 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>.
12. C. Jutla. Encryption modes with almost free message integrity. In B. Pfitzmann, editor, *EUROCRYPT 2001*, volume 2045 of LNCS, pages 529–544. Springer-Verlag, May 2001.
13. J. Katz and M. Yung. Unforgeable encryption and chosen ciphertext secure modes of operation. In B. Schneier, editor, *FSE 2000*, volume 1978 of LNCS, pages 284–299. Springer-Verlag, Apr. 2000.
14. T. Kohno, J. Viega, and D. Whiting. CWC: A high-performance conventional authenticated encryption mode, 2003. Full version of this paper, available at <http://eprint.iacr.org/2003/106/>.
15. H. Krawczyk. The order of encryption and authentication for protecting communications (or: How secure is SSL?). In J. Kilian, editor, *CRYPTO 2001*, volume 2139 of LNCS, pages 310–331. Springer-Verlag, Aug. 2001.
16. H. Lipmaa. AES/Rijndael: speed, 2003. Available at <http://www.tcs.hut.fi/~helger/aes/rijndael.html>.
17. M. Luby and C. Rackoff. How to construct pseudorandom permutations from pseudorandom functions. *SIAM J. Computation*, 17(2), Apr. 1988.
18. D. McGrew. Integer counter mode, Oct. 2002. Available at <http://www.ietf.org/internet-drafts/draft-irtf-cfrg-icm-01.txt>.
19. 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>.
20. D. McGrew. The universal security transform, Oct. 2002. Available at <http://www.ietf.org/internet-drafts/draft-irtf-cfrg-ust-01.txt>.
21. W. Nevelsteen and B. Preneel. In J. Stern, editor, *EUROCRYPT '99*, volume 1592 of LNCS, pages 24–41. Springer-Verlag, 1999.
22. P. Rogaway. Authenticated encryption with associated data. In *Proc. of the 9th CCS*, Nov. 2002.
23. P. Rogaway, M. Bellare, J. Black, and T. Krovetz. OCB: A block-cipher mode of operation for efficient authenticated encryption. In *Proc. of the 8th CCS*, pages 196–205. ACM Press, 2001.
24. P. Rogaway and D. Wagner. A critique of CCM, Apr. 2003. Available at <http://eprint.iacr.org/2003/070/>.
25. V. Shoup. On fast and provably secure message authentication based on universal hashing. In N. Kobitz, editor, *CRYPTO '96*, volume 1109 of LNCS, pages 313–328. Springer-Verlag, Aug. 1996.
26. 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.
27. 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.