The Software Performance of Authenticated-Encryption Modes

Ted Krovetz[1] and Phillip Rogaway[2]

[1] Computer Science, California State University, Sacramento, CA 95819 USA
[2] Computer Science, University of California, Davis, CA 95616 USA

Abstract. We study software performance of authenticated-encryption modes CCM, GCM, and OCB. Across a variety of platforms, we find OCB to be substantially faster than either alternative. For example, on an Intel i5 ("Clarkdale") processor, good implementations of CCM, GCM, and OCB encrypt at around 4.2 cpb, 3.7 cpb, and 1.5 cpb, while CTR mode requires about 1.3 cpb. Still we find room for algorithmic improvements to OCB, showing how to trim one blockcipher call (most of the time, assuming a counter-based nonce) and reduce latency. Our findings contrast with those of McGrew and Viega (2004), who claimed similar performance for GCM and OCB.

Keywords: authenticated encryption, cryptographic standards, encryption speed, modes of operation, CCM, GCM, OCB.

1 Introduction

BACKGROUND. Over the past few years, considerable effort has been spent constructing schemes for *authenticated encryption* (AE). One reason is recognition of the fact that a scheme that delivers both privacy and authenticity may be more efficient than the straightforward amalgamation of separate privacy and authenticity techniques. A second reason is the realization that an AE scheme is less likely to be incorrectly used than a privacy-only encryption scheme.

While other possibilities exist, it is natural to build AE schemes from blockciphers, employing some *mode of operation*. There are two approaches. In a *composed* ("two-pass") AE scheme one conjoins essentially separate privacy and authenticity modes. For example, one might apply CTR-mode encryption and then compute some version of the CBC MAC. Alternatively, in an *integrated* ("one-pass") AE scheme the parts of the mechanism responsible for privacy and for authenticity are tightly coupled. Such schemes emerged around a decade ago, with the work of Jutla [21], Katz and Yung [23], and Gligor and Donescu [11].

Integrated AE schemes were invented to improve performance of composed ones, but it has not been clear if they do. In the only comparative study to date [32], McGrew and Viega found that their composed scheme, GCM, was about as fast as, and sometimes faster than, the integrated scheme OCB [36] (hereinafter OCB1, to distinguish it from a subsequent variant we'll call OCB2 [35]). After McGrew and Viega's 2004 paper, no subsequent performance study

scheme	ref	date	ty	high-level description
EtM	[1]	2000	C	Encrypt-then-MAC (and other) generic comp. schemes
RPC	[23]	2000	I	Insert counters and sentinels in blocks, then ECB
IAPM	[21]	2001	I	Seminal integrated scheme. Also IACBC
XCBC	[11]	2001	I	Concurrent with Jutla's work. Also XECB
✓ OCB1	[36]	2001	I	Optimized design similar to IAPM
TAE	[29]	2002	I	Recasts OCB1 using a tweakable blockcipher
✓ CCM	[40]	2002	C	CTR encryption + CBC MAC
CWC	[24]	2004	C	CTR encryption + $GF(2^{127}-1)$-based CW MAC
✓ GCM	[32]	2004	C	CTR encryption + $GF(2^{128})$-based CW MAC
EAX	[2]	2004	C	CTR encryption + CMAC, a cleaned-up CCM
✓ OCB2	[35]	2004	I	OCB1 with AD and alleged speed improvements
CCFB	[30]	2005	I	Similar to RPC [23], but with chaining
CHM	[18]	2006	C	Beyond-birthday-bound security
SIV	[37]	2006	C	Deterministic/misuse-resistant AE
CIP	[17]	2008	C	Beyond-birthday-bound security
HBS	[20]	2009	C	Deterministic AE. Single key
BTM	[19]	2009	C	Deterministic AE. Single key, no blockcipher inverse
✓ OCB3	new	2010	I	Refines the prior versions of OCB

Fig. 1. Authenticated-encryption schemes built from a blockcipher. Checks indicate schemes included in our performance study. Column **ty** (type) specifies if the scheme is integrated (I) or composed (C). Schemes EtM, CCM, GCM, EAX, OCB2, and SIV are in various standards (ISO 19772, NIST 800-38C and 800-38D, RFC 5297).

was ever published. This is unfortunate, as there seems to have been a major problem with their work: reference implementations were compared against optimized ones, and none of the results are repeatable due to the use of proprietary code. In the meantime, CCM and GCM have become quite important to cryptographic practice. For example, CCM underlies modern WiFi (802.11i) security, while GCM is supported in IPsec and TLS.

McGrew and Viega identified two performance issues in the design of OCB1. First, the mode uses $m + 2$ blockcipher calls to encrypt a message of $m = \lceil |M|/128 \rceil$ blocks. In contrast, GCM makes do with $m + 1$ blockcipher calls. Second, OCB1 twice needs one AES result before another AES computation can proceed. Both in hardware and in software, this can degrade performance. Beyond these facts, existing integrated modes cannot exploit the "locality" of counters in CTR mode—that high-order bits of successive input blocks are usually unchanged, an observation first exploited, for software speed, by Hongjun Wu [4].

Given all of these concerns, maybe GCM really is faster than OCB—and, more generally, maybe composed schemes are the fastest way to go. The existence of extremely high-speed MACs supports this possibility [3,5,25].

CONTRIBUTIONS. We begin by refining the definition of OCB to address the performance concerns just described. When the provided nonce is a counter, the mode that we call OCB3 shaves off one AES encipherment per message encrypted about 98% of the time. In saying that the nonce is a counter we mean that, in a given session, its top portion stays fixed, while, with each successive message, the bottom portion gets bumped by one. This is the approach recommended in RFC 5116 [31, Section 3.2] and, we believe, the customary way to use an AE scheme. We do not introduce something like a $GF(2^{128})$ multiply to compensate for the usually-eliminated blockcipher call, and no significant penalty is paid, compared to OCB1, if the provided nonce is not a counter (one just fails to save the blockcipher call). We go on to eliminate the latency that used to occur when computing the "checksum" and processing the AD (associated data).

Next we study the relative software performance of CCM, GCM, and the different versions of OCB. We employ the fastest publicly available code for Intel x86, both with and without Intel's new instructions for accelerating AES and GCM. For other platforms—ARM, PowerPC, and SPARC—we use a refined and popular library, OpenSSL. We test the encryption speed on messages of every byte length from 1 byte to 1 Kbyte, plus selected lengths beyond. The OCB code is entirely in C, except for a few lines of inline assembly on ARM and compiler intrinsics to access byteswap, trailing-zero count, and SSE/AltiVec functionality.

We find that, across message lengths and platforms, OCB, in any variant, is well faster than CCM and GCM. While the performance improvements from our refining OCB are certainly measurable, those differences are comparatively small. Contrary to McGrew and Viega's findings, the speed differences we observe between GCM and OCB1 are large and favor OCB1.

As an example of our experimental findings, for 4 KB messages on an Intel i5 ("Clarkdale") processor, we clock CCM at 4.17 CPU cycles per byte (cpb), GCM at 3.73 cpb, OCB1 at 1.48 cpb, OCB2 at 1.80 cpb, and OCB3 at 1.48 cpb. As a baseline, CTR mode runs at 1.27 cpb. See Fig. 2. These implementations exploit the processor's AES New Instructions (AES-NI), including "carryless multiplication" for GCM. The OCB3 *authentication overhead*—the time the mode spends in excess of the time to encrypt with CTR—is about 0.2 cpb, and the difference between OCB and GCM overhead is about a factor of 10. Even written in C, our OCB implementations provide, on this platform, the fastest reported times for AE.

The means for refining OCB are not complex, but it took much work to understand what optimization would and would not help. First we wanted to arrange that nonces agreeing on all but their last few bits be processed using the same blockcipher call. To accomplish this in a way that minimizes runtime state and key-setup costs, we introduce a new hash-function family, a *stretch-then-shift* xor-universal hash. The latency reductions are achieved quite differently, by changes in how the mode defines and operates on the Checksum.
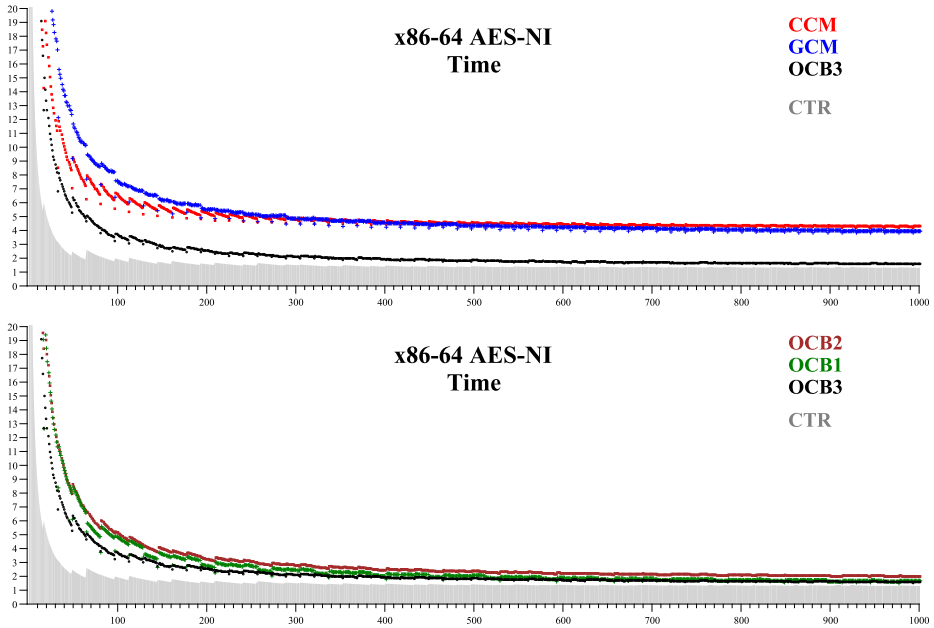
Fig. 2. Top: Performance of CCM, GCM, and OCB3 on an x86 with AES-NI. The
x-coordinate is the message length, in bytes; the y-coordinate is the measured number
of cycles per byte. From top-to-bottom on the right-hand side, the curves are for
CCM, GCM, and OCB3. The shaded region shows the time for CTR mode. This and
subsequent graphs are best viewed in color. **Bottom:** Performance of OCB variants
on an x86 with AES-NI. From top-to-bottom, the curves are for OCB2, OCB1, and
OCB3. The shaded region shows the time for CTR mode.

One surprising finding is that, on almost all platforms, OCB2 is slightly slower
than OCB1. To explain, recall that most integrated schemes involve computing
an *offset* for each blockcipher call. With OCB1, each offset is computed by xoring
a key-dependent value, an approach going back to Jutla [21]; with OCB2, each
offset is computed by a "doubling" in $GF(2^{128})$. The former approach turns out
to be faster.

During our work we investigated novel ways to realize a maximal period,
software-efficient, 128-bit LFSR; such constructions can also be used to make
the needed offsets. A computer-aided search identified constructions like sending
$A \parallel B \parallel C \parallel D$ to $C \parallel D \parallel B \parallel ((A \lll 1) \oplus (A \ggg 1) \oplus (D \lll 15))$ (where A, B, C, D are
32 bits); see Appendix A. While very fast, such maps are still slower than xoring
a precomputed value. Our findings thus concretize Chakraborty and Sarkar's
suggestion [6] to improve OCB using a fast, 128-bit, word-oriented LFSR—but,
in the end, we conclude that the idea doesn't really help. Of course software-
optimized 128-bit LFSRs may have other applications.

All code and data used in this paper, plus a collection of clickable tables and
graphs, are available from the second author's webpage.

2   The Mode OCB3

PRELIMINARIES. We begin with a few basics. A *blockcipher* is a deterministic algorithm $E\colon \mathcal{K} \times \{0,1\}^n \to \{0,1\}^n$ where $\mathcal{K}$ is a finite set and $n \geq 1$ is a number, the *key space* and *blocklength*. We require $E_K(\cdot) = E(K, \cdot)$ be a permutation for all $K \in \mathcal{K}$. Let $D = E^{-1}$ be the map from $\mathcal{K} \times \{0,1\}^n$ to $\{0,1\}^n$ defined by $D_K(Y) = D(K, Y)$ being the unique point X such that $E_K(X) = Y$.

Following recent formalizations [1,23,34,36], a scheme for (nonce-based) authenticated encryption (with associated-data) is a three-tuple $\Pi = (\mathcal{K}, \mathcal{E}, \mathcal{D})$. The *key space* $\mathcal{K}$ is a finite, nonempty set. The *encryption algorithm* $\mathcal{E}$ takes in a *key* $K \in \mathcal{K}$, a *nonce* $N \in \mathcal{N} \subseteq \{0,1\}^*$, a *plaintext* $M \in \mathcal{M} \subseteq \{0,1\}^*$, and *associated data* $A \in \mathcal{A} \subseteq \{0,1\}^*$. It returns, deterministically, either a ciphertext $C = \mathcal{E}_K^{N,\,A}(M) \in \mathcal{C} \subseteq \{0,1\}^*$ or the distinguished value INVALID. Sets $\mathcal{N}$, $\mathcal{M}$, $\mathcal{C}$, and $\mathcal{A}$ are called the *nonce space*, *message space*, *ciphertext space*, and *AD space* of Π. The *decryption algorithm* $\mathcal{D}$ takes a tuple $(K, N, A, C) \in \mathcal{K} \times \mathcal{N} \times \mathcal{A} \times \mathcal{C}$ and returns, deterministically, either INVALID or a string $M = \mathcal{D}_K^{N,\,A}(C) \in \mathcal{M} \subseteq \{0,1\}^*$. We require that $\mathcal{D}_K^{N,\,A}(C) = M$ for any string $C = \mathcal{E}_K^{N,\,A}(M)$ and that $\mathcal{E}$ and $\mathcal{D}$ return INVALID if provided an input outside of $\mathcal{K} \times \mathcal{N} \times \mathcal{A} \times \mathcal{M}$ or $\mathcal{K} \times \mathcal{N} \times \mathcal{A} \times \mathcal{C}$, respectively. We require $|\mathcal{E}_K^{N,A}(M)| = |\mathcal{E}_K^{N,A}(M')|$ when the encryptions are strings and $|M| = |M'|$. If this value is always $|M| + \tau$ we call τ the *tag length* of the scheme.

DEFINITION OF OCB3. Fix a blockcipher $E\colon \mathcal{K} \times \{0,1\}^{128} \to \{0,1\}^{128}$ and a tag length $\tau \in [0\,..\,128]$. In Fig. 3 we define the AE scheme $\Pi = \text{OCB3}[E, \tau] = (\mathcal{K}, \mathcal{E}, \mathcal{D})$. The nonce space $\mathcal{N}$ is the set of all binary strings with fewer than 128 bits.[1] The message space $\mathcal{M}$ and AD-space $\mathcal{A}$ are all binary strings. The ciphertext space $\mathcal{C}$ is the set of all strings whose length is at least τ bits. Fig. 3's procedure Setup is implicitly run on or before the first call to $\mathcal{E}$ or $\mathcal{D}$. The variables it defines are understood to be global. In the protocol definition we write $\text{ntz}(i)$ for the number of trailing zeros in the binary representation of positive integer i (eg, $\text{ntz}(1) = \text{ntz}(3) = 0$, $\text{ntz}(4) = 2$), we write $\text{msb}(X)$ for the first (most significant) bit of X, we write $A \wedge B$ for the bitwise-and of A and B, and we write $A \ll i$ for the shift of A by i positions to the left (maintaining string length, leftmost bits falling off, zero-bits entering at the right). At lines 111 and 311 we regard Bottom as a number instead of a string.

DESIGN RATIONALE. We now explain some of the design choices made for OCB3. While not a large departure from OCB1 or OCB2, the refinements do help.

Trimming a blockcipher call. OCB1 and OCB2 took $m + 2$ blockcipher calls to encrypt an m-block string M: one to map the nonce N into an initial offset Δ; one for each block of M; one to encipher the final Checksum. The first of these is easy to eliminate if one is willing to replace the $E_K(N)$ computation by, say,

---

[1] In practice one would either restrict nonces to byte strings of 1–15 bytes, or else demand that nonces have a fixed length, say exactly 12-bytes. Under RFC 5116, a conforming AE scheme *should* use a 12-byte nonce.

```
101  algorithm 𝓔_K^{N A}(M)                          301  algorithm 𝓓_K^{N A}(𝒞)
102  if |N| ≥ 128 return INVALID                      302  if |N| ≥ 128 or |𝒞| < τ return INVALID
103  M_1 ··· M_m M_* ← M where each                   303  C_1 ··· C_m C_* T ← 𝒞 where each
104      |M_i| = 128 and |M_*| < 128                  304      |C_i| = 128 and |C_*| < 128 and |T| = τ
105  Checksum ← 0^{128};  C ← ε                        305  Checksum ← 0^{128};  M ← ε
106  Nonce ← 0^{127−|N|} 1 N                           306  Nonce ← 0^{127−|N|} 1 N
107  Top ← Nonce ∧ 1^{122} 0^6                         307  Top ← Nonce ∧ 1^{122} 0^6
108  Bottom ← Nonce ∧ 0^{122} 1^6                      308  Bottom ← Nonce ∧ 0^{122} 1^6
109  Ktop ← E_K(Top)                                  309  Ktop ← E_K(Top)
110  Stretch ← Ktop ‖ (Ktop ⊕ (Ktop≪8))              310  Stretch ← Ktop ‖ (Ktop ⊕ (Ktop≪8))
111  Δ ← (Stretch ≪ Bottom)[1..128]                   311  Δ ← (Stretch ≪ Bottom)[1..128]
112  for i ← 1 to m do                                312  for i ← 1 to m do
113      Δ ← Δ ⊕ L[ntz(i)]                            313      Δ ← Δ ⊕ L[ntz(i)]
114      C ← E_K(M_i ⊕ Δ) ⊕ Δ                         314      M ← D_K(C_i ⊕ Δ) ⊕ Δ
115      Checksum ← Checksum ⊕ M_i                    315      Checksum ← Checksum ⊕ M_i
116  if M_* ≠ ε then                                  316  if C_* ≠ ε then
117      Δ ← Δ ⊕ L_*                                  317      Δ ← Δ ⊕ L_*
118      Pad ← E_K(Δ)                                 318      Pad ← E_K(Δ)
119      C ← M_* ⊕ Pad[1..|M_*|]                      319      M ← M_* ⊕ C_* ⊕ Pad[1..|C_*|]
120      Checksum ← Checksum ⊕ M_* 10^*               320      Checksum ← Checksum ⊕ M_* 10^*
121  Δ ← Δ ⊕ L_$                                      321  Δ ← Δ ⊕ L_$
122  Final ← E_K(Checksum ⊕ Δ)                        322  Final ← E_K(Checksum ⊕ Δ)
123  Auth ← Hash_K(A)                                 323  Auth ← Hash_K(A)
124  Tag ← Final ⊕ Auth                               324  Tag ← Final ⊕ Auth
125  T ← Tag[1..τ]                                    325  T' ← Tag[1..τ]
126  return C ‖ T                                     326  if T = T' then return M
                                                      327          else return INVALID


201  algorithm Setup(K)                               401  algorithm Hash_K(A)
202  L_* ← E_K(0^{128})                               402  A_1 ··· A_m A_* ← A where each
203  L_$ ← double(L_*)                                403      |A_i| = 128 and |A_*| < 128
204  L[0] ← double(L_$)                               404  Sum ← 0^{128}
205  for i ← 1, 2, ··· do L[i] ← double(L[i−1])       405  Δ ← 0^{128}
206  return                                           406  for i ← 1 to m do
                                                      407      Δ ← Δ ⊕ L[ntz(i)]
                                                      408      Sum ← Sum ⊕ E_K(A_i ⊕ Δ)
                                                      409  if A_* ≠ ε then
                                                      410      Δ ← Δ ⊕ L_*
211  algorithm double(X)                              411      Sum ← Sum ⊕ E_K(A_* 10^* ⊕ Δ)
212  return (X≪1) ⊕ (msb(X) · 135)                    412  return Sum
```

Fig. 3. Definition of OCB3[E,τ]. Here $E \colon \mathcal{K} \times \{0,1\}^{128} \to \{0,1\}^n$ is a blockcipher and $\tau \in [0..128]$ is the tag length. Algorithms $\mathcal{E}$ and $\mathcal{D}$ are called with arguments $K \in \mathcal{K}$, $N \in \{0,1\}^{\leq 127}$, and $M, C \in \{0,1\}^*$.

$K_1 \cdot N$, the product in $\mathrm{GF}(2^{128})$ of nonce N and a variant K_1 of K. The idea has been known since Halevi [14]. But such a change would necessitate implementing a $\mathrm{GF}(2^{128})$ multiply for just this one step. Absent hardware support, one would need substantial precomputation and enlarged internal state to see any savings; not a net win. We therefore compute the initial offset Δ using a different xor-universal hash function: $\Delta = H_K(N) = (\text{Stretch} \ll \text{Bottom})[1..128]$ where Bottom is the last six bits of N and the $(128+64)$-bit string Stretch is made by a process involving enciphering N with its last six bits zeroed out. This *stretch-then-shift* hash will be proven xor-universal in Section 4.1. Its use ensures that, when the nonce N is a counter, the initial offset Δ can be computed without a new blockcipher call $63/64 \approx 98\%$ of the time. In this way we reduce cost from

$m + 2$ blockcipher calls to an amortized $m+1.016$ blockcipher calls, plus tiny added time for the hash.

Reduced latency. Assume the message being encrypted is not a multiple of 128 bits; there is a final block M_* having 1–127 bits. In prior versions of OCB one would need to wait on the penultimate blockcipher call to compute the Checksum and, from it, the final blockcipher call. Not only might this result in pipeline stalls [32], but if the blockcipher's efficient implementation needs a long string to ECB, then the lack of parallelizability translates to extra work. For example, Käsper and Schwabe's bit-sliced AES [22] ECB-encrypts eight AES blocks in a single shot. Using this in OCB1 or OCB2 would result in enciphering 24 blocks to encrypt a 100-byte string—three times more than what "ought" to be needed—since twice one must wait on AES output to form the next AES input. In OCB3 we restructure the algorithm so that the Checksum never depends on any ciphertext. Concretely, Checksum $= M_1 \oplus M_2 \oplus M_{m-1} \oplus M_m 10^*$ for a short final block, and Checksum $= M_1 \oplus M_2 \oplus M_{m-1} \oplus M_m$ for a full final block. The fact that you can get the same Checksum for distinct final blocks is addressed by using different offsets in these two cases.

Incrementing offsets. In OCB1, each noninitial offset is computed from the prior one by xoring some key-derived value; the ith offset is constructed by $\Delta \leftarrow \Delta \oplus L[\mathrm{ntz}(i)]$. In OCB2, each noninitial offset is computed from the prior one by multiplying it, again in $\mathrm{GF}(2^{128})$, by a constant: $\Delta \leftarrow (\Delta \ll 1) \oplus (\mathrm{msb}(\Delta) \cdot 135)$, an operation that has been called *doubling*. Not having to go to memory or attend to the index i, doubling was thought to be faster than the first method. In our experiments, it is not. While doubling can be coded in five Intel x86-64 assembly instructions, it still runs more slowly. In some settings, doubling loses big: it is expensive on 32-bit machines, and some compilers do poorly at turning C/C++ code for doubling into machine code that exploits the available instructions. On Intel x86, the 128-bit SSE registers lack the ability to be efficiently shifted one position to the left. Finally, the doubling operation is not endian neutral: if we must create a bit pattern in memory to match the sequence generated by doubling (and AES implementations generally do expect their inputs to live in memory) we will effectively favor big-endian architectures. We can trade this bias for a little-endian one by redefining double() to include a byteswap. But one is still favoring one endian convention over the other, and not just at key-setup time. See Appendix A for some of the alternatives to repeated doubling that we considered.

Further design issues. Unlike OCB1 and OCB2, each 128-bit block of plaintext is now processed in the same way whether or not it is the final 128 bits. This change facilitates implementing a clean incremental API, since one is able to output each 128-bit chunk of ciphertext after receiving the corresponding chunk of plaintext, even if it is not yet known if the plaintext is complete.

All AD blocks can now be processed concurrently; in OCB2, the penultimate block's output was needed to compute the final block's input, potentially creating

pipeline stalls or inefficient use of a blockcipher's multi-block ECB interface. Also, each 128-bit block of AD is treated the same way if it is or isn't the message's end, simplifying the incremental provisioning of AD.

We expect the vast majority of processors running OCB3 will be little-endian; still, the mode's definition does nothing to favor this convention. The issue arises each time "register oriented" and "memory oriented" values interact. These are the same on big-endian machines, but are opposite on little-endian ones. One could, therefore, favor little-endian machines by building into the algorithm byte swaps that mimic those that would occur naturally each time memory and register oriented data interact. We experimentally adapted our implementation to do this but found that it made very little performance difference. This is due, first, to good byte reversal facilities on most modern processors (eg, `pshufb` can reverse 16 bytes on our x86 in a single cycle). It is further due to the fact that OCB3's table-based approach for incrementing offsets allows for the table to be endian-adjusted at key setup, removing most endian-dependency on subsequent encryption or decryption calls. Since it makes little difference to performance, and big-endian specifications are conceptually easier, OCB3 does not make any gestures toward little-endian orientation.

A low-level choice where OCB and GCM part ways is in the representation of field points. In GCM the polynomial $a_{127}\mathtt{x}^{127} + \cdots a_1\mathtt{x} + a_0$ corresponds to string $a_0 \ldots a_{127}$ rather than $a_{127} \ldots a_0$. McGrew and Viega call this the little-endian representation, but, in fact, this choice has nothing to do with endianness. The usual convention on machines of all kinds is that the msb is the left-most bit of any register. Because of this, GCM's "reflected-bit" convention can result in extra work to be performed even on Intel chips having instructions specifically intended for accelerating GCM [12,13]. Among the advantages of following the msb-first convention is that a left shift by one can be implemented by adding a register to itself, an operation often faster than a logical shift.

SECURITY OF OCB3. First we provide our definitions. Let $\Pi = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ be an AE scheme. Given an adversary (algorithm) $\mathcal{A}$, we let $\mathbf{Adv}_{\Pi}^{\mathrm{priv}}(\mathcal{A}) = \Pr[K \overset{\$}{\leftarrow} \mathcal{K} : \mathcal{A}^{\mathcal{E}_K(\cdot, \cdot, \cdot)} \Rightarrow 1] - \Pr[\mathcal{A}^{\$(\cdot, \cdot, \cdot)} \Rightarrow 1]$ where queries of $\$(N, A, M)$ return a uniformly random string of length $|\mathcal{E}_K^{N,A}(M)|$. We demand that $\mathcal{A}$ never asks two queries with the same first component (the N-value), that it never ask a query outside of $\mathcal{N} \times \mathcal{A} \times \mathcal{M}$, and that it never repeats a query. Next we define authenticity. For that, let $\mathbf{Adv}_{\Pi}^{\mathrm{auth}}(\mathcal{A}) = \Pr[K \overset{\$}{\leftarrow} \mathcal{K} : \mathcal{A}^{\mathcal{E}_K(\cdot, \cdot, \cdot)} \text{ forges}]$ where we say that the adversary *forges* if it outputs a value $(N, A, C) \in \mathcal{N} \times \mathcal{A} \times \mathcal{C}$ such that $\mathcal{D}_K^{N,A}(C) \neq \text{INVALID}$ yet there was no prior query (N, A, M') that returned C. We demand that $\mathcal{A}$ never asks two queries with the same first component (the N-value), never asks a query outside of $\mathcal{N} \times \mathcal{A} \times \mathcal{M}$, and never repeats a query.

When $E : \mathcal{K} \times \{0,1\}^n \rightarrow \{0,1\}^n$ is a blockcipher define $\mathbf{Adv}_E^{\pm\mathrm{prp}}(\mathcal{A}) = \Pr[\mathcal{A}^{E_K(\cdot),\, E_K^{-1}(\cdot)} \Rightarrow 1] - \Pr[\mathcal{A}^{\pi(\cdot),\, \pi^{-1}(\cdot)} \Rightarrow 1]$ where K is chosen uniform from $\mathcal{K}$ and $\pi(\cdot)$ is a uniform permutation on $\{0,1\}^n$. Define $\mathbf{Adv}_E^{\mathrm{prp}}(\mathcal{A}) = \Pr[\mathcal{A}^{E_K(\cdot)} \Rightarrow 1] - \Pr[\mathcal{A}^{\pi(\cdot)} \Rightarrow 1]$ by removing the decryption oracle. The *ideal* blockcipher of

blocksize n is the blockcipher $\mathrm{Bloc}[n] \colon \mathcal{K} \times \{0,1\}^n \to \{0,1\}^n$ where each key K names a distinct permutation.

The security of OCB3 is given by the following theorem. We give the result in its information-theoretic form. Passing to the complexity-theoretic setting, where the idealized blockcipher $\mathrm{Bloc}[n]$ is replaced by a conventional blockcipher secure as a strong-PRP, is standard.

Theorem 1. *Fix $n = 128$, $\tau \in [0\,..\,n]$, and let $\Pi = \mathrm{OCB3}[E, \tau]$ where $E = \mathrm{Bloc}[n]$ is the ideal blockcipher on n bits. If $\mathcal{A}$ asks encryption queries that entail σ total blockcipher calls, then $\mathbf{Adv}_{\Pi}^{\mathrm{priv}}(\mathcal{A}) \leq 6\,\sigma^2/2^n$. Alternatively, if $\mathcal{A}$ asks encryption queries then makes a forgery attempt that together entail σ total blockcipher calls, then $\mathbf{Adv}_{\Pi}^{\mathrm{auth}}(\mathcal{A}) \leq 6\sigma^2/2^n + (2^{n-\tau})/(2^n - 1)$.* $\hfill\square$

When we speak of the number of blockcipher calls entailed we are adding up the (rounded-up) blocklength for all the different strings output by the adversary and adding in $q + 2$ (q =number of queries), to upper-bound blockcipher calls for computing L_* and the initial Δ values. Main elements of the proof are described in Section 4; see the full paper for further details [26].

3   Experimental Results

SCOPE AND CODEBASE. We empirically study the software performance of OCB3, and compare this with state-of-the-art implementations of GCM, which delivers the fastest previously reported AE times. Both modes are further compared against CTR, the fastest privacy-only mode, which makes a good baseline for answering how much extra one pays for authentication. Finally, we consider CCM, the first NIST-approved AE scheme, and also OCB1 and OCB2, which are benchmarked to show how the evolution of OCB has affected performance.

Intensively optimized implementations of CTR and GCM are publicly available for the x86. Käsper and Schwabe hold the speed record for 64-bit code with no AES-NI, reporting peak rates of 7.6 and 10.7 CPU cycles per byte (cpb) for CTR and GCM [22]. With AES-NI, developmental versions of OpenSSL achieve 1.3 cpb for CTR [33] and 3.3 cpb for GCM.[2] These various results use different x86 chips and timing mechanisms. Here we use the Käsper-Schwabe AES, CTR, and GCM, the OpenSSL CTR, CCM, and GCM, augment the collection with new code for OCB, and compare performance on a single x86 and use a common timing mechanism, giving the fairest comparison to date.

The only non-proprietary, architecture-specific non-x86 implementations for AES and GCM that we could find are those in OpenSSL. Although these implementations are hand-tuned assembly, they are designed to be timing-attack resistant, and are therefore somewhat slow. This does not make comparisons with them irrelevant. OCB is timing-attack resistant too (assuming the underlying blockcipher is), making the playing field level. We adopt the OpenSSL

---

[2] Andy Polyakov, personal communication, August 27, 2010. The fastest published AES-NI time for GCM is 3.5 cpb on 8KB messages, from Gueron and Kounavis [13].

implementations for non-x86 comparisons and emphasize that timing-resistant implementations are being compared, not versions written for ultimate speed.

The OCB1 and OCB2 implementations are modifications of our OCB3 implementation, and therefore are similarly optimized. These implementations are in C, calling out to AES. No doubt further performance improvements can be obtained by rewriting the OCB code in assembly.

Hardware and software environments. We selected five representative instruction-set architectures: (1) 32-bit x86, (2) 64-bit x86, (3) 32-bit ARM, (4) 64-bit PowerPC, and (5) 64-bit SPARC. Collectively, these architectures dominate the workstation, server, and portable computing marketplace. The x86 processor used for both 32- and 64-bit tests is an Intel Core i5-650 "Clarkdale" supporting the AES-NI instructions. The ARM is a Cortex-A8. The PowerPC is a 970fx. The SPARC is an UltraSPARC IIIcu. Each runs Debian Linux 6.0 with kernel 2.6.35 and GCC 4.5.1. Compilation is done with `-O3` optimization, `-mcpu` or `-march` set according to the host processor, and `-m64` to force 64-bit compilation when needed.

Testing methodology. The number of CPU cycles needed to encrypt a message is divided by the length of the message to arrive at the cost per byte to encrypt messages of that length. This is done for every message length from 1 to 1024 bytes, as well as 1500 and 4096 bytes. So as not to have performance results overly influenced by the memory subsystem of a host computer, we arrange for all code and data to be in level-1 cache before timing begins. Two timing strategies are used: C clock and x86 time-stamp counter. In the clock version, the ANSI C `clock()` function is called before and after repeatedly encrypting the same message, on sequential nonces, for a little more than one second. The clock difference determines how many CPU cycles were spent on average per processed byte. This method is highly portable, but it is time-consuming when collecting an entire dataset. On x86 machines there is a "time-stamp counter" (TSC) that increments once per CPU cycle. To capture the average cost of encryption—including the more expensive OCB3 encryptions that happen once every 64 calls—the TSC is used to time encryption of the same message 64 times on successive counter-based nonces. The TSC method is not portable, working only on x86, but is fast. Both methods have their potential drawbacks. The clock method depends on the hardware having a high-resolution timer and the OS doing a good job of returning the time used only by the targeted process. The TSC read instruction might be executed out of order, in some cases it has high latency, and it continues counting when other processes run.[3] In the end, we found that both timing methods give similar results. For example, in the

---

[3] To lessen these problems we read the TSC once before and after encrypting the same message 65 times, then read the TSC once before and after encrypting the same message once more. Subtracting the second timing from the first gives us the cost for encrypting the message 64 times, and mitigates the out-of-order and latency problems. To avoid including context-switches, we run experiments multiple times and keep only the median timing.

x86-64 AES-NI					x86-32 AES-NI					x86-64 Käsper-Schwabe			
Mode	T_{4K}	T_{IPI}	Size	Init	Mode	T_{4K}	T_{IPI}	Size	Init	Mode	T_{4K}	T_{IPI}	Size	Init
CCM	4.17	4.57	512	265	CCM	4.18	4.70	512	274	GCM	22.4	26.7	1456	3780
GCM	3.73	4.53	656	337	GCM	3.88	4.79	656	365	GCM-8K	10.9	15.2	9648	2560
OCB1	1.48	2.08	544	251	OCB1	1.60	2.22	544	276	OCB1	8.28	13.4	3008	3390
OCB2	1.80	2.41	448	185	OCB2	1.79	2.42	448	197	OCB2	8.55	13.6	2912	3350
OCB3	1.48	1.87	624	253	OCB3	1.59	2.04	624	270	OCB3	8.05	9.24	3088	3480
CTR	1.27	1.37	244	115	CTR	1.39	1.52	244	130	CTR	7.74	8.98	1424	1180

| ARM Cortex-A8					PowerPC 970					UltraSPARC III			
Mode	T_{4K}	T_{IPI}	Size	Init	Mode	T_{4K}	T_{IPI}	Size	Init	Mode	T_{4K}	T_{IPI}	Size	Init
CCM	51.3	53.7	512	1390	CCM	75.7	77.8	512	1510	CCM	49.4	51.7	512	1280
GCM	50.8	53.9	656	1180	GCM	53.5	56.2	656	1030	GCM	39.3	41.5	656	904
OCB1	29.3	31.5	672	1920	OCB1	38.2	41.0	672	2180	OCB1	25.5	27.7	672	1720
OCB2	28.5	31.8	576	1810	OCB2	38.1	41.1	576	2110	OCB2	24.8	27.0	576	1700
OCB3	28.9	30.9	784	1890	OCB3	37.5	39.6	784	2240	OCB3	25.0	26.5	784	1730
CTR	25.4	25.9	244	236	CTR	37.5	37.8	244	309	CTR	24.1	24.4	244	213
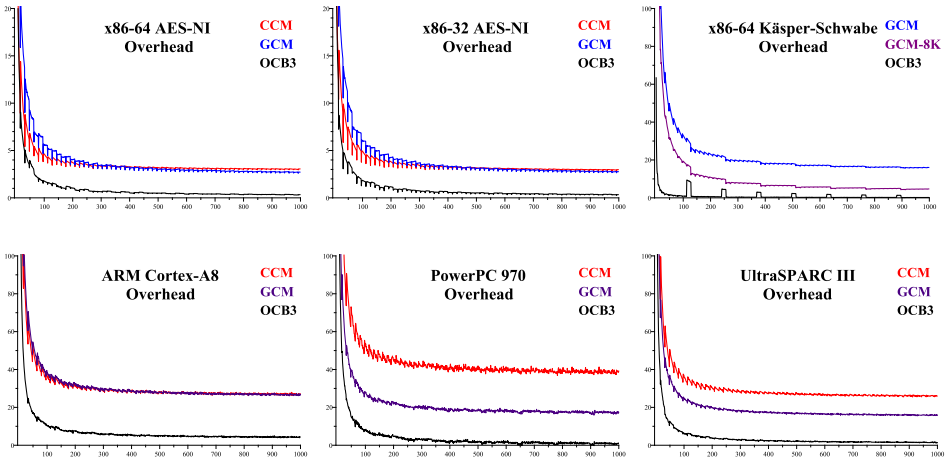
Fig. 4. Empirical performance of AE modes. For each architecture we give time to encrypt 4KB messages (in CPU cycles per byte), time to encrypt a weighted basket of message lengths (IPI, also in cpb), size of the implementation's context (in bytes), and time to initialize key-dependent values (in CPU cycles). Next we graph the same data, subtracting the CTR time and dropping the curves for OCB1 and OCB2, which may be visually close to that of OCB3. The CCM and GCM curves are visually hard to distinguish in the x86-64 AES NI, x86-32 AES NI, and ARM Cortex-A8 graphs.

eighteen x86 test runs done for this paper, the Internet Performance Index values computed by the two methods varied by no more than 0.05 cpb 10 times, no more than 0.10 cpb 15 times, and no more than 0.20 cpb all 18 times.

RESULTS. Summary findings are presented in Figs. 2 and 4. On all architectures and message lengths, OCB3 is significantly faster than GCM and CCM. Except on very short messages, it is nearly as fast as CTR. On x86, GCM's most competitive platform, OCB3's authentication overhead (its cost beyond CTR encryption) is 4–16%, with or without AES-NI, on both an Internet Performance Index (IPI)[4] and 4KB message length basis. In all our tests, CCM never has IPI or 4KB rates better than GCM, coming close only when small registers make GCM's multiplications expensive, or AES-NI instructions speed CCM's block encipherments. Results are similar on other architectures. The overhead of OCB3 does not exceed 12% that of GCM or CCM on PowerPC or SPARC, or 18% on ARM, when looking at either IPI or 4KB message encryption rates.

To see why OCB3 does so well, consider that there are four phases in OCB3 encryption: initial offset generation, encryption of full blocks, encryption of a partial final block (if there is one), and tag generation. On all but the shortest messages, full-block processing dominates overall cost per byte. Here OCB3, and OCB1, are particularly efficient. An unrolled implementation of, say, four blocks per iteration, will have, as overhead on top of the four blockcipher calls and the reads and writes associated to them: 16 xor operations (each on 16-byte words), 1 ntz computation, and 1 table lookup of a 16-byte value. On x86, summing the latencies of these 18 operations—which ignores the potential for instruction-level parallelism (ILP)—the operations require 23 cycles, or 0.36 cpb. In reality, on 64-bit x64 using AES-NI, we see CTR taking 1.27 cpb on 4KB messages while OCB3 uses 1.48, an overhead of 0.21 cpb, the savings coming from the ILP.

Short messages are optimized for too. When there is little or no full-block processing, it is the other three phases of encryption that determine performance. One gets a sense of the cost of these by looking at the cost to encrypt a single byte. On x86, OpenSSL's AES-NI based CTR implementation does this in 86 cycles, while CCM, GCM, and OCB3 use 257, 354, and 249 cycles, respectively. CCM remains competitive with OCB3 only for very short strings. On 64-bit x86 without AES-NI, using Käsper-Schwabe's bit-sliced AES that processes eight blocks at once, OCB3's performance lead is much greater, as its two blockcipher calls can be computed concurrently, unlike CCM and GCM. In this scenario, single-byte encryption rates for CCM, GCM, OCB3, CTR are 2600, 2230, 1080, 1010 cycles. On the other three architectures we see the following single-byte encryption times for (CCM, GCM, OCB3; CTR): ARM (1770, 1950, 1190; 460), PowerPC (2520, 1860, 1450; 309), and SPARC (1730, 1520, 1770; 467).

With hardware support making AES very cheap, authentication overhead becomes more prominent. AES-NI instructions enable AES-128 throughput of around 20 cycles per block. VIA's `xcrypt` assembly instruction is capable of 10

---

[4] The IPI is a weighted average of timings for messages of 44 bytes (5%), 552 bytes (15%), 576 bytes (20%), and 1500 bytes (60%) [32]. It is based on Internet backbone studies from 1998. We do not suggest that the IPI reflects a contemporary, real-world distribution of message lengths, only that it is useful to have *some* metric that attends to shorter messages and those that are not a multiple of 16 bytes. Any metric of this sort will be somewhat arbitrary in its definition.

cycles per block on long ECB sequences [39]. Speeds like these can make authentication overhead more expensive than encryption. With the Käsper-Schwabe code (no AES-NI), for example, on an IPI basis, OCB3 overhead is only 3% of encryption cost, but under AES-NI it rises to 27%. Likewise, GCM overhead rises from 41% to 70%. One might think CCM would do well using AES-NI since its overhead is mostly blockcipher calls, but its use of (serial) CBC for authentication reduces AES throughput to around 60 cycles per block, causing authentication overhead of about 70%.[5]

As expected, OCB1 and OCB3 long-message performance is the same due to having identical full-block processing. OCB2 is slower on long messages on all tested platforms but SPARC (computing ntz is slow on SPARC). With a counter-based nonce, OCB3 computes its initial encryption offset using a few bitwise shifts of a cached value rather than generating it with a blockcipher as both OCB1 and OCB2 do. This results in significantly improved average performance for encryption of short messages. The overall effect is that on an IPI basis on, say, 64-bit x86 using AES-NI, OCB3's authentication overhead is only 65% of that for OCB1 and only 40% of that for OCB2. When the provided nonce is *not* a counter, OCB3 performance is, in most of our test environments, indistinguishable from that of OCB1.

4   Proof of Security for OCB3

We describe three elements in the proof of OCB3's security: (1) the new xor-universal hash function it employs; (2) the definition and proof for a simple TBC (tweakable blockcipher) based generalization of OCB3; and (3) the proof that the particular TBC used by OCB3 is good.

4.1   Stretch-then-Shift Universal Hash

A new hash function H underlies the mapping of the low-order bits of the nonce to a 128-bit string (lines 108, 110, and 111 of Fig. 3). While an off-the-shelf hash would have worked alright, we were able to do better for this step. We start with the needed definitions.

DEFINITION. Let $\mathcal{K}$ be a finite set and let $H : \mathcal{K} \times \mathcal{X} \to \{0,1\}^n$ be a function. We say that H is *strongly xor-universal* if for all distinct $x, x' \in \mathcal{X}$ we have that $H_K(x) \oplus H_K(x')$ is uniformly distributed in $\{0,1\}^n$ and, also, $H_K(x)$ is uniformly distributed in $\{0,1\}^n$ for all $x \in \mathcal{X}$. The first requirement is the usual definition for H being *xor-universal*; the second we call *universal-1*.

---

[5] Intel released their Sandy Bridge microarchitecture January 2011, too late for a thorough update of this paper. Sandy Bridge increases both AES throughput and latency. Under Sandy Bridge, OCB and CTR will be substantially faster (likely under 1.0 cpb on long messages) because their work is dominated by parallel AES invocations. GCM will be just a little faster because most of its time is spent in authentication, which does not benefit from Sandy Bridge. CCM will be slower because longer latencies negatively affect CBC authentication.

THE TECHNIQUE. We aim to construct strongly xor-universal hash-functions $H \colon \mathcal{K} \times \mathcal{X} \to \{0,1\}^n$ where $\mathcal{K} = \{0,1\}^{128}$, $\mathcal{X} = [0 \mathinner{..} \mathrm{domSize} - 1]$, and $n = 128$. We want domSize to be at least some modest-size number, say domSize ≥ 64, and intend that computing $H_K(x)$ be almost as fast as doing a table lookup. Fast computation of H should not require any large table, nor the preprocessing of K. Our desire for extreme speed in the absence of preprocessing and big tables rules out methods based on $\mathrm{GF}(2^{128})$ multiplication, the obvious first attempt.

The method we propose is to stretch the key K into a longer string $stretch(K)$, and then extract its bits $x+1$ to $x+128$. Symbolically, $H_K(x) = (stretch(K))[x+1 \mathinner{..} x + 128]$ where $S[a \mathinner{..} b]$ denotes bits a through b of S, indexing beginning with 1. Equivalently, $H_K(x) = (stretch(K) \ll x)[1 \mathinner{..} 128]$. We call this a *stretch-then-shift* hash.

How to stretch K? It seems natural to have $stretch(K)$ begin with K, so let's assume that $stretch(K) = K \parallel s(K)$ for some function s. It's easy to see that $s(K) = K$ and $s(K) \ll c$ won't work, but $s(K) = K \oplus (K \ll c)$, for some constant c, looks plausible for accommodating modest-sized domain. We now demonstrate that, for well-chosen c, this function does the job.

ANALYSIS. To review, we are considering $H_K^c(x) = (\mathrm{Stretch} \ll x)[1 \mathinner{..} 128]$ where $\mathrm{Stretch} = stretch(K) = K \parallel (K \oplus (K \ll c))$ and $c \in [0 \mathinner{..} 127]$. We'd like to know the maximal value of domSize for which $H_K(x)$ is xor-universal on the domain $\mathcal{X} = [0 \mathinner{..} \mathrm{domSize}(c) - 1]$. This can be calculated by a computer program, as we now explain. Fix c and consider the 256×128 entry matrix $A = \binom{I}{J}$ where I is the 128×128 identity matrix and J is the 128×128-bit matrix for which $J_{ij} = 1$ iff $j = i$ or $j = i + c$. Let A_i denote the 128×128 submatrix of A that includes only A's rows i to $i + 127$. Then $H_K^c(x) = A_{x+1}K$, the product in $\mathrm{GF}(2)$ of the matrix A_{i+1} and the column vector K. Let $B_{i,j} = A_i + A_j$ be the indicated 128×128 matrix, the matrix sum over $\mathrm{GF}(2)$. We would like to ensure that, for arbitrary $0 \leq i < j < \mathrm{domSize}(c)$ and a uniform $K \in \{0,1\}^{128}$ that the 128-bit string $H_K^c(i) + H_K^c(j)$ is uniform—which is to say that $A_{i+1}K + A_{j+1}K = (A_{i+1} + A_{j+1})K = B_{i+1,j+1}K$ is uniform. This will be true if and only if $B_{i,j}$ is invertible in $\mathrm{GF}(2)$ for all $1 \leq i < j \leq \mathrm{domSize}(c)$. Thus domSize$(c)$ can be computed as the largest number domSize(j) such that $B_{i,j}$ is full rank, over $\mathrm{GF}(2)$, for all $1 \leq i < j \leq \mathrm{domSize}(j)$. Recalling the universal-1 property we also demand that A_i have full rank for all $1 \leq i \leq \mathrm{domSize}(c)$. Now for any c, the number of matrices $A_{i,j}$ to consider is at most 2^{13}, and finding the rank in $\mathrm{GF}(2)$ of that many 128×128 matrices is a feasible calculation.

Our results are tabulated in Fig. 5. The most interesting cases are H^5 and H^8, which are strongly xor-universal on $\mathcal{X} = [0 \mathinner{..} 123]$ and $\mathcal{X} = [0 \mathinner{..} 84]$, respectively. We offer no explanation for why these functions do well and various other H^c do not. As both H^5 and H^8 work on $[0 \mathinner{..} 63]$ we select the latter map for use in OCB3 and single out the following result:

Lemma 1. *Let $H_K(x)$ be the first 128 bits of $\mathrm{Stretch} \ll x$ where $\mathrm{Stretch} = K \parallel (K \oplus (K \ll 8))$, $|K| = 128$, $x \in [0 \mathinner{..} 63]$. Then H is strongly xor-universal.* $\square$

c	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
$\mathrm{domSize}(c)$	3	15	7	3	124	7	3	85	120	3	118	63	3	31	63	3	7	31	3	7

Fig. 5. Stretch-then-shift hash. Largest $\mathcal{X} = [0 .. \mathrm{domSize}(c) - 1]$ such that $H_K^c(x) = (\mathrm{Stretch}(K) \lll x)[1 .. 128]$ is strongly xor-universal when $c \in [1 .. 16]$, $K \in \{0,1\}^{128}$, $x \in \mathcal{X}$, and $\mathrm{Stretch}(K) = K \parallel (K \oplus (K \lll c))$.

EFFICIENCY. On 64-bit computers, assuming $K \parallel (K \oplus (K \lll 8))$ is precomputed and in memory, the value of $H_K(x)$ can be computed by three memory loads and two multiprecision shifts, requiring fewer than ten cycles on most architectures. If only K is in memory then the first 64 bits of $K \oplus (K \lll 8)$ can be computed with three additional assembly instructions. In the absence of a preprocessed table or special hardware-support, a method based on $\mathrm{GF}(2^{128})$ multiplies would not fare nearly as well.

Computing successive H_K^c values requires a single extended-precision shift, making stretch-then-shift a reasonable approach for incrementing offsets. Unfortunately, it is not endian-neutral.

4.2   The TBC-Based Generalization of OCB3

Following the insight of Liskov, Rivest, and Wagner [29], OCB3 can be understood as an instantiation of an AE scheme that depends on a *tweakable blockcipher* (TBC). This is a deterministic algorithm $\widetilde{E}$ having signature $\widetilde{E} : \mathcal{K} \times \mathcal{T} \times \{0,1\}^n \to \{0,1\}^n$ where $\mathcal{K}$ and $\mathcal{T}$ are sets and $n \geq 1$ is a number—the *key space*, *tweak space*, and *blocklength*, respectively. We require $\widetilde{E}_K^T(\cdot) = \widetilde{E}(K, T, \cdot)$ be a permutation for all $K \in \mathcal{K}$ and $T \in \mathcal{T}$. Write $\widetilde{D} = \widetilde{E}^{-1}$ for the map from $\mathcal{K} \times \mathcal{T} \times \{0,1\}^n$ to $\{0,1\}^n$ defined by $\widetilde{D}_K^T(Y) = \widetilde{D}(K, T, Y)$ being the unique X such that $\widetilde{E}_K^T(X) = Y$. The *ideal* TBC for a tweak set $\mathcal{T}$ and block-size n is the blockcipher $\mathrm{Bloc}[\mathcal{T}, n] \colon \mathcal{K} \times \mathcal{T} \times \{0,1\}^n \to \{0,1\}^n$ where the keys name distinct permutations for each tweak T. For $\mathcal{T} = \mathcal{T}^{\pm} \cup \mathcal{T}^+$, $\mathcal{T}^{\pm} \cap \mathcal{T}^+ = \emptyset$, let $\mathbf{Adv}_{\widetilde{E}}^{\mathrm{prp}[\mathcal{T}^{\pm}]}(\mathcal{A}) = \Pr[K \xleftarrow{\$} \mathcal{K} \colon \mathcal{A}^{\widetilde{E}_K(\cdot,\cdot),\, \widetilde{D}_K(\cdot,\cdot)} \Rightarrow 1] - \Pr[\mathcal{A}^{\pi(\cdot,\cdot),\, \pi^{-1}(\cdot,\cdot)} \Rightarrow 1]$ where π is chosen uniformly from $\mathrm{Bloc}[\mathcal{T}, n]$ and adversary $\mathcal{A}$ is only allowed to ask decryption queries (T, Y) with $T \in \mathcal{T}^{\pm}$. Write $\mathbf{Adv}_{\widetilde{E}}^{\pm \mathrm{prp}}(\mathcal{A})$ for $\mathbf{Adv}_{\widetilde{E}}^{\mathrm{prp}[\mathcal{T}]}(\mathcal{A})$ and $\mathbf{Adv}_{\widetilde{E}}^{\mathrm{prp}}(\mathcal{A})$ for $\mathbf{Adv}_{\widetilde{E}}^{\mathrm{prp}[\emptyset]}(\mathcal{A})$. Our definition unifies PRP and strong-PRP security, allowing forward queries for all tweaks and backwards queries for those in $\mathcal{T}^{\pm}$. A conventional blockcipher can be regarded as a TBC with a singleton tweak space.

THE ΘCB3 SCHEME. Fix an arbitrary set of nonces $\mathcal{N}$; for concreteness, say $\mathcal{N} = \{0,1\}^{<128}$. Define from this set the corresponding tweak space $\mathcal{T}$ by

$$\mathcal{T} = \mathcal{N} \times \mathbb{N}_1 \ \cup \ \mathcal{N} \times \mathbb{N}_0 \times \{*\} \ \cup \ \mathcal{N} \times \mathbb{N}_0 \times \{\$\} \ \cup \ \mathcal{N} \times \mathbb{N}_0 \times \{*\$\} \ \cup \ \mathbb{N}_1 \ \cup \ \mathbb{N}_0 \times \{*\}$$

where $\mathbb{N}_1$ and $\mathbb{N}_0$ are the positive and nonnegative integers, respectively. Tweaks, it can be seen, are of six mutually exclusive "types." Tweaks of the first type are

```
101  algorithm 𝓔ₖ^{N,A}(M)                    201  algorithm 𝓓ₖ^{N,A}(𝓒)
102  if N ∉ 𝒩 then return INVALID             202  if N ∉ 𝒩 or |𝓒| < τ then return INVALID
103  M₁ ··· Mₘ M∗ ← M where each              203  C₁ ··· Cₘ C∗ T ← 𝓒 where each
104      |Mᵢ| = n and |M∗| < n                204      |Cᵢ| = n, |C∗| < n, and |T| = τ
105  Checksum ← 0ⁿ,  C∗ ← ε                    205  Checksum ← 0ⁿ,  M∗ ← ε
106  for i ← 1 to m do                         206  for i ← 1 to m do
107      Cᵢ ← Ẽₖ^{N i}(Mᵢ)                     207      Mᵢ ← D̃ₖ^{N i}(Cᵢ)
108      Checksum ← Checksum ⊕ Mᵢ              208      Checksum ← Checksum ⊕ Mᵢ
109  if M∗ = ε then Final ← Ẽₖ^{N m $}(Checksum)  209  if C∗ = ε then Final ← Ẽₖ^{N m $}(Checksum)
111  else  Pad ← Ẽₖ^{N m ∗}(0ⁿ)                211  else  Pad ← Ẽₖ^{N m ∗}(0ⁿ)
111      C∗ ← M∗ ⊕ Pad[1..|M∗|]               211      M∗ ← C∗ ⊕ Pad[1..|C∗|]
112      Checksum ← Checksum ⊕ M∗ 10∗         212      Checksum ← Checksum ⊕ M∗ 10∗
113      Final ← Ẽₖ^{N m ∗ $}(Checksum)       213      Final ← Ẽₖ^{N m ∗ $}(Checksum)
114  Auth ← Hashₖ(A)                           214  Auth ← Hashₖ(A)
115  Tag ← Final ⊕ Auth                        215  Tag ← Final ⊕ Auth
116  T ← Tag[1..τ]                             216  T′ ← Tag[1..τ]
117  return C₁ ··· Cₘ C∗ ‖ T                   217  if T = T′ then return M₁ ··· Mₘ M∗
                                               218             else return INVALID

301  algorithm Hashₖ(A)
302  Sum ← 0ⁿ
303  A₁ ··· Aₘ A∗ ← A for |Aᵢ| = n, |A∗| < n
304  for i ← 1 to m do
305      Sum ← Sum ⊕ Ẽₖ^{i}(Aᵢ)
306  if A∗ ≠ ε then
307      Sum ← Sum ⊕ Ẽₖ^{m ∗}(A∗ 10∗)
308  return Sum
```

Fig. 6. Definition of $\Theta\mathrm{CB3}[\widetilde{E},\tau]$. Here $\widetilde{E}\colon \mathcal{N} \times \mathcal{T} \times \{0,1\}^n \to \{0,1\}^n$ is a tweakable blockcipher and $\tau \in [0..n]$ is the tag length. We have that $\mathrm{OCB3}[E,\tau] = \Theta\mathrm{CB3}[\widetilde{E},\tau]$ for an appropriately chosen $\widetilde{E}$.

in the set $\mathcal{T}^{\pm} = \mathcal{N} \times \mathbb{N}_1$. Omitting parenthesis and commas when writing tweaks, TBC calls will look like $\widetilde{E}_K^{N\,i}(X)$, $\widetilde{E}_K^{N\,i\,*}(X)$, $\widetilde{E}_K^{N\,i\,\$}(X)$, $\widetilde{E}_K^{N\,i\,*\,\$}(X)$, $\widetilde{E}_K^{i}(X)$, or $\widetilde{E}_K^{i\,*}(X)$. Now given such a TBC $\widetilde{E}\colon \mathcal{K} \times \mathcal{T} \times \{0,1\}^n \to \{0,1\}^n$ and given a tag length $\tau \in [0..n]$, we construct the AE scheme $\Pi = \Theta\mathrm{CB3}[\widetilde{E},\tau] = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ as defined in Fig. 6. The scheme's nonce space is $\mathcal{N}$, the message space is $\mathcal{M} = \{0,1\}^*$, the AD space is $\mathcal{A} = \{0,1\}^*$, and the ciphertext space is $\mathcal{C} = \{0,1\}^{\geq \tau}$. The scheme is illustrated in Fig. 7.

We now describe the security of $\Theta\mathrm{CB3}$ when using an ideal TBC. The proof is in the full paper [26].

Lemma 2. *Let* $\Pi = \Theta\mathrm{CB3}[\widetilde{E},\tau]$ *where* $\widetilde{E} = \mathrm{Bloc}[\mathcal{T},n]\colon \mathcal{K} \times \mathcal{T} \times \{0,1\}^n \to \{0,1\}^n$ *is ideal. Let* $\mathcal{A}$ *be an adversary. Then* $\mathbf{Adv}_{\Pi}^{\mathrm{priv}}(\mathcal{A}) = 0$ *and* $\mathbf{Adv}_{\Pi}^{\mathrm{auth}}(\mathcal{A}) \leq (2^{n-\tau})/(2^n - 1)$. $\qquad\square$

4.3 Instantiating the TBC

Continuing to assume that $n = 128$ and $\mathcal{N} = \{0,1\}^{<n}$, map each blockcipher $E\colon \mathcal{K} \times \{0,1\}^n \to \{0,1\}^n$ to the TBC $\widetilde{E} = \mathrm{Tw}[E]$, $\widetilde{E}\colon \mathcal{K} \times \mathcal{T} \times \{0,1\}^n \to \{0,1\}^n$, where $\mathcal{T} = \mathcal{N} \times \mathbb{N}_1 \ \cup\ \mathcal{N} \times \mathbb{N}_0 \times \{*\} \ \cup\ \mathcal{N} \times \mathbb{N}_0 \times \{\$\} \ \cup\ \mathcal{N} \times \mathbb{N}_0 \times \{*\$\} \ \cup\ \mathbb{N}_1 \ \cup\ \mathbb{N}_0 \times \{*\}$ by the construction of Fig. 8. There, multiplication is in $\mathrm{GF}(2^{128})$ using the irreducible polynomial $\mathtt{x}^{128} + \mathtt{x}^7 + \mathtt{x}^7 + \mathtt{x}^2 + \mathtt{x} + 1$. We use the standard facts on the Gray code sequence $a\colon \mathbb{N}_0 \to \mathbb{N}_0$ that it is a permutation and $0 \leq a(i) \leq 2i$. It follows that coefficients $\Lambda = \{\lambda_i,\ \lambda_j^*,\ \lambda_j^\$,\ \lambda_j^{*\$}\colon\ 1 \leq i \leq 2^{120},$
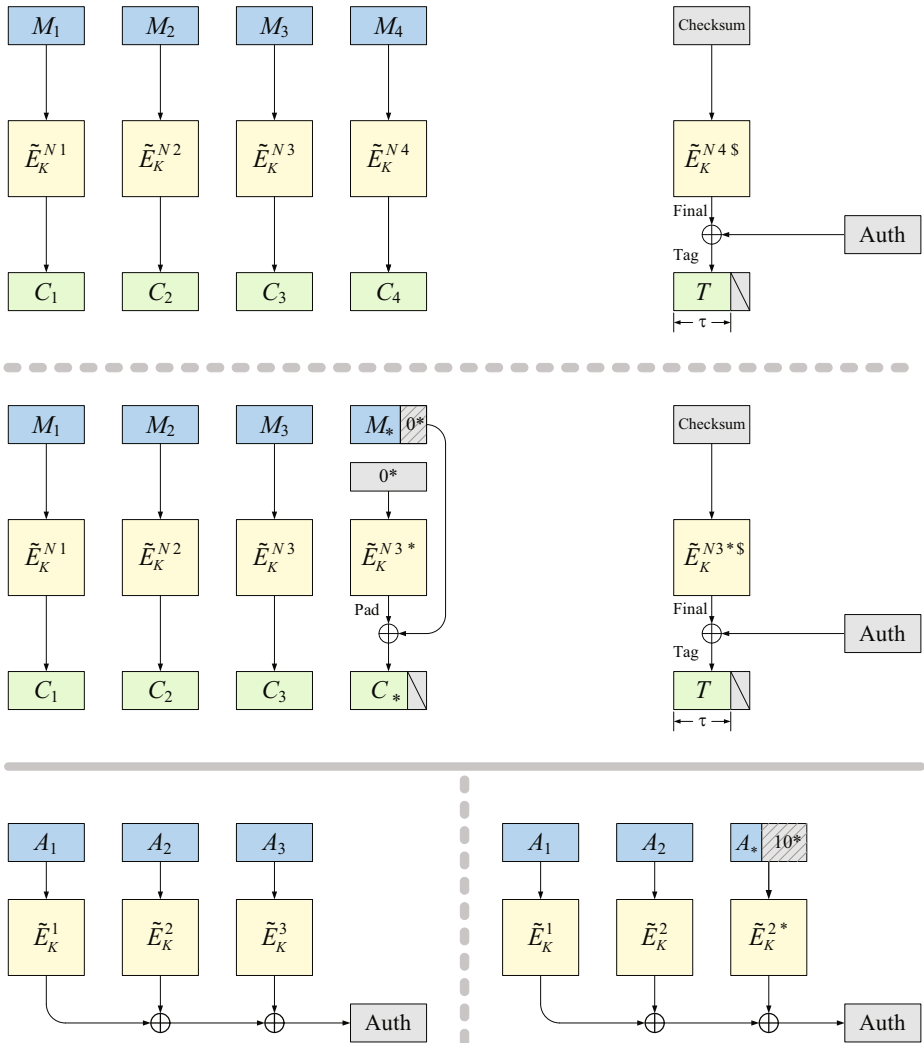
Fig. 7. Illustration of ΘCB3. The scheme depends on tweakable blockcipher $\widetilde{E} \colon \mathcal{N} \times \mathcal{T} \times \{0,1\}^n \to \{0,1\}^n$ and tag length $\tau \in [0 \mathinner{\ldotp\ldotp} n]$. The top figure shows the treatment of a message M having a full final block ($|M_4| = n$) (Checksum $= M_1 \oplus M_2 \oplus M_3 \oplus M_4$) while the middle picture shows the treatment of a message M having a short final block ($1 \le |M_*| < n$) (Checksum $= M_1 \oplus M_2 \oplus M_3 \oplus M_* 10^*$). The bottom-left picture shows the processing of a three-block AD; on bottom-right, an AD with two full blocks and a short one. Algorithm OCB3$[E, \tau]$ coincides with ΘCB3$[\widetilde{E}, \tau]$ for a particular TBC $\widetilde{E} = \mathrm{Tw}[E]$ constructed from E.

$$\widetilde{E}_K^{\,N\,i}(X) = E_K(X \oplus \Delta) \oplus \Delta \quad \text{with} \quad \Delta = \text{Initial} \,\oplus\, \lambda_i\, L \qquad \text{for } i \geq 1$$

$$\widetilde{E}_K^{\,N\,i\,*}(X) = E_K(X \oplus \Delta) \qquad \text{with} \quad \Delta = \text{Initial} \,\oplus\, \lambda_i^*\, L \qquad \text{for } i \geq 0$$

$$\widetilde{E}_K^{\,N\,i\,\$}(X) = E_K(X \oplus \Delta) \qquad \text{with} \quad \Delta = \text{Initial} \,\oplus\, \lambda_i^\$\, L \qquad \text{for } i \geq 0$$

$$\widetilde{E}_K^{\,N\,i\,*\,\$}(X) = E_K(X \oplus \Delta) \qquad \text{with} \quad \Delta = \text{Initial} \,\oplus\, \lambda_i^{*\$}\, L \qquad \text{for } i \geq 0$$

$$\widetilde{E}_K^{\,i}(X) = E_K(X \oplus \Delta) \qquad \text{with} \quad \Delta = \lambda_i\, L \qquad\qquad\quad \text{for } i \geq 1$$

$$\widetilde{E}_K^{\,i\,*}(X) = E_K(X \oplus \Delta) \qquad \text{with} \quad \Delta = \lambda_i^*\, L \qquad\qquad\quad \text{for } i \geq 0$$

where

$$\text{Nonce} = 0^{127-|N|}\, 1\, N$$
$$\text{Top} = \text{Nonce} \wedge 1^{122}\, 0^6$$
$$\text{Bottom} = \text{Nonce} \wedge 0^{122}\, 1^6$$
$$\text{Ktop} = E_K(\text{Top})$$
$$\text{Stretch} = \text{Ktop} \parallel (\text{Ktop} \oplus (\text{Ktop} \ll 8))$$
$$\text{Initial} = (\text{Stretch} \ll \text{Bottom})[1..128]$$

$$L = E_K(0^{128})$$
$$\lambda_i = 4\, a(i)$$
$$\lambda_i^* = 4\, a(i) + 1$$
$$\lambda_i^\$ = 4\, a(i) + 2$$
$$\lambda_i^{*\$} = 4\, a(i) + 3$$
$$a(0) = 0 \quad //\text{Grey code seq } 0, 1, 3, 2, 6, \dots$$
$$a(i) = a(i-1) \oplus 2^{\text{ntz}(i)} \text{ if } i \geq 1$$

Fig. 8. Definition of $\widetilde{E} = \text{Tw}[E]$, **the tweakable blockcipher built from** E

$0 \leq j \leq 2^{120}\}$ are distinct and nonzero points of $\text{GF}(2^{128})$. The reader can check that $\text{OCB3}[E, \tau] = \Theta\text{CB3}[\text{Tw}[E], \tau]$.

SECURITY OF THE CONSTRUCTED TBC. We show that $\widetilde{E} = \text{Tw}[E]$ is a good TBC if E is a good blockcipher. In formalizing this, forward queries may be asked throughout $\mathcal{T}$, but backwards queries must be of the form $\widetilde{E}_K^{\,N\,i}$.

Lemma 3. *Let* $n = 128$ *and let* $E = \text{Bloc}[n]$ *be the ideal blockcipher on* n *bits. Let* $\widetilde{E} = \text{Tw}[E]$, *the tweak space being* $\mathcal{T}$, *and let* $\mathcal{T}^\pm = \mathcal{N} \times \mathbb{N}_1$. *Let* $\mathcal{A}$ *be an adversary that asks at most* q *queries, non employing an* i-*value in excess of* 2^{120}. *Then* $\mathbf{Adv}_{\widetilde{E}}^{\text{prp}[\mathcal{T}^\pm]}(\mathcal{A}) \leq 6q^2/2^n$. $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\square$

The proof is in the full version [26]. Combining it with Lemma 2 gives Theorem 1.

Acknowledgments

References

1. Bellare, M., Namprempre, C.: Authenticated encryption: relations among notions and analysis of the generic composition paradigm. J. Cryptology 21(4), 469–491 (2008); Earlier version in ASIACRYPT 2000
2. Bellare, M., Rogaway, P., Wagner, D.: The EAX mode of operation. In: Roy, B., Meier, W. (eds.) FSE 2004. LNCS, vol. 3017, pp. 389–407. Springer, Heidelberg (2004)

324 T. Krovetz and P. Rogaway

3. Bernstein, D.: The Poly1305-AES message-authentication code. In: Gilbert, H., Handschuh, H. (eds.) FSE 2005. LNCS, vol. 3557, pp. 32–49. Springer, Heidelberg (2005)
4. Bernstein, D.J., Schwabe, P.: New AES software speed records. In: Chowdhury, D.R., Rijmen, V., Das, A. (eds.) INDOCRYPT 2008. LNCS, vol. 5365, pp. 322–336. Springer, Heidelberg (2008)
5. Black, J., Halevi, S., Krawczyk, H., Krovetz, T., Rogaway, P.: UMAC: fast and secure message authentication. In: Wiener, M. (ed.) CRYPTO 1999. LNCS, vol. 1666, pp. 216–233. Springer, Heidelberg (1999)
6. Chakraborty, D., Sarkar, P.: A general construction of tweakable block ciphers and different modes of operations. IEEE Trans. on Information Theory 54(5) (May 2008)
7. Dworkin, M.: Recommendation for block cipher modes of operation: the CCM mode for authentication and confidentiality. NIST Special Publication 800-38C (May 2004)
8. Dworkin, M.: Recommendation for block cipher modes of operation: Galois/Counter Mode (GCM) and GMAC. NIST Special Publication 800-38D (November 2007)
9. Ekdahl, P., Johansson, T.: A new version of the stream cipher SNOW. In: Nyberg, K., Heys, H.M. (eds.) SAC 2002. LNCS, vol. 2595, pp. 47–61. Springer, Heidelberg (2003)
10. Ferguson, N., Whiting, D., Schneier, B., Kelsey, J., Lucks, S., Kohno, T.: Helix: fast encryption and authentication in a single cryptographic primitive. In: Johansson, T. (ed.) FSE 2003. LNCS, vol. 2887, pp. 330–346. Springer, Heidelberg (2003)
11. Gligor, V., Donescu, P.: Fast encryption and authentication: XCBC encryption and XECB authentication modes. In: Matsui, M. (ed.) FSE 2001. LNCS, vol. 2355, pp. 92–108. Springer, Heidelberg (2002)
12. Gueron, S.: Intel's New AES instructions for enhanced performance and security. In: Dunkelman, O. (ed.) FSE 2009. LNCS, vol. 5665, pp. 51–66. Springer, Heidelberg (2009)
13. Gueron, S., Kounavis, M.: Intel carry-less multiplication instruction and its usage for computing the GCM mode (revision 2) (May 2010), White paper, available from http://www.intel.com
14. Halevi, S.: An observation regarding Jutla's modes of operation. Cryptology ePrint report 2001/015, April 2 (2001)
15. IEEE Standard 802.11i-2004. Part 11: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications: Medium Access Control (MAC) Security Enhancements (2004)
16. ISO/IEC 19772. Information technology – Security techniques – Authenticated encryption, 1st edn. (February 15, 2009)
17. Iwata, T.: Authenticated encryption mode for beyond the birthday bound security. In: Vaudenay, S. (ed.) AFRICACRYPT 2008. LNCS, vol. 5023, pp. 125–142. Springer, Heidelberg (2008)
18. Iwata, T.: New blockcipher modes of operation with beyond the birthday bound security. In: Robshaw, M.J.B. (ed.) FSE 2006. LNCS, vol. 4047, pp. 310–327. Springer, Heidelberg (2006)
19. Iwata, T., Yasuda, K.: BTM: A single-key, inverse-cipher-free mode for deterministic authenticated encryption. In: Jacobson Jr., M.J., Rijmen, V., Safavi-Naini, R. (eds.) SAC 2009. LNCS, vol. 5867, pp. 313–330. Springer, Heidelberg (2009)

20. Iwata, T., Yasuda, K.: HBS: a single-key mode of operation for deterministic authenticated encryption. In: Dunkelman, O. (ed.) FSE 2009. LNCS, vol. 5665, pp. 394–415. Springer, Heidelberg (2009)
21. Jutla, C.: Encryption modes with almost free message integrity. In: Pfitzmann, B. (ed.) EUROCRYPT 2001. LNCS, vol. 2045, pp. 529–544. Springer, Heidelberg (2001)
22. Käsper, E., Schwabe, P.: Faster and timing-attack resistant AES-GCM. In: Clavier, C., Gaj, K. (eds.) CHES 2009. LNCS, vol. 5747, pp. 1–17. Springer, Heidelberg (2009)
23. Katz, J., Yung, M.: Unforgeable encryption and chosen ciphertext secure modes of operation. In: Schneier, B. (ed.) FSE 2000. LNCS, vol. 1978, p. 284. Springer, Heidelberg (2001)
24. Kohno, T., Viega, J., Whiting, D.: CWC: A high-performance conventional authenticated encryption mode. In: Roy, B., Meier, W. (eds.) FSE 2004. LNCS, vol. 3017, pp. 408–426. Springer, Heidelberg (2004)
25. Krovetz, T.: Message authentication on 64-bit architectures. In: Biham, E., Youssef, A.M. (eds.) SAC 2006. LNCS, vol. 4356, pp. 327–341. Springer, Heidelberg (2007)
26. Krovetz, T., Rogaway, P.: The software performance of authenticated-encryption modes. Full version of this paper (January 2011)
27. Leiserson, C., Prokop, H., Randall, K.: Using de Bruijn sequences to index a 1 in a computer word (July 7, 1998) (unpublished manuscript)
28. Lidl, R., Niederreiter, H.: Introduction to finite fields and their applications (Revised Edition). Cambridge University Press, Cambridge (1994)
29. Liskov, M., Rivest, R., Wagner, D.: Tweakable block ciphers. In: Yung, M. (ed.) CRYPTO 2002. LNCS, vol. 2442, pp. 31–46. Springer, Heidelberg (2002)
30. Lucks, S.: Two-pass authenticated encryption faster than generic composition. In: Gilbert, H., Handschuh, H. (eds.) FSE 2005. LNCS, vol. 3557, pp. 284–298. Springer, Heidelberg (2005)
31. McGrew, D.: An interface and algorithms for authenticated encryption. IETF RFC 5116 (January 2008)
32. McGrew, D., Viega, J.: The security and performance of the Galois/Counter Mode (GCM) of operation. In: Canteaut, A., Viswanathan, K. (eds.) INDOCRYPT 2004. LNCS, vol. 3348, pp. 343–355. Springer, Heidelberg (2004); Also Cryptology ePrint report 2004/193, with somewhat different performance results
33. OpenSSL: The Open Source Toolkit for SSL/TLS, http://www.openssl.org/
34. Rogaway, P.: Authenticated-encryption with associated-data. In: CCS 2002. ACM Press, New York (2002)
35. Rogaway, P.: Efficient instantiations of tweakable blockciphers and refinements to modes OCB and PMAC. In: Lee, P.J. (ed.) ASIACRYPT 2004. LNCS, vol. 3329, pp. 16–31. Springer, Heidelberg (2004)
36. Rogaway, P., Bellare, M., Black, J.: OCB: A block-cipher mode of operation for efficient authenticated encryption. ACM Trans. on Information and System Security 6(3), 365–403 (2003); Earlier version, with T. Krovetz, in CCS 2001
37. Rogaway, P., Shrimpton, T.: A provable-security treatment of the key-wrap problem. In: Vaudenay, S. (ed.) EUROCRYPT 2006. LNCS, vol. 4004, pp. 373–390. Springer, Heidelberg (2006)
38. Tsaban, B., Vishne, U.: Efficient linear feedback shift registers with maximal period. Finite Fields and Their Applications 8(2), 256–267 (2002), Also CoRR cs.CR/0304010 (2003)
39. VIA Technologies. VIA Padlock programming guide (2005)

40. Whiting, D., Housley, R., Ferguson, N.: AES encryption & authentication using CTR mode & CBC-MAC. IEEE P802.11 doc 02/001r2 (May 2002)
41. Whiting, D., Housley, R., Ferguson, N.: Counter with CBC-MAC (CCM). IETF RFC 3610 (September 2003)
42. Zeng, G., Han, W., He, K.: High efficiency feedback shift register: σ-LFSR. Cryptology ePrint report 2007/114 (2007)

A   New Word-Oriented LFSRs

Recall that in OCB2 each 128-bit offset is computed from the prior one by multiplying it, in $GF(2^{128})$, by the constant $\mathsf{x} = 2 = 0^{126}10$. Concretely, the point $X \in \{0,1\}^{128}$ is stepped (or "incremented" or "doubled") by applying the map $S(X) = (X \lll 1) \oplus (\mathrm{msb}(X) \cdot 135)$. The constant 135 (decimal) represents (without the x^{128} term) the primitive polynomial $g(\mathsf{x}) = \mathsf{x}^{128} + \mathsf{x}^7 + \mathsf{x}^2 + \mathsf{x} + 1$.

Chakraborty and Sarkar suggested [6] that there might be an incrementing function more efficient than S; they suspected that one might achieve efficiency gains with a *word-oriented* LFSR [38], as exemplified by the blockcipher SNOW [9]. After all, multiplication by x and reducing mod $g(\mathsf{x})$ is just the "Galois configuration" of a particular 128-bit LFSR [28], and one that has not been optimized for software performance. Some other 128-bit LFSRs might run faster.

To develop this idea, let S be an $n \times n$ binary matrix that is invertible over $GF(2)$. Then we may regard S as the feedback matrix of an LFSR that transforms the row vector $X \in \{0,1\}^n$ into the row vector $X \cdot \mathsf{S}$, a process we refer to as *stepping* the string X under S. The t-fold stepping of X by S is realized by matrix S^t. If the characteristic polynomial of S is primitive (over $GF(2)$) then the order of S in the general linear group $GL(n, GF(2))$ will be $2^n - 1$ and the map $X \mapsto X \cdot \mathsf{S}$ will have two cycles: the length-1 cycle from 0^n to itself and the cycle of length $2^n - 1$ passing through all remaining n-bit strings [28]. The matrices $\langle \mathsf{S} \rangle = \{\mathsf{S}^i : 1 \le i \le 2^{n-1} - 1\}$, along with the matrix $n \times n$ zero matrix, can be regarded as a representation of $GF(2^n)$ under the operations of matrix multiplication and matrix addition, both mod 2.

Based on the paragraph above, the following is a simple way to obtain maximal and fast-to-compute 128-bit LFSRs. Generate candidate LFSRs by randomly combining a small number of shifts, ands, xors, using small or random constants. Represent each scheme by its feedback matrix. For each candidate matrix, check if it has a primitive characteristic polynomial. This is roughly the same approach taken by Zeng, Han, and He [42] to devise some software-efficient maximal-period shift registers intended for stream-cipher use. Using it, we generated and tested thousands of 128-bit stepping functions. Some efficient-to-compute schemes giving rise to maximal LFSRs are as follows:

$$X\,Y \mapsto Y\ \ ((X \lll 1) \oplus (\mathrm{msb}(X) \cdot 10^{120}1010001) \oplus Y) \tag{1}$$

$$X\,Y \mapsto Y\ \ ((X \lll 1) \oplus (X \ggg 1) \oplus (Y \wedge 148)) \tag{2}$$

$$A\,B\,C\,D \mapsto C\ \ D\ \ B\ \ ((A \lll 1) \oplus (\mathrm{msb}(A) \cdot 831) \oplus B \oplus D) \tag{3}$$

$$A\,B\,C\,D \mapsto C\ \ D\ \ B\ \ ((A \lll 1) \oplus (A \ggg 1) \oplus (D \wedge 107)) \tag{4}$$

$$A\,B\,C\,D \mapsto C\ \ D\ \ B\ \ ((A \lll 1) \oplus (A \ggg 1) \oplus (D \lll 15)) \tag{5}$$

Here $|X| = |Y| = 64$ and $|A| = |B| = |C| = |D| = 32$. Our experience searching for such maximal LFSRs suggests that they are rather finicky and sparse.

Implementing the candidate LFSRs on a variety of platforms revealed no clear winner. Beyond this, we found that none of the stepping functions were competitive with xoring in a pre-computed 128-bit value. All of the candidate stepping function introduce endian favoritism. In the end, then, we decided against using an LFSR stepping function to update offsets, going back to the OCB1 approach, instead.