## PMAC

Proposal to NIST for a parallelizable message authentication code

April 1, 2001

Submitter:
Phillip Rogaway
rogaway@cs.ucdavis.edu
http://www.cs.ucdavis.edu/~rogaway

Department of Computer Science
University of California at Davis
Engineering II Building
Davis, California 95616 USA

$$
\begin{gathered}
+15307527583 \text { office } \\
+15307524767 \text { fax }
\end{gathered}
$$

Department of Computer Science
Faculty of Science
Chiang Mai University
Chiang Mai 50200 Thailand
+66 18818290 cell
+6653943433 fax

Inventor and Owner:
Same as submitter

Auxiliary submitter:
John Black
jrb@cs.unr.edu
Department of Computer Science
University of Nevada
Reno, Nevada 89557 USA

## Contents

1 Introduction ..... 1
2 Mathematical Preliminaries ..... 1
3 Specification ..... 4
3.1 Definition of the Mode ..... 4
3.2 Conformance Criteria ..... 4
3.3 An Equivalent Description ..... 6
4 Discussion ..... 8
4.1 Properties ..... 8
4.2 Design Rationale ..... 10
4.3 Limitations ..... 11
4.4 Related Work ..... 11
4.5 Design History ..... 12
5 Theorems ..... 13
5.1 Security Definitions ..... 13
5.2 Theorem Statements ..... 13
5.3 Structure of the Proof ..... 14
6 Performance ..... 16
7 Intellectual Property Statement and Disclosures ..... 17
Acknowledgments ..... 17
References ..... 18
A Proofs ..... 20
A. 1 Proof of the Structure Lemma (Lemma 1) ..... 20
A. 2 Proof of the Pairwise-Collision Bound (Lemma 2) ..... 24
B Test Vectors ..... 25
C Document History ..... 25

## 1 Introduction

Popular ways to authenticate a message, like the CBC MAC [27] and HMAC [1], are inherently sequential: one cannot process the $i$-th message block until all previous message blocks have been processed. The sequential nature of these algorithms can limit performance. While this has not been too big an obstacle in the past, the issue can be expected to increase in importance, both for hardware and for software. Special-purpose hardware will get limited by the latency of the underlying cryptographic primitive, while performance on commodity processors will be limited by an inability to fully exploit the multiple instruction pipes provided. Thus it seems a ripe time to develop a fully parallelizable message authentication code (MAC).

This submission describes such a message authentication code, PMAC (which stands for Parallelizable MAC). Using an $n$-bit block cipher, PMAC authenticates an arbitrary string $M \in\{0,1\}^{*}$ using $\lceil|M| / n\rceil$ block-cipher calls. (The empty string is an exception; it requires one block-cipher call.) Overhead beyond the block-cipher calls is low-about $8 \%$ more than with the basic CBC MAC when one is in a non-parallelizable environment. PMAC works correctly across messages of arbitrary and varying bit lengths. It uses a single key for the underlying block cipher. The length of the computed MAC is an arbitrary number of bits $\tau \in[1 . . n]$, with the corresponding forging probability being about $2^{-\tau}$. PMAC is stateless and deterministic: MAC generation does not require a nonce or random value.

We prove PMAC secure, in the sense of reduction-based cryptography. Specifically, we prove that PMAC is a good variable-input-length pseudorandom function (PRF), and is therefore a good MAC, as long as the underlying block cipher is good as a pseudorandom permutation (PRP) [5, 22]. The actual results are quantitative; the security analysis is in the concrete-security paradigm.

Parallelizable MACs related to PMAC are the XOR MAC of Bellare, Guérin and Rogaway [4], the variant of this construction due to Bernstein [9], and the XECB-MAC of Gligor and Donescu [13]. For details on these and other works, see Section 4.4.

## 2 Mathematical Preliminaries

Notation. If $a$ and $b$ are integers, $a \leq b$, then $[a . . b]$ is the set $\{a, a+1, \ldots, b\}$. If $i \geq 1$ is an integer then $n t z(i)$ is the number of trailing 0 -bits in the binary representation of $i$ (equivalently, $\mathrm{ntz}(i)$ is the largest integer $z$ such that $2^{z}$ divides $i$ ). So, for example, $\operatorname{ntz}(7)=0$ and $\mathrm{ntz}(8)=3$.

A string is a finite sequence of symbols, each symbol being 0 or 1 . The string of length 0 is called the empty string and is denoted $\varepsilon$. Let $\{0,1\}^{*}$ denote the set of all strings. Let $i, n$ be nonnegative integers. Then $0^{i}$ and $1^{i}$ denote the strings of $i 0$ 's and 1 's, respectively. Let $\{0,1\}^{n}$ denote the set of all strings of length $n$. If $A \in\{0,1\}^{*}$ then $|A|$ denotes the length of $A$, in bits, while $\|A\|_{n}=\max \{1,\lceil|A| / n\rceil\}$ denotes the length of $A$ in $n$-bit blocks, where the empty string counts as one block. If $A, B \in\{0,1\}^{*}$ then $A B$, or $A \| B$, is their concatenation. If $A \in\{0,1\}^{*}$ and $A \neq \varepsilon$ then $\operatorname{firstbit}(A)$ is the first bit of $A$ and lastbit $(A)$ is the last bit of $A$. If $A \in\{0,1\}^{*}$ and $i \in[0 . .|A|]$ then $A$ [first $i$ bits $]$ and $A[$ last $i$ bits $]$ denote the strings containing the first $i$ bits of $A$ and the last $i$ bits of $A$, respectively. Both of these values are the empty string if $\tau=0$. If $A=a_{n-1} \cdots a_{1} a_{0} \in\{0,1\}^{n}$ is a string (each $a_{i} \in\{0,1\}$ ) then $\operatorname{str} 2 \operatorname{num}(A)$ is the number $\sum_{i=0}^{n-1} 2^{i} a_{i}$. If $a \in\left[0 . .2^{n}-1\right]$ then num $2 \operatorname{str}_{n}(a)$ is the $n$-bit string $A$ such that $\operatorname{str} 2 n u m(A)=a$. Let $\operatorname{len}_{n}(A)=$ num $2 \operatorname{str}_{n}(|A|)$. We omit the subscript when $n$ is understood. If $A, B \in\{0,1\}^{*}$ then $A \oplus B$ is the bitwise xor of $A$ [first $\ell$ bits] and $B$ [first $\ell$ bits], where $\ell=\min \{|A|,|B|\}$ (where $\varepsilon \oplus \varepsilon=\varepsilon$ ). So, for example, $1001 \oplus 11=01$. If $A \in\{0,1\}^{*}$ and $|A|<n$ then $\operatorname{pad}_{n}(A)$ is the string $A 10^{n-|A|-1}$. If $A \in\{0,1\}^{n}$ then $\operatorname{pad}_{n}(A)=A$. With $n$ understood we write $\operatorname{pad}(A)$ for $\operatorname{pad}_{n}(A)$.

If $A=a_{n-1} a_{n-2} \cdots a_{1} a_{0} \in\{0,1\}^{n}$ then $A \ll 1=a_{n-2} a_{n-3} \cdots a_{1} a_{0} 0$ is the $n$-bit string which is a left shift of $A$ by 1 bit (the first bit of $A$ disappearing and a zero coming into the last bit), while $A \gg 1=0 a_{n-1} a_{n-2} \ldots a_{2} a_{1}$ is the $n$-bit string which is a right shift of $A$ by one bit (the last bit disappearing and a zero coming into the first bit).

In pseudocode we write "Partition $M$ into $M[1] \cdots M[m]$ " as shorthand for "Let $m=\|M\|_{n}$ and let $M[1], \ldots, M[m]$ be strings such that $M[1] \cdots M[m]=M$ and $|M[i]|=n$ for $1 \leq i<m$."

The field with $2^{n}$ points. Recall that a finite field is a finite set together with an addition operation and a multiplication operation, each defined to take a pair of points in the field to another point in the field. The operations must obey certain basic axioms. (For example, there must be a point 0 in the field such that $a+0=0+a=a$ for every $a$; there must be a point 1 in the field such that $a \cdot 1=1 \cdot a=a$ for every $a$; and for every $a \neq 0$ there must be a point $a^{-1}$ in the field such that $a \cdot a^{-1}=a^{-1} \cdot a=1$.) If one fixes a positive integer $n$, then there turns out to be a unique finite field (up to the naming of the points) that has $2^{n}$ elements. It is called the Galois field of size $2^{n}$, and it is denoted $\operatorname{GF}\left(2^{n}\right)$.

Example 1 The field $\mathrm{GF}(2)$ has two points, 0 and 1, and operations $\oplus$ (addition) and $\cdot$ (multiplication) are defined by $0 \oplus 0=0,0 \oplus 1=1,1 \oplus 0=1,1 \oplus 1=0,0 \cdot 0=0,0 \cdot 1=0,1 \cdot 0=0$, and $1 \cdot 1=1$.

We interchangeably think of a point $a$ in $\operatorname{GF}\left(2^{n}\right)$ in any of the following ways: (1) as an abstract point in a field; (2) as an $n$-bit string $a_{n-1} \ldots a_{1} a_{0} \in\{0,1\}^{n}$; (3) as a formal polynomial $a(\mathrm{x})=a_{n-1} \mathrm{x}^{n-1}+\cdots+a_{1} \mathrm{x}+a_{0}$ with binary coefficients; (4) as a nonnegative integer between 0 and $2^{n}-1$, where the string $a \in\{0,1\}^{n}$ corresponds to the number str2num $(a)$. For example, one can regard the string $a=0^{125} 101$ as a 128 -bit string, as the number 5 , as the polynomial $\mathrm{x}^{2}+1$, or as a particular point in the finite field $\mathrm{GF}\left(2^{128}\right)$. We write $a(\mathrm{x})$ instead of $a$ if we wish to emphasize that we are thinking of $a$ as a polynomial.

To add two points in $\operatorname{GF}\left(2^{n}\right)$, take their bitwise xor. We denote this operation by $a \oplus b$.
Before we can say how to multiply two points we must fix some irreducible polynomial $p_{n}(\mathrm{x})$ having binary coefficients and degree $n$. (Saying that $p_{n}(\mathrm{x})$ is irreducible means that if $q(\mathrm{x})$ and $q^{\prime}(\mathrm{x})$ are polynomials over $\mathrm{GF}(2)$ which multiply to give $p_{n}(\mathrm{x})$, then one of these polynomials is 1 and the other is $p_{n}(\mathrm{x})$.) For PMAC, choose the lexicographically first polynomial among the irreducible degree $n$ polynomials having a minimum number of coefficients. For $n=128$, the indicated polynomial is

$$
p_{128}(\mathrm{x})=\mathrm{x}^{128}+\mathrm{x}^{7}+\mathrm{x}^{2}+\mathrm{x}+1
$$

A few other $p_{n}(x)$-values are $x^{64}+x^{4}+x^{3}+x+1$ and $x^{96}+x^{10}+x^{9}+x^{6}+1$ and $x^{160}+x^{5}+x^{3}+x^{2}+1$ and $x^{192}+x^{7}+x^{2}+x+1$ and $x^{224}+x^{9}+x^{8}+x^{3}+1$ and $x^{256}+x^{10}+x^{5}+x^{2}+1$.

To multiply points $a, b \in \operatorname{GF}\left(2^{n}\right)$, which we denote $a \cdot b$, regard $a$ and $b$ as polynomials $a(\mathrm{x})=$ $a_{n-1} \mathrm{x}^{n-1}+\cdots+a_{1} \mathrm{x}+a_{0}$ and $b(\mathrm{x})=b_{n-1} \mathrm{x}^{n-1}+\cdots+b_{1} \mathrm{x}+b_{0}$, form their product $c(\mathrm{x})$ where one adds and multiplies coefficients in $\operatorname{GF}(2)$ (the coefficient of degree $j$ in $c(\mathrm{x})$, where $j \in[0 . .2 n-2]$, is $\left.c_{j}=\oplus_{i=0}^{j}\left(b_{i} \cdot a_{j-i}\right)\right)$ and take the remainder one gets when dividing $c(\mathrm{x})$ by the polynomial $p_{n}(\mathrm{x})$.

By convention, the multiplication operator has higher precedence than addition operator so, for example, $\gamma_{1} \cdot L \oplus R$ means $\left(\gamma_{1} \cdot L\right) \oplus R$.

Example 2 Assume $n=128$. Suppose one multiplies $a(\mathrm{x})=\mathrm{x}^{127}+\mathrm{x}+1$ by $b(\mathrm{x})=\mathrm{x}+1$. The result is $c(\mathrm{x})=\mathrm{x}^{128}+\mathrm{x}^{2}+\mathrm{x}+\mathrm{x}^{127}+\mathrm{x}+1=\mathrm{x}^{128}+\mathrm{x}^{127}+\mathrm{x}^{2}+1$. If one divides $c(\mathrm{x})$ by $p(\mathrm{x})$ one gets a quotient of $q(\mathrm{x})=1$ and a remainder (which is the answer) of $r(\mathrm{x})=\mathrm{x}^{127}+\mathrm{x}^{7}+\mathrm{x}$. In string notation, $10^{125} 11 \cdot 0^{126} 11=10^{119} 10000010$.

It is particularly easy to multiply a point $a \in\{0,1\}^{n}$ by x . We illustrate the method for $n=128$, where $p(\mathrm{x})=\mathrm{x}^{128}+\mathrm{x}^{7}+\mathrm{x}^{2}+\mathrm{x}+1$. Then multiplying $a=a_{n-1} \cdots a_{1} a_{0}$ by x yields a product $a_{n-1} \mathrm{x}^{n}+a_{n-2} \mathrm{x}^{n-1}+a_{1} \mathrm{x}^{2}+a_{0} \mathrm{x}$. Thus, if the first bit of $a$ is 0 , then $a \cdot \mathrm{x}=a \ll 1$. If the first bit of $a$ is 1 then we must add $\mathrm{x}^{128}$ to $a \ll 1$. Since $\mathrm{x}^{128}+\mathrm{x}^{7}+\mathrm{x}^{2}+\mathrm{x}+1=0$ we know that $\mathrm{x}^{128}=\mathrm{x}^{7}+\mathrm{x}^{2}+\mathrm{x}+1$, so adding $\mathrm{x}^{128}$ means to xor by $0^{120} 10000111$. In summary, when $n=128$,

$$
a \cdot \mathrm{x}= \begin{cases}a \ll 1 & \text { if firstbit }(a)=0 \\ (a \ll 1) \oplus 0^{120} 10000111 & \text { if firstbit }(a)=1\end{cases}
$$

Example 3 Let us again compute $10^{125} 11 \cdot 0^{126} 11$. Since the latter string is $\mathrm{x}+1$, we should multiply the first string by x and then add it to (xor it with) the first string. As the first bit of $10^{125} 11$ is 1 , multiplying this point by x yields $0^{125} 110 \oplus 0^{120} 10000111=0^{120} 10000001$, and xoring this with $10^{125} 11$ gives a final answer of $10^{119} 10000010$, as before.

If $L \in\{0,1\}^{n}$ and $i \geq 0$, we write $L(i)$ as shorthand for $L \cdot \mathrm{x}^{i}$. We have an easy way to compute $L(1), L(2), \ldots, L(\mu)$-values, where $\mu$ is a small number. Namely, set $L(0)=L$ and compute $L(i)=L(i-1) \cdot \mathrm{x}$ for all $i \in[1 . . \mu]$.

If $a \neq 0$ is a point in $\{0,1\}^{n}$, we can divide $a$ by x , meaning that one multiplies $a$ by the multiplicative inverse of x in the field: $a \cdot \mathrm{x}^{-1}$. It is easy to compute $a \cdot \mathrm{x}^{-1}$. To illustrate, again assume that $n=128$. Then if the last bit of $a$ is 0 , then $a \cdot \mathrm{x}^{-1}$ is $a \gg 1$. If the last bit of $a$ is 1 , then we must add (xor) to $a \gg 1$ the value $x^{-1}$. Since $x^{128}=x^{7}+x^{2}+x+1$ we have $x^{127}=x^{6}+x+1+x^{-1}$ and so $\mathrm{x}^{-1}=\mathrm{x}^{127}+\mathrm{x}^{6}+\mathrm{x}+1=10^{120} 1000011$. In summary, for $n=128$,

$$
a \cdot \mathrm{x}^{-1}=\left\{\begin{array}{lr}
a \gg 1 & \text { if lastbit }(a)=0 \\
(a \gg 1) \oplus 10^{120} 1000011 & \text { if lastbit }(a)=1
\end{array}\right.
$$

We point out that, for any $n=128$, the value huge $=\mathrm{x}^{-1}$ will be an enormous number (when viewed as a number); in particular, huge starts with a 1 bit, so $2^{n-1} \leq$ huge. For the remainder of this submission, we will use huge as a synonym for $\mathrm{x}^{-1}$ whenever this seems to add to clarity. We will later assume that any messages $M=M[1] \cdots M[m]$ to be MACed has block length $m<$ huge, for otherwise our theorem statements assert a non-result. Thus for any message $M=M[1] \cdots M[m]$ to be MACed, each of $\gamma_{1}, \gamma_{2}, \ldots, \ldots, \gamma_{m}$ is different from huge.

Gray codes. For $\ell \geq 1$, a Gray code is an ordering $\gamma^{\ell}=\gamma_{0}^{\ell} \quad \gamma_{1}^{\ell} \ldots \gamma_{2^{\ell}-1}^{\ell}$ of $\{0,1\}^{\ell}$ such that successive points differ (in the Hamming sense) by just one bit. For $n$ a fixed number, PMAC makes use of the "canonical" Gray code $\gamma=\gamma^{n}$ constructed by

$$
\gamma^{1}=0 \quad 1
$$

while, for $\ell>0$,

$$
\gamma^{\ell+1}=0 \gamma_{0}^{\ell} \quad 0 \gamma_{1}^{\ell} \quad \cdots \quad 0 \gamma_{2^{\ell}-2}^{\ell} \quad 0 \gamma_{2^{\ell}-1}^{\ell} \quad 1 \gamma_{2^{\ell}-1}^{\ell} \quad 1 \gamma_{2^{\ell}-2}^{\ell} \quad \cdots \quad 1 \gamma_{1}^{\ell} \quad 1 \gamma_{0}^{\ell}
$$

It is easy to see that $\gamma$ is a Gray code. What is more, for $1 \leq i \leq 2^{n}-1, \gamma_{i}=\gamma_{i-1} \oplus\left(0^{n-1} 1<n \mathrm{ntz}(i)\right)$. This makes it easy to compute successive points.

Example 4 The canonical Gray code with 2 points is $\gamma^{1}=\gamma_{0}^{1} \gamma_{1}^{1}=01$. The canonical Gray code with 4 points is obtained by writing this once forward, then once backwards, prefixing each string in the first half by 0 and prefixing each string in the second half by 1 : that is, $\gamma^{2}=$ $\gamma_{0}^{2} \gamma_{1}^{2} \gamma_{2}^{2} \gamma_{3}^{2}=00011110=0132$. Repeating the process, the canonical Gray code with 8 points
is $\gamma^{3}=\gamma_{0}^{3} \gamma_{1}^{3} \gamma_{2}^{3} \gamma_{3}^{3} \gamma_{4}^{3} \gamma_{5}^{3} \gamma_{6}^{3} \gamma_{7}^{3}=000001011010110111101100=0132675$ 4. In PMAC we use the Gray code $\gamma=\gamma^{n}$ having $2^{n}$ points: $\gamma=\gamma_{0} \gamma_{1} \gamma_{2} \gamma_{3} \cdots \gamma_{2^{n}-1}=0132675412 \cdots 2^{n-1}$. To calculate $\gamma_{i}$ from $\gamma_{i-1}$, xor $\gamma_{i-1}$ by $0^{n-1} 1 \ll \mathrm{ntz}(i)$. For example, $\gamma_{8}=12$ can be computed from $\gamma_{7}=4$ by xoring 4 with $0^{n-1} 1 \ll 3$.

We emphasize the following characteristics of the Gray-code values $\gamma_{1}, \gamma_{2}, \ldots, \gamma_{2^{n}-1}$. First, they are distinct and different from 0 . Second, that $\gamma_{1}=1$. Third, that $\gamma_{i} \leq 2 i$.

Let $L \in\{0,1\}^{n}$ and consider the problem of successively forming the strings $\gamma_{1} \cdot L, \gamma_{2} \cdot L$, $\gamma_{3} \cdot L, \ldots, \gamma_{m} \cdot L$. Of course $\gamma_{1} \cdot L=1 \cdot L=L$. Now, for $i \geq 2$, assume one has already produced $\gamma_{i-1} \cdot L$. Since $\gamma_{i}=\gamma_{i-1} \oplus\left(0^{n-1} 1 \ll \operatorname{ntz}(i)\right)$ we know that

$$
\begin{aligned}
\gamma_{i} \cdot L & =\left(\gamma_{i-1} \oplus\left(0^{n-1} 1 \ll \mathrm{ntz}(i)\right)\right) \cdot L \\
& =\left(\gamma_{i-1} \cdot L\right) \oplus\left(0^{n-1} 1 \ll \mathrm{ntz}(i)\right) \cdot L \\
& =\left(\gamma_{i-1} \cdot L\right) \oplus\left(L \cdot \mathrm{x}^{\mathrm{ntz}(i)}\right) \\
& =\left(\gamma_{i-1} \cdot L\right) \oplus L(\mathrm{ntz}(i))
\end{aligned}
$$

That is, the $i$ th word in the sequence $\gamma_{1} \cdot L, \gamma_{2} \cdot L, \gamma_{3} \cdot L, \ldots$ is obtained by xoring the previous word with $L(\mathrm{ntz}(i))$.

## 3 Specification

### 3.1 Definition of the Mode

Parameters. To use PMAC one must specify two parameters: a block cipher and a tag length.

- The block cipher $E$ is a function $E: \mathcal{K} \times\{0,1\}^{n} \rightarrow\{0,1\}^{n}$, for some number $n$, where each $E(K, \cdot)=E_{K}(\cdot)$ is a permutation on $\{0,1\}^{n}$. Here $\mathcal{K}$ is the set of possible keys and $n$ is the block length. Both are arbitrary, though we insist that $n \geq 64$, and we discourage $n<128$.
- The tag length $\tau$ is an integer between 1 and $n$. By trivial means, the adversary will be able to forge a valid ciphertext with probability $2^{-\tau}$.

The popular block cipher to use with PMAC is likely to be AES, but any other block cipher is fine. As for the tag length, a suggested default of $\tau=64$ is reasonable, though both shorter and longer tags are likely to be common. Note that tags of $\tau=32$ bits have been standard for retail banking for many years, while tags of $\tau=80$ bits are used in IPSec. Using a tag of more than 80 bits adds questionable security benefit, though it does entail extra bits being transmitted or stored.

With $E: \mathcal{K} \times\{0,1\}^{n} \rightarrow\{0,1\}^{n}$ and $\tau \in[1 . . n]$, we let PMAC[ $\left.E, \tau\right]$ denote the PMAC mode of operation using block cipher $E$ and tag length $\tau$. This is a function from a key $K \in \mathcal{K}$ and a message $M \in\{0,1\}^{*}$ to a string in $\{0,1\}^{\tau}$.

Definition. Letting $E: \mathcal{K} \times\{0,1\}^{n} \rightarrow\{0,1\}^{n}$ and $\tau \in[1 . . n]$, the definition of the function $\operatorname{PMAC}[E, \tau]$ is given in Figure 1 and illustrated in Figure 2.

### 3.2 Conformance Criteria

An implementation of PMAC is said to conform to the specification if some specified subset MsgSpace $\subseteq\{0,1\}^{*}$ of messages can be presented to be MACed, and any message $M$ carrying a tag Tag is deemed invalid if $\operatorname{PMAC}_{K}(M) \neq$ Tag. For example,

```
Algorithm \(\mathrm{PMAC}_{K}(M)\)
    \(L \leftarrow E_{K}\left(0^{n}\right)\)
    Partition \(M\) into \(M[1] \cdots M[m]\)
    for \(i \leftarrow 1\) to \(m-1\) do
        \(Y[i] \leftarrow E_{K}\left(M[i] \oplus \gamma_{i} \cdot L\right)\)
    \(\Sigma \leftarrow Y[1] \oplus Y[2] \oplus \cdots \oplus Y[m-1] \oplus \operatorname{pad}(M[m])\)
    if \(|M[m]|=n\) then \(X[m]=\Sigma \oplus L \cdot \mathbf{x}^{-1}\)
        else \(X[m] \leftarrow \Sigma\)
    Tag \(=E_{K}(X[m])\) [first \(\tau\) bits]
    return Tag
```

Figure 1: Definition of PMAC. Constants $\gamma_{1}, \gamma_{2}, \ldots$ and the meaning of the multiplication operator are defined in the text, but $\gamma_{i} \cdot L$ is simply $\gamma_{i-1} \cdot L \oplus L(\operatorname{ntz}(i))$ (where $L(j)$ is defined in the text and easily computed from $L$ ). We let $\operatorname{pad}(A)=A 10^{n-|A|-1}$ if $|A|<n$ and $\operatorname{pad}(A)=A$ if $|A|=n$.


Figure 2: Illustration of PMAC. The message to MAC is $M$ and the key is $K$. Message $M$ is written as $M=M[1] \cdots M[m]$, where $m=\max \{1,\lceil|M| / n\rceil\}$ and $|M[1]|=|M[2]|=\cdots=|M[m-1]|=n$. Value $L=E_{K}\left(0^{n}\right)$ is derived from $K$.

- A conforming implementation might only be able to MAC byte strings, these strings limited to $2^{36}$ bytes.
- A message $M$ might be deemed invalid for reasons which go beyond its presenting an incorrect tag. Replay detection would be the most typical reason.


### 3.3 An Equivalent Description

The following description of PMAC may help to clarify what a typical implementation might choose to do. In what follows, fix a block length $n$, block cipher $E: \mathcal{K} \times\{0,1\}^{n} \rightarrow\{0,1\}^{n}$, and a tag length $\tau$. A PMAC implementation may work as follows.

Key generation. Choose a random key $K \stackrel{R}{\leftarrow} \mathcal{K}$ for the block cipher. The key $K$ is provided to both the party that generates MACs and the party that verifies them.

Key setup. Once the key $K$ is known, the following may be precomputed.

1. Set up the block-cipher key. Both the party that generates the MACs and the party that verifies the MACs do any key setup useful for applying the block-cipher in its forward direction.
2. Precompute $L$. Compute $L \leftarrow E_{K}\left(0^{n}\right)$.
3. Precompute $L(i)$-values. Let $m^{*}$ bound the maximum number of $n$-bit blocks for any message which will be MACed. Let $\mu \leftarrow\left\lceil\log _{2} m^{*}\right\rceil$. Let $L(0) \leftarrow L$ and, for $i \in[1 . . \mu]$, compute $L(i) \leftarrow L(i-1) \cdot \mathrm{x}$ using a shift and a conditional xor, as described in Section 2. Compute $L(-1) \leftarrow L \cdot \mathrm{x}^{-1}$ using a shift and a conditional xor, as described in Section 2. Save the values $L(-1), L(0), L(1), L(2), \ldots, L(\mu)$ in a table. (Note: alternatively, if one wishes to save space, compute only the first few $L(i)$-values now, and compute any further ones only when needed, "on the fly.")

MAC generation. To generate the MAC Tag for a message $M \in\{0,1\}^{*}$, do the following steps.

1. Partition the message. Let $m \leftarrow\lceil|M| / n\rceil$. If $m=0$ then let $m \leftarrow 1$. Let $M[1], \ldots, M[m]$ be strings such that $M[1] \cdots M[m]=M$ and $|M[i]|=n$ for $i \in[1 . . m-1]$.
2. Initialize variables. Let Offset $\leftarrow 0^{n}$. Let $\Sigma \leftarrow 0^{n}$.
3. Encipher all blocks but the last one. For $i \leftarrow 1$ to $m-1$, do the following:

Let Offset $\leftarrow$ Offset $\oplus L(\mathrm{ntz}(i))$.
Let $Y[i] \leftarrow E_{K}(M[i] \oplus$ Offset $)$.
Let $\Sigma \leftarrow \Sigma \oplus Y[i]$.
4. Compute the MAC.

Let $\Sigma \leftarrow \Sigma \oplus \operatorname{pad}(M[m])$.
If $|M[m]|=n$ then let $Y[m] \leftarrow E_{K}(\Sigma \oplus L(-1))$ else let $Y[m] \leftarrow E_{K}(\Sigma)$.
Let Tag be the first $\tau$ bits of $Y[m]$.
Return Tag as the computed MAC.
MAC verification. To test if ( $M, \mathrm{Tag}^{\prime}$ ) is authentic, do the following.

1. Re-MAC the message. Generate the MAC Tag for the message $M$ using the MAC generation procedure just described.
2. Compare the presented and the re-computed MACs. If Tag $=T a g^{\prime}$ then regard the message $M$ as authentic. If $T a g \neq T a g^{\prime}$ then regard the message $M$ as inauthentic.

| Security <br> Function | Message Authentication Code. The computed MAC would normally be used to provide message authenticity. But, more generally, PMAC is a pseudorandom function (PRF) having variable input-length $\left(\{0,1\}^{*}\right)$ and fixed output-length $\left(\{0,1\}^{\tau}\right)$. Besides their use as MACs, PRFs can be used, for example, for key separation and within entity-authentication and key-distribution protocols. |
| :---: | :---: |
| Error Propagation | Not applicable. (However, any accidental modification to ( $M$, Tag) will be detected during MAC verification with probability about $1 / 2^{\tau}$.) |
| Synchronization | Not applicable. |
| Parallelizability | Fully parallelizable. All block-cipher invocations (except the last one) may be computed at the same time. |
| Keying Material | One block-cipher key. One needs a single key, $K$, which keys all invocations of the underlying block cipher. |
| Ctr/IV/Nonce Requirements | None. No counter/IV/nonce is used. |
| Memory Requirements | Very modest. About $6 n$ bits beyond the key are sufficient for internal calculations. Implementations may choose whether or not to save $L(i)$ values, offering some tradeoff between memory and simplicity/speed. |
| Pre-processing Capability | Limited. During key-setup the string $L$ would typically be pre-computed (one block cipher call), as would a few $L(i)$-values. Additional precomputation is not possible. |
| Message-Length Requirements | Any bit string allowed. Any string $M \in\{0,1\}^{*}$ may be MACed, including the empty string and strings which are not an integral number of bytes. The length of the string need not be known in advance. |
| Ciphertext <br> Expansion | Not applicable. (But the MAC itself is of minimal length: $\tau$ bits are used to obtain a forging probability of approximately $1 / 2^{\tau}$.) |
| Other Characteristics | Efficiency: PMAC uses $\lceil\|M\| / n\rceil$ block-cipher calls when $M \neq \varepsilon$ (uses one block-cipher call for $M=\varepsilon$ ). Overhead beyond block-cipher calls is low. Endian neutral: Can be implemented equally efficiently on bigendian and little-endian machines (assuming the underlying block cipher has this property). No decryption. PMAC uses only the forward direction of the block cipher. Incrementality. PMAC is incremental, in the sense of [3], with respect to replacement, truncation, and concatenation. Parsimonious PRF. This property lets one easily construct from PMAC a parallelizable block cipher which acts on any message $M \in\{0,1\}^{*}$ of at least $n$ bits. Provable security: PMAC provably meets its goals, assuming the underlying block cipher meets a standard cryptographic assumption (security in the sense of a pseudorandom permutation). |

Figure 3: Summary of properties of PMAC. Further discussion of these and other properties is given in Section 4.1.

## 4 Discussion

### 4.1 Properties

PMAC has been designed to have a variety of desirable properties. These properties are summarized in Figure 3. We now expand on some of the points referenced in that table.

Security Function. PMAC is a (variable-length input, fixed-length output) pseudorandom function (PRF): as long as the underlying block cipher $E$ is secure, an adversary will be unable to distinguish $\operatorname{PMAC}_{K}(\cdot)$, for a random but hidden key $K$, from a random function $\rho$ from $\{0,1\}^{*}$ to $\{0,1\}^{\tau}$. It is a well-known observation, dating to the introduction of PRFs [15], that a pseudorandom function is necessarily a good MAC. Of course the converse is not true, and PMAC, being a good PRF, has uses which go beyond its being a good MAC.

Parallelizability. In settings where there is adequate opportunity for parallelism, PMAC will be faster than the CBC MAC or HMAC, which are inherently non-parallelizable. Parallelizability is becoming important for obtaining good performance from both high-speed hardware and commodity processors. In the former case, one may want to authenticate messages at speeds in excess of 10 Gbits/second - an impossible task for the CBC MAC or HMAC (with today's technology). In the latter case, there is an architectural trend towards highly-pipelined machines with multiple instruction pipes and lots of registers. Optimally exploiting such features necessitates algorithms with plenty to do in parallel.

Keying Material. Conceptually the key is ( $K, L$ ), but $L$ is defined from the underlying key $K$, and then key $K$ is still used. Normally such "lazy key-derivation" would get one in trouble. For PMAC we prove that it does not. Avoiding multiple block-cipher keys is desirable for saving on memory and key-setup time.

Ctr/IV/Nonce Requirements. Conventional MACs have always been deterministic. Maintaining this characteristic results in shorter MACs and a scheme more robust against usage errors.

Message-Length Requirements. Any string $M \in\{0,1\}^{*}$ can be MACed, and messages which are not a multiple of the block length are handled without the need for obligatory padding, which would increase the number of block-cipher calls. MAC generation is "on line," meaning that one does not need to know the length of the message $M$ in advance. Instead, the message can be MACed as one goes along, continuing until there is an indication that the message is now complete. An incremental interface (as is popular for cryptographic hash functions) would be used to support this functionality.

Efficiency. Shaving off a few block-cipher calls may not seem important. But often one is dealing with short messages; for example, roughly a third of the messages on the Internet backbone are 43 bytes. If one is MACing messages of such short lengths, one should be careful about extra computational work since, by percentage, the inefficiencies can be large.

Endian Neutrality. In contrast to a scheme based on $\bmod p$ arithmetic or based on $\bmod 2^{n}$ arithmetic, there is almost no endian-favoritism implicit in the definition of PMAC. (The exception is that the one left shift used for forming $L(i+1)$ from $L(i)$ is more convenient under a big-endian convention, as is the one right shift used for forming $L(-1)=L \cdot \mathrm{x}^{-1}$ from $L$.)

Incrementality. The concept of incrementality for a cryptographic primitive was introduced by [15]. For a MAC, the idea is that, having already computed the MAC Tag of some (possibly long) message $M$, if $M$ is modified into some similar message $M^{\prime}$, the time to compute the MAC Tag ${ }^{\prime}$ for the new message $M^{\prime}$ should be proportional to the amount of change that $M$ underwent when being modified to $M^{\prime}$ (as opposed to just computing the MAC of $M^{\prime}$ from scratch, which would typically take $\Theta\left(\left|M^{\prime}\right|\right)$ time.) Assuming that $\tau=n$ (or that one retains a constant amount of extra information), PMAC is incremental with respect to three common operations for changing $M$ to $M^{\prime}$, namely, append, truncate, and replace. These operations are defined as follows: $\operatorname{append}(M, x)=M \| x$; truncate $(M$, ndrop $)=M[$ first $|M|-$ ndrop bits], where $|M| \geq$ ndrop; and replace $(M$, posn,$x)=M$ [first posn - 1 bits] $\|x\| M[$ last $|M|$ - posn - $|x|+1$ bits], where $|M| \geq \operatorname{posn}+|x|-1)$. For these three operations it is easy to see how to update the MAC of $M$ in time proportional to $|x|$, ndrop, or $|x|$, respectively.

Incrementality is useful when, for example, one wishes to continually maintain the MAC of a large file which is being edited, or the MAC of a file system. Various tree-based schemes offer solutions with logarithmic space overhead. For PMAC, there is zero or constant overhead.

Parsimonious PRF. Partition $M$ into $M=M[1] \cdots M[m-1] M[m]$ and assume $|M| \geq n$. Assume a tag length of the full $\tau=n$ bits. We make the following observation: for any $i \in[1 . . m]$ such that $M[i]$ is a full block (that is, $|M[i]|=n$ ), there is a simple algorithm to recover $M[i]$ given the key $K$, the message $M^{\prime}=M[1] \ldots M[i-1] M[i+1] \cdots M[m]$ which omits block $M[i]$, and the MAC Tag $=\operatorname{PMAC}_{K}(M)$. This property of a PRF was identified in [7], where a PRF having this property was said to be parsimonious. As shown in [7], a parsimonious PRF can be combined with a block cipher $E$ to yield a length-preserving pseudorandom permutation (that is, a "variable-input-length block cipher") that acts on messages of any number of bits $\ell$ greater than or equal to $n$. In particular, given $K 1, K 2 \in \mathcal{K}$, one can encipher $M$ under key $K=K 1 \| K 2$ by first setting IV $=\operatorname{PMAC}_{K 1}(M)$ and then encrypting $M$ using CTR-mode using key $K 2$ and an initial counter value of IV. This yields a ciphertext $\mathcal{C}=\mathrm{IV} \| C[1] \cdots C[m-1] C[m]$. One can use $\mathcal{C}$ directly, as a deterministic (but not length-preserving) encryption of $M$, or one can drop $n$ of the bits from $\mathcal{C}$, the enciphered string being $C=\mathrm{IV} \| C[1] \ldots C[m-2] C[m]$, say. In this way one has constructed a length-preserving pseudorandom permutation (PRP) that is fully parallelizable and that uses a number of block-cipher calls which is roughly twice that of CBC encryption.

Provable Security. In recent years provable security has become a popular goal for practical protocols. This is for good reason; demonstrating provable security is the best way to gain assurance that a cryptographic scheme does what it is supposed to do. For a scheme which enjoys provable security, one does not need to consider attacks, since successful ones imply successful attacks on some simpler object (e.g., the algorithm of the AES). Provable security represents a major departure from iterated, attack-directed design.

Fix a block cipher $E$ and a tag length $\tau$. When we say that PMAC is provably secure we are asserting the existence of a particular theorem (namely, Theorem 1, or one of its corollaries). The theorem shows that if an adversary $A$ could do a good job at distinguishing $\operatorname{PMAC}_{K}(\cdot)$, for a random but hidden key $K$, from a random function $\rho(\cdot)$, where $\rho:\{0,1\}^{*} \rightarrow\{0,1\}^{\tau}$, then there would be an adversary $B$, about as efficient as $A$, that does a good job at distinguishing block cipher $E_{K}(\cdot)$, for a random key $K$, from a random permutation $\pi(\cdot)$, where $\pi:\{0,1\}^{n} \rightarrow\{0,1\}^{n}$. A theorem of this sort is called a reduction. In cryptography, provable security means giving reductions (along with the associated definitions).

Provable security begins with Goldwasser and Micali [16]. The style of provable security which we use here - where the primitive is a block cipher, the scheme is a usage mode, and the analysis
is concrete (no asymptotics) - is the approach of Bellare and Rogaway $[2,4,5]$.
It is not really enough to know that there is a provable-security result for the scheme in question: one should also understand the definitions and the bounds. We have already sketched the definitions. When we speak of the bounds we are addressing "how effective is the adversary $B$ in terms of the efficacy of adversary $A$ " (where $A$ and $B$ are the adversaries above). For PMAC, the bounds can be summarized as follows: the maximal advantage an adversary can get in attacking the constructed PRF is limited to about $\sigma^{2} / 2^{n}$ more than what a similarly powerful adversary can get in attacking the underlying block cipher. The advantage is a real number, between 0 and 1 , with 0 meaning that the adversary is doing terribly and 1 meaning that the adversary is doing great. Here $\sigma$ is the total number of messages the adversary sees the MACs of. The conclusion is that one is safe using PMAC as long as $E$ is a good block cipher and $\sigma$ is small compared to $2^{n / 2}$. This is the same security degradation one observes for the CBC MAC [5, 23]. This kind of security loss was the main motivation for choosing an AES block length of $n=128$ bits, and it is the reason that we discourage the use of PMAC with $n<128$.

### 4.2 Design Rationale

Avoiding state and randomness. A message authentication code is stateful if the authenticating party maintains state, typically a counter, across MAC-generation invocations. It is probabilistic if the authenticating party uses random bits to generate a MAC. (In both cases, MAC-verification is stateless and deterministic.) As in [4], construction of a stateful message authentication code would have facilitated obtaining a better concrete security bound than PMAC delivers. A randomized scheme would also have had a better bound. But PMAC anticipates the use of a block cipher with a block length of $n \geq 128$, and so the bound that we get - the customary $\Theta\left(\sigma^{2} 2^{-n}\right)$ boundis quite acceptable. By insisting on a stateless and deterministic scheme one saves bits in the MAC compared to a probabilistic or stateful scheme, one eliminates the sender's need to maintain state across MAC-generation invocations, and one largely eliminates concern about usage errors, since one does not have to worry about what happens if the sender reuses a nonce or provides a non-random value where a random value is called for.

Not fixing the tag length. The number of bits that are necessary and appropriate for the tag length vary according to the application. In a context where the adversary obtains something quite valuable from a single successfully forgery, one may wish to choose a tag length of 80 bits or more. In contexts such as authenticating a video stream, where an adversary would have to forge a signficant fraction of the frames even to have a noticeable effect on the image, an 8 -bit tag may be appropriate. With no universally appropriate value to choose, it is best to leave this parameter unspecified.

We comment that short tags seem to be more appropriate for PMAC than for some other MAC, particularly Carter-Wegman MACs $[12,28]$. Many Carter-Wegman MACs have the property that if you can forge one message with probability $\delta$, than you can forge an arbitrary set of (all correct) messages with probability $\delta$. This does not appear to be true for PMAC.

Avoiding mod $2^{n}$ addition. Our earlier designs included a scheme based on modular $2^{n}$ addition ("addition" for the remainder of this paragraph) [25]. This is an interesting approach due to Gligor and Donescu [13]. Compared to our GF $\left(2^{n}\right)$-based approach ("xor" for the remainder of this paragraph), an addition-based scheme is quicker to understand a specification for, and may be easier to implement. But the use of addition (where $n \geq 128$ ) has signficant disadvantages:

- The bit-asymmetry of the addition operator implies that the resulting scheme will have a bias
towards big-endian architectures or little-endian architectures; there will be no way to achieve an endian-neutral scheme. The AES algorithm was constructed to be endian-neutral. We did not want to lose this nice attribute with our mode of operation.
- Modular addition of $n$-bit words is particularly unpleasant when programming in a high-level language, where one does not have access to the underlying add-with-carry instruction.
- Modular addition of $n$-bit words is not parallelizable. As a consequence, dedicated hardware will perform this operation more slowly than xor, and, correspondingly, modern processors can xor two $n$-bit quantities faster than they can add them.
- The concrete security bound is worse with an addition-based scheme: the degradation in the bound appears to be $\Theta\left(\lg m^{*}\right)$, where $m^{*}$ is the maximal message length.
- A correct proof for the addition scheme is substantially more complex than a correct proof for the xor-scheme.
For all of these reasons, we eventually decided to reject the addition-based approach in favor of the xor-based one.


### 4.3 Limitations

We note the following limitations. None of these points are specific to PMAC - they apply to any MAC - but they are still important enough to single out.

- As with all modes of operation, the user should be careful not to use the MAC key $K$ for any other purpose. Standard key-separation techniques should be used to derive a multiplicity of keys when keys are needed for a multiplicity of purposes. The key $K$ itself must not be used to derive additional keys.
- A MAC does not, by itself, provide for replay detection. When replay detection is desired, it should be added, using well-known techniques, by a higher-level, protocol.
- One should be careful in combining an encryption scheme and a MAC: the keys must be distinct, and it is generally preferable to MAC the ciphertext than to encrypt the MAC (and when one wants privacy one should not transmit, in the clear, the MAC of a plaintext). See [6] for an analysis.
- The utility of a MAC is eliminated if one encrypts using an authenticated-encryption scheme, such as [13, 20, 26].


### 4.4 Related Work

PMAC springs from ideas in [4], [9], and [13]. The lineage of ideas is from [4], to [9] and [13], and on to PMAC. Let us sketch representative constructions from [4, 9,13$]$.

The XOR MAC. Bellare, Guérin and Rogaway defined a fully parallelizable MAC they called the XOR MAC [4]. In this MAC the message $M$ is divided into pieces $M[1] \cdots M[\ell]$ of length less than the block length. For concreteness, think of each $M[i]$ as having 64 bits and the block length as having $n=128$ bits. Each $M[i]$ is preceded by the number $i$, encoded as the 64bit string $\underline{i}$, and $E_{K}$ is applied to $\underline{i} \| M[i]$. Xor together all $\ell$ results to get an $n$-bit block $\Sigma=E_{K}(\underline{1} \| M[1]) \oplus \cdots \oplus E_{K}(\underline{\ell} \| M[\ell])$. Block cipher $E_{K}$ is applied to one more block, $\underline{0} \|$ IV, where IV is a counter or random value. This gives the ciphertext block $P=E_{K}(\underline{0} \| \mathrm{IV})$. The MAC is (IV, $\Sigma \oplus P$ ).

The XOR MAC requires $\ell+1 \approx 2 m+1$ block-cipher invocations to authenticate a message of $m$ $n$-bit blocks. So one gets parallelizability at the cost of about a factor of two in serial speed. One also pays in the need for randomness or state, and in the length of the MAC (which is longer in
order to communicate this randomness or state). These two limitations-reduced serial efficiency and the use of randomness or state - are overcome individually in [13] and [9], respectively.

Deterministic variants. Bernstein suggested the following variant of the XOR MAC [9]. He starts not with a block cipher but with a pseudorandom function $F$ where, to illustrate, let us say that $F_{K}$ maps 384 bits to 256 bits. In such a case, break the message $M$ into 256 -bit blocks $M[1], \ldots, M[\ell]$ and apply $F_{K}$ to each of $\underline{1}\|M[1], \underline{2}\| M[2], \ldots, \underline{\ell} \| M[\ell]$, where $\underline{i}$ now means the 128 -bit string corresponding to the number $i$. Xor the resulting ciphertext blocks to get a value $\Sigma=F_{K}(\underline{1} \| M[1]) \oplus \cdots \oplus F_{K}(\underline{\ell} \| M[\ell])$. Then apply $F_{K}$ to $\underline{0} \| \Sigma$ to give the MAC $T a g=F_{K}(\underline{0} \| \Sigma)$.

One could consider the following modification to Bernstein's scheme, where one uses a block cipher $E$ instead of the PRF $F$. Form $\Sigma$ exactly as with the XOR MAC, but then form the MAC as $T a g=E_{K^{\prime}}(\Sigma)$, where $K^{\prime}$ is part of the MAC key Key $=K \| K^{\prime}$.

An alternative to $\underline{i} \|$ • annotation. Gligor and Donescu suggested the following approach, which they called the XECB-MAC [13]. Let $E: \mathcal{K} \times\{0,1\}^{n} \rightarrow\{0,1\}^{n}$ be a block cipher and let $K \in \mathcal{K}$. Given a counter ctr, let $R=E_{K}(\operatorname{ctr})$. Let $M=M[1] \cdots M[m]$ be the string we want to MAC, partitioned into $n$-bit words. Then encipher the following $m$ strings using $E_{K}: M[1]+R$, $M[2]+2 R, \ldots, M[m]+m R$. The arithmetic is done modulo $2^{n}$. Xor together these $m$ ciphertexts to get a value $\Sigma=E_{K}(M[1]+R) \oplus \cdots \oplus E_{K}(M[m]+m R)$. The MAC is ctr together with $\Sigma \oplus E_{K}$ (ctr). The advantage of this approach is that one does not "waste bits" for $\underline{i} \|$-encoding; instead, symmetry has been broken by different means - by adding to the $i$ th block of the message the $i$ th multiple of a number unknown to the adversary.

Further afield. A completely different way for improving parallelizability is to generically construct a more parallelizable MAC from an arbitrary one. For example, one could break the message $M[1] \cdots M[2 m]$ into pieces $M^{\prime}=M[1] M[3] M[5] \cdots M[2 m-1]$ and $M^{\prime \prime}=M[2] M[4] M[6] \cdots M[2 m]$, and separately compute $\operatorname{MAC}_{K}\left(M^{\prime}\right) \| \operatorname{MAC}_{K}\left(M^{\prime \prime}\right)$, and then MAC these MACs, under a separate key. But such an approach requires one to anticipate the maximal amount of parallelism that one aims to extract (the example here shows how to get a factor of two). We are not interested in such schemes; we want a MAC that is "fully parallelizable."

Another approach for making a parallelizable MAC is the Carter-Wegman paradigm [12, 28], as in $[11,18,21]$, making sure to select a universal hash-function family that is fully parallelizable. In fact, most universal hash-functions that have been suggested are fully parallelizable. This is perfectly workable, but fast constructions for universal hash-functions have proven to be complex to specify or to implement well $[9,11]$.

### 4.5 Design History

The initial version of PMAC, described in [25], came in three "forms"- one based on $\bmod 2^{n}$ addition, one based on $\operatorname{GF}\left(2^{n}\right)$ addition, and one based on $\bmod p$ addition. However, we eventually settled on the $\mathrm{GF}\left(2^{n}\right)$ scheme, for the reasons discussed in Section 4.2. What is described in the current submission differs from the $\operatorname{GF}\left(2^{n}\right)$-scheme in [25] only in minor ways.

This project was carried out in parallel with the design of an authenticated-encryption scheme, OCB [26]. We have attempted to give these two algorithms a similar flavor. However, the authenticated-encryption goal turns out to be much more complex than the message-authentication goal, and the design of OCB was much harder than the design of PMAC.

## 5 Theorems

This section gives our security results on PMAC. The proofs of the two lemmas used are deferred to Appendix A. We begin with the requisite definitions, which are standard; see, for example, [5].

### 5.1 Security Definitions

A block cipher is a function $E: \mathcal{K} \times\{0,1\}^{n} \rightarrow\{0,1\}^{n}$ where $\mathcal{K}$ is a finite set of strings and each $E_{K}(\cdot)=E(K, \cdot)$ is a permutation on $\{0,1\}^{n}$. Let $\operatorname{Perm}(n)$ denote the set of all permutations on $\{0,1\}^{n}$. This set can be regarded as a block cipher by imagining that each permutation is named by a unique string. Let $A$ be an adversary (a probabilistic algorithm) with access to an oracle, and suppose that $A$ always outputs a bit. Define

$$
\operatorname{Adv}_{E}^{\operatorname{prp}}(A)=\operatorname{Pr}\left[K \stackrel{R}{\stackrel{R}{\leftarrow}} \mathcal{K}: A^{E_{K}(\cdot)}=1\right]-\operatorname{Pr}\left[\pi \stackrel{R}{\leftarrow} \operatorname{Perm}(n): A^{\pi(\cdot)}=1\right]
$$

The above is the probability that adversary $A$ outputs 1 when given an oracle for $E_{K}(\cdot)$, minus the probability that $A$ outputs 1 when given an oracle for $\pi(\cdot)$, where $K$ is selected at random from $\mathcal{K}$ and $\pi$ is selected at random from $\operatorname{Perm}(n)$.

A function family from $n$-bits to $n$-bits is a map $F: \mathcal{K} \times\{0,1\}^{n} \rightarrow\{0,1\}^{n}$ where $\mathcal{K}$ is a finite set of strings. We write $F_{K}(\cdot)$ for $F(K, \cdot)$. Let $\operatorname{Rand}(n)$ denote the set of all functions from $\{0,1\}^{n}$ to $\{0,1\}^{n}$. This set can be regarded as a function family by imagining that each function in $\operatorname{Rand}(n)$ is named by a unique string. Define

$$
\operatorname{Adv}_{F}^{\operatorname{prf}}(A)=\operatorname{Pr}\left[K \stackrel{R}{\stackrel{R}{\leftarrow}} \mathcal{K}: A^{F_{K}(\cdot)}=1\right]-\operatorname{Pr}\left[\rho \stackrel{R}{\leftarrow} \operatorname{Rand}(n): A^{\rho(\cdot)}=1\right]
$$

A function family from $\{0,1\}^{*}$ to $\{0,1\}^{\tau}$ is a map $f: \mathcal{K} \times\{0,1\}^{*} \rightarrow\{0,1\}^{\tau}$ where $\mathcal{K}$ is a set with an associated distribution. We write $f_{K}(\cdot)$ for $f(K, \cdot)$. Let $\operatorname{Rand}(*, \tau)$ denote the set of all functions from $\{0,1\}^{*}$ to $\{0,1\}^{\tau}$. This set is given a probability measure by asserting that a random element $\rho$ of $\operatorname{Rand}(*, \tau)$ associates to each string $x \in\{0,1\}^{*}$ a random string $\rho(x) \in\{0,1\}^{\tau}$. Define

$$
\operatorname{Adv}_{f}^{\operatorname{prf}}(A)=\operatorname{Pr}\left[K \stackrel{R}{\leftarrow} \mathcal{K}: A^{f_{K}(\cdot)}=1\right]-\operatorname{Pr}\left[g \stackrel{R}{\gtrless} \operatorname{Rand}(*, \tau): A^{g(\cdot)}=1\right]
$$

### 5.2 Theorem Statements [DRAFT]

We give the following information-theoretic bound on the security of PMAC.
Theorem $1[\operatorname{PMAC}[\operatorname{Perm}(n), \tau] \approx \operatorname{Rand}(*, \tau)]$ Fix PMAC parameters $n$ and $\tau$. Let $A$ be an adversary with an oracle. Suppose that $A$ asks its oracle $q$ queries, these queries having aggregate length of $\sigma$ blocks. Let $\bar{\sigma}=\sigma+q+1$. Then

$$
\operatorname{Adv}_{\operatorname{PMAC}[\operatorname{Perm}(n), \tau]}^{\mathrm{prf}}(A) \leq \frac{1.5 \bar{\sigma}^{2}}{2^{n}}
$$

In the theorem statement, and from now on, the aggregate length of messages $M_{1}, \ldots, M_{q}$ asked by $A$ of its oracle is the number $\sigma=\sum_{r=1}^{q}\left\|M_{r}\right\|_{n}$.

From the theorem above it is standard to pass to a complexity-theoretic analog. One gets the following. Fix PMAC parameters $n$ and $\tau$, and a block cipher $E: \mathcal{K} \times\{0,1\}^{n} \rightarrow\{0,1\}^{n}$. Let $A$ be an adversary with an oracle. Suppose these $q$ queries have aggregate length of $\sigma$ blocks. Let $\bar{\sigma}=\sigma+q+1$. Let $\delta=\operatorname{Adv}_{\mathrm{PMAC}[E, \tau]}^{\mathrm{prf}}(A)-1.5 \bar{\sigma}^{2} / 2^{n}$. Then there is an adversary $B$ for attacking

```
\(L \stackrel{R}{\leftarrow}\{0,1\}^{n}\)
for \(i \leftarrow 1\) to \(m-1\) do \(\left\{X[i] \leftarrow M[i] \oplus \gamma_{i} \cdot L ; \quad Y[i]{ }^{R}{ }^{R}\{0,1\}^{n}\right\}\)
\(\Sigma \leftarrow Y[1] \oplus \cdots \oplus Y[m-1] \oplus \operatorname{pad}(M[m])\)
if \(|M[m]|=n\) then \(X[m] \leftarrow \Sigma \oplus\) huge \(\cdot L\) else \(X[m] \leftarrow \Sigma\)
Equal \(\leftarrow\{i \in[1 . . \min \{m, \bar{m}\}-1]: M[i]=\bar{M}[i]\}\)
Unequal \(\leftarrow[1 . . \bar{m}] \backslash\) Equal
for \(i \leftarrow 1\) to \(\bar{m}-1\) do
    if \(i \in\) Equal then \(\{\bar{X}[i] \leftarrow X[i] ; \quad \bar{Y}[i] \leftarrow Y[i]\}\)
    if \(i \in\) Unequal then \(\left\{\bar{X}[i] \leftarrow \bar{M}[i] \oplus \gamma_{i} \cdot L ; \quad \bar{Y}[i] \stackrel{R}{\leftarrow}\{0,1\}^{n}\right\}\)
\(\bar{\Sigma} \leftarrow \bar{Y}[1] \oplus \cdots \oplus \bar{Y}[\bar{m}-1] \oplus \operatorname{pad}(\bar{M}[\bar{m}])\)
if \(|\bar{M}[\bar{m}]|=n\) then \(\bar{X}[\bar{m}] \leftarrow \bar{\Sigma} \oplus\) huge \(\cdot L\) else \(\bar{X}[\bar{m}] \leftarrow \bar{\Sigma}\)
\(D \leftarrow\left\{0^{n}, X[1], \ldots, X[m]\right\} \cup\{\bar{X}[i]: i \in\) Unequal \(\}\)
if there is a repetition in \(D\) then bad \(\leftarrow\) true
```

Figure 4: The pairwise-collision probability of $M$ and $\bar{M}$ is defined using this game. In line 21, $D$ is understood to be a multiset. The flag bad is set to true when a nontrivial collision occurs. Recall that huge is a synonym for $\mathrm{x}^{-1}$.
block cipher $E$ that achieves advantage $\operatorname{Adv}_{E}^{\text {prp }}(B) \geq \delta$. Adversary $B$ asks at most $q^{\prime}=\sigma+1$ oracle queries and has a running time which is equal to $A$ 's running time plus the time to compute $E$ on $q^{\prime}$ points, plus additional time which is $c n \sigma$, where the constant $c$ depends only on details of the model of computation.

It is a standard and easy result that being secure in the sense of a PRF implies an inability to forge with good probability. See $[5,17]$.

### 5.3 Structure of the Proof [DRAFT]

The proof of Theorem 1 combines two lemmas. The first lemma, the structure lemma, measures the pseudorandomness of PMAC in terms of a function we call the pairwise-collision probability, $\operatorname{coll}_{n}(m, \bar{m})$. This function captures the probability of a "nontrivial collision" between a pair of distinct messages $M$ and $\bar{M}$ having lengths $m$ and $\bar{m}$, respectively. The second lemma, the pairwisecollision bound, gives an upperbound on $\operatorname{coll}_{n}(m, \bar{m})$.

We begin by defining the pairwise-collision probability, $\operatorname{coll}_{n}(\cdot, \cdot)$. Fix $n$ and choose distinct messages $M$ and $\bar{M}$, and partition them into $M[1] \ldots M[m]$ and $\bar{M}[1] \ldots \bar{M}[\bar{m}]$. Consider the experiment of Figure 4. When bad is set to true in the last line of this experiment, we say that there has been a nontrivial collision between $M$ and $\bar{M}$. Let $\operatorname{coll}_{n}(M, \bar{M})$ denote the probability of a nontrivial collision between $M$ and $\bar{M}$, and let $\operatorname{coll}_{n}(m, \bar{m})$ denote the maximal value of $\operatorname{coll}_{n}(M, \bar{M})$ over all distinct $M$ and $\bar{M}$ having $m$ and $\bar{m}$ blocks, respectively.

We can now state the structure lemma.
Lemma 1 [Structure lemma] Fix $n$ and $\tau$. Let $A$ be an adversary who asks $q$ queries, these having aggregate length of $\sigma$ blocks. Let $\operatorname{coll}_{n}(\cdot, \cdot)$ denote the pairwise-collision probability. Then

$$
\operatorname{Adv}_{\operatorname{PMAC}[\operatorname{Perm}(n), \tau]}^{\operatorname{prf}}(A) \leq \max _{\substack{m_{1}, \ldots, m_{q} \\ \sigma=\sum m_{i}}}\left\{\sum_{1 \leq r<s \leq q} \operatorname{coll}_{n}\left(m_{r}, m_{s}\right)\right\}+\frac{(\sigma+1)^{2}}{2} \cdot \frac{1}{2^{n}}
$$

The proof of this lemma is in Appendix A.1.

Explanation. Informally, $\operatorname{coll}_{n}(m, \bar{m})$ bounds the probability of running into trouble when the adversary asks some two different oracle queries, $M$ and $\bar{M}$, having lengths $m$ and $\bar{m}$, respectively. Trouble (a nontrivial collision) is defined as follows. Consider the $m+\bar{m}+1$ points at which the block cipher is applied in processing $M$ and $\bar{M}$. There are $m$ points $X[1], \ldots, X[m]$, another $\bar{m}$ points $\bar{X}[1], \ldots, \bar{X}[\bar{m}]$, and then there is the point $0^{n}$ (the block cipher was applied at this point to define $L$ ). Some pairs of these $m+\bar{m}+1$ points could coincide for a "trivial" reason: namely, we know that $X[i]=\bar{X}[i]$ if $i<m$ and $i<\bar{m}$ and $M[i]=\bar{M}[i]$. We say that there is a nontrivial collision if any other pair of block-cipher inputs coincide. Rather pessimistically, we are going to "give up" in the analysis any time there is a nontrivial collision. The value $\operatorname{coll}_{n}(m, \bar{m})$ measures the probability of a nontrivial collision.

The structure lemma provides a simple recipe for measuring the maximal advantage of an adversary who attacks the pseudorandomness of PMAC: namely, bound the pairwise-collision probability and then use the formula. The lemma simplifies the analysis of PMAC in two ways. First, it allows one to excise adaptivity as a concern. Dealing with adaptivity is a major complicating factor in proofs of this type. Second, it allows one to concentrate on what happens to fixed pairs of messages. It is easier to think about what happens with two messages than what is happening with all $q$ of them.

Bounding the pairwise-collision probability. The following lemma indicates that nontrivial collisions rarely occur. Its proof is given in Appendix A.2.

Lemma 2 [Pairwise-collision bound] Let $\operatorname{coll}_{n}(\cdot, \cdot)$ denote the pairwise-collision probability. Then

$$
\operatorname{coll}_{n}(m, \bar{m}) \leq\binom{ m+\bar{m}+1}{2} \cdot \frac{1}{2^{n}}
$$

Concluding the theorem. The pseudorandomness of PMAC follows by combining Lemmas 1 and 2. Namely,

$$
\begin{aligned}
\operatorname{Adv}_{\mathrm{PMAC}[\operatorname{Perm}(n), \tau]}^{\operatorname{prf}} & \leq \max _{\substack{m_{1}, m_{q}, m_{q} \\
\sigma=\sum m_{i}}}\left\{\sum_{1 \leq r<s \leq q} \operatorname{coll}_{n}\left(m_{r}, m_{s}\right)\right\}+\frac{(\sigma+1)^{2}}{2^{n+1}} \\
& \leq \max _{\substack{m_{1}, \ldots m_{q} \\
\sigma=\Sigma m_{i}}}\left\{\sum_{1 \leq r<s \leq q}\binom{m_{r}+m_{s}+1}{2} \cdot \frac{1}{2^{n}}\right\}+\frac{(\sigma+1)^{2}}{2^{n+1}}
\end{aligned}
$$

which can be bounded by setting each $m_{i}$ to the (possibly non-integral) $m_{i}=\sigma / q$. (When $a$ is non-integral, define $\binom{a}{2}=a(a-1) / 2$.) This yields

$$
\begin{aligned}
\operatorname{Adv}_{\operatorname{PMAC}[\operatorname{Perm}(n), \tau]}^{\operatorname{prf}} & \leq\binom{ q}{2}\binom{2 \sigma / q+1}{2} \cdot \frac{1}{2^{n}}+\frac{(\sigma+1)^{2}}{2^{n+1}} \\
& \leq \frac{q^{2}}{2} \cdot \frac{4 \sigma^{2} / q^{2}+4 \sigma / q+1}{2^{n+1}}+\frac{(\sigma+1)^{2}}{2^{n+1}} \\
& \leq \frac{2 \sigma^{2}+2 \sigma q+0.5 q^{2}+\sigma^{2}+2 \sigma+1}{2^{n+1}} \\
& \leq \frac{1.5 \sigma^{2}}{2^{n}}+\frac{\sigma q+\sigma+0.25 q^{2}+0.5}{2^{n}}
\end{aligned}
$$

| Algorithm | 16 B | 128 B | 2 KB |
| :--- | ---: | ---: | ---: |
| PMAC | $\mathbf{2 2 . 1}$ | $\mathbf{1 8 . 7}$ | $\mathbf{1 8 . 4}$ |
| CBC MAC | 18.9 | 17.4 | 17.1 |

Figure 5: Performance results. Numbers are in cycles per byte, on a Pentium 3, for three message lengths, the code written in assembly. The underlying block cipher is AES-128. CBC MAC refers to the "basic" CBC MAC—no padding is performed and nothing extra is done to get security across messages of varying lengths.

Using the fact that $(\sigma+\Delta)^{2}-\sigma^{2} \geq 2 \sigma \Delta$ and $(\sigma+\Delta)^{2}-\sigma^{2} \geq \Delta^{2}$, if we increase $\sigma$ by $2 / 3(q / 2+$ $1 / 2+q / 2+3 / 4)=2 / 3(q+5 / 4) \leq q+1$, this will be enough to at least match the lower order terms. That is, letting $\bar{\sigma}=\sigma+q+1$, we have that the quantity above is at most $1.5 \bar{\sigma}^{2} / 2^{n}$, completing the proof of Theorem 1.

## 6 Performance

Abstract accounting. PMAC uses $\lceil|M| / n\rceil$ block-cipher invocations for any nonempty message $M$. (The empty string takes one block-cipher invocation). We compare with the CBC MAC:

- The "basic" CBC MAC, which assumes that the message is a nonzero multiple of the block length and which is only secure when all messages to be MACed are of one fixed length, uses the same number of block cipher calls: $|M| / n$.
- The version of the CBC MAC described in [10], which removes the two restrictions just mentioned, uses the same number of calls as PMAC, $\lceil|M| / n\rceil$.
- Obligatory padding (to support short-final-block messages) and standard methods to process the final block (double or triple encryption, to achieve security across variable-length messages) can raise the number of block-cipher calls to as much as $\lceil|M+1| / n\rceil+2$.

Thus PMAC saves between 0 and 3 block-cipher calls compared to the various versions of the CBC MAC.

As with any mode, there is further overhead beyond the block-cipher calls. Per block, this overhead is about three $n$-bit xor operations plus associated logic. The work for this associated logic will vary according to whether or not one precomputed $L(i)$-values, whether or not there is an $n t z()$ instruction available, and on other factors.

Though some of the needed $L(i)$-values are likely to be pre-computed, calculating these values "on the fly" is not too expensive. Starting with $0^{n}$ we form successive offsets by xoring the previous offset with $L, 2 \cdot L, L, 4 \cdot L, L, 2 \cdot L, L, 8 \cdot L$, and so forth. So half the time we use $L$ itself; a quarter of the time we use $2 \cdot L$; one eighth of the time we use $4 \cdot L$; and so forth. Thus the expected number of $a \cdot \mathrm{x}$-operations to compute an offset is at most $\sum_{i=1}^{\infty} i / 2^{i+1}=1$. Each $a \cdot \mathrm{x}$ instruction requires an $n$-bit xor and a conditional 32 -bit xor. Said differently, for any $m>0$, the total number of $a \cdot \mathrm{x}$ operations needed to compute $\gamma_{1} \cdot L, \gamma_{2} \cdot L, \ldots, \gamma_{m} \cdot L$ is $\sum_{i=1}^{m} n t z(i)$, which is less than $m$. The above assumed that one does not retain or precompute any $L(i)$ value beyond $L=L(0)$. Suppose that one retains a few values: $L(0), L(1), L(2), L(3)$. Computing and storing these three additional values is less overhead than computing $L$ itself, which required an application of $E_{K}$. But now the desired multiple of $L$ has already been computed $1 / 2+1 / 4+1 / 8+1 / 16 \approx 94 \%$ of the time. When it has not been pre-computed it must be calculated, starting from $8 \cdot L$, so the amortized number of doubling steps has thus been reduced from 1 to $\sum_{i=1}^{\infty}=i / 2^{i+4}=0.125$.

Experimental results. Our colleague Ted Krovetz implemented PMAC using AES-128 as the
underlying block cipher. He compared its performance in an entirely sequential setting to that of the CBC MAC. By the CBC MAC we mean the "basic" CBC MAC-nothing extra done to take care of length-variability or the possibility of strings which are not a multiple of the block length. The code was written in assembly. The OS was Windows 2000 sp 1 and the compiler was Visual C++6.0 sp4. All data fit into L1 cache.

Note that the basic CBC MAC needs to be modified to correctly handle length-variability and the possibility of short final blocks, and doing this is often done in a way that entails additional block-cipher calls. Ignoring these factors tends to make the performance comparison conservative.

Disregarding the one-block message in Figure 6, we see that PMAC adds about $8 \%$ overhead compared to the basic CBC MAC. We emphasize that this is for an entirely serial execution environment with a limited number of registers. In an environment with plenty of registers and multiple instruction pipes, PMAC, properly implemented, will of course be faster than the CBC MAC.

## 7 Intellectual Property Statement and Disclosures

Patent applications covering the ideas of this proposal were filed (Rogaway as inventor) on 13 September 2000, 12 October 2000, and 9 February 2001.

The inventor hereby releases IP rights covering PMAC for all non-commercial, non-governmental applications. For commercial applications, the inventor will license PMAC under a non-exclusive license, on a non-discriminatory basis, based on reasonable terms and conditions.

IBM has a patent covering the XOR MAC of [4]; this is US Patent \#5,673,318 (September 30, 1997). Though the submitter is in no position to render a legal opinion, it appears to him that the claims of US Patent $\# 5,673,318$ do not read upon the use of PMAC. Virgil Gligor indicates that he has made patent filings on 31 January 2000, 31 March 2000, and 24 August 2000. We know no more details.

The IP status specified in this submission will be updated when more is known.

## Acknowledgments

Thanks to Virgil Gligor, who described [13] to Rogaway at CRYPTO '00. This conversation inspired the development of PMAC. Special thanks to Ted Krovetz, who wrote reference code, prepared test vectors, collected performance data, and gave valuable interaction. Thanks to Michael Amling for a careful proofreading.

The first NIST modes-of-operation workshop, held in October 2000, was the motivation for putting together our first draft of PMAC [25]. The present writeup was prepared in response to a later call for proposals by NIST. We are pleased with NIST's leadership in promoting the development of modern modes of operation, and thank Elaine Barker, Morris Dworkin, and Jim Foti for their work in this connection.

This submission was prepared while Rogaway was on leave of absence from UC Davis, visiting the Department of Computer Science, Faculty of Science, Chiang Mai University. No external funding was used for this particular piece of research, but the submitter gratefully acknowledges a recent gift by Cisco Systems. Many thanks to Cisco for their support of my research.

## References

[1] M. Bellare, R. Canetti, and H. Krawczyk. Keying hash functions for message authentication. Advances in Cryptology - CRYPTO '96. Lecture Notes in Computer Science, vol. 1109, Springer-Verlag, pp. 1-15, 1996.
[2] M. Bellare, A. Desai, E. Jokipir, and P. Rogaway. A concrete security treatment of symmetric encryption: Analysis of the DES modes of operation. Proceedings of 38th Annual Symposium on Foundations of Computer Science (FOCS 97), IEEE, 1997. http://www.cs.ucdavis.edu/~rogaway/
[3] M. Bellare, S. Goldwasser, and O. Goldreich. Incremental cryptography and applications to virus protection. Proceedings of the 27th Annual ACM Symposium on the Theory of Computing (STOC '95). ACM Press, pp. 45-56, 1995.
[4] M. Bellare, R. Guérin and P. Rogaway, "XOR MACs: New methods for message authentication using finite pseudorandom functions." Advances in Cryptology - CRYPTO '95. Lecture Notes in Computer Science Vol. 963, Springer-Verlag, D. Coppersmith, Ed., pp. 15-28, 1995. http://www.cs.ucdavis.edu/~rogaway/
[5] M. Bellare, J. Kilian, and P. Rogaway. The security of the cipher block chaining message authentication code. Journal of Computer and System Sciences, vol. 61, no. 3, Dec 2000. (Full version of paper from Advances in Cryptology - CRYPTO '94. Lecture Notes in Computer Science, vol. 839, pp. 340-358, 1994.) http://www.cs.ucdavis.edu/~rogaway/
[6] M. Bellare and C. Namprempre. Authenticated encryption: Relations among notions and analysis of the generic composition paradigm. Advances in Cryptology ASIACRYPT '00. Lecture Notes in Computer Science, T. Okamoto., ed., Springer-Verlag, 2000. http://www-cse.ucsd.edu/users/mihir/
[7] M. Bellare and P. Rogaway. Encode-then-encipher encryption: How to exploit nonces or redundancy in plaintexts for efficient encryption. Advances in Cryptology ASIACRYPT '00. Lecture Notes in Computer Science, T. Okamoto., ed., Springer-Verlag, 2000. http://www.cs.ucdavis.edu/~rogaway/
[8] A. Berendschot, B. den Boer, J.P. Boly, A. Bosselaers, J. Brandt, D. Chaum, I. Damgård, M. Dichtl, W. Fumy, M. van der Ham, C.J.A. Jansen, P. Landrock, B. Preneel, G. Roelofsen, P. de Rooij, and J. Vandewalle. Final report of race integrity primitives. Lecture Notes in Computer Science, vol. 1007, Springer-Verlag,
[9] D. Bernstein. How to stretch random functions: the security of protected counter sums. Journal of Cryptology, 12(3): pp. 185-192 (1999). http://cr.yp.to/djb.html
[10] J. Black and P. Rogaway. CBC MACs for arbitrary-length messages: the three-key constructions. Full version of paper from Advances in Cryptology - CRYPTO 2000. Lecture Notes in Computer Science, vol. 1880, pp. 197-215, 2000. http://www.cs.ucdavis.edu/~rogaway/
[11] J. Black, S. Halevi, H. Krawczyk, T. Krovetz, and P. Rogaway. UMAC: Fast and secure message authentication. Advances in Cryptology - CRYPTO '99. Lecture Notes in Computer Science, Springer-Verlag, 1999.
[12] L. Carter and M. Wegman. Universal hash functions. J. of Computer and System Sciences. vol. 18, pp. 143-154, 1979.
[13] V. Gligor and P. Donescu. Fast encryption and authentication: XCBC encryption and XECB authentication modes. Manuscript. August 18, 2000. http://www.eng.umd.edu/~gligor/
[14] V. Gligor and P. Donescu. Fast encryption and authentication: XCBC encryption and XECB authentication modes. Manuscript. October 2000. http://csrc.nist.gov/encryption/aes/modes/
[15] O. Goldreich, S. Goldwasser, and S. Micali. How to construct random functions. Journal of the ACM, vol. 33, no. 4, pp. 210-217, 1986.
[16] S. Goldwasser and S. Micali. Probabilistic encryption. Journal of Computer and System Sciences, vol. 28, April 1984, pp. 270-299.
[17] S. Goldwasser, S. Micali, and R. Rivest. A digital signature scheme secure against adaptive chosen-message attacks. SIAM Journal of Computing, vol. 17, no. 2, April 1988, pp. 281-308.
[18] S. Halevi and H. Krawczyk. MMH: Software message authentication in the Gbit/second rates. Proceedings of the 4th Workshop on Fast Software Encryption, Lecture Notes in Computer Science, vol. 1267, Springer-Verlag, pp. 172-189, 1997.
[19] Iso/Iec 9797-1. Information technology - Security techniques - Data integrity mechanism using a cryptographic check function employing a block cipher algorithm. Second edition. International Organization for Standards (ISO), Geneva, Switzerland, 1999.
[20] C. Jutla. Encryption modes with almost free message integrity. Manuscript. August 1, 2000. http://eprint.iacr.org/ (reference number 2000/039). Proceedings version to appear in Eurocrypt 2001.
[21] H. Krawczyk. LFSR-based hashing and authentication. Advances in Cryptology-CRYPTO '94 Lecture Notes in Computer Science, vol. 839, Springer-Verlag, pp 129-139, 1994.
[22] M. Luby and C. Rackoff. How to construct pseudorandom permutations from pseudorandom functions. SIAM J. Computation, vol. 17, no. 2, April 1988.
[23] E. Petrank and C. Rackoff. CBC MAC for real-time data sources. Manuscript 97-10 (1997) in http://philby.ucsd.edu/cryptolib.html 1995.
[24] B. Preneel. Cryptographic primitives for information authentication - State of the art. State of the Art in Applied Cryptography, COSIC '97, LNCS 1528, B. Preneel and V. Rijmen, eds., Springer-Verlag, pp. 49-104, 1998.
[25] P. Rogaway. PMAC: A parallelizable message authentication code. Contribution to NIST. October 16, 2000. Earlier version of the present paper. http://csrc.nist.gov/encryption/modes/workshop1/
[26] P. Rogaway (submitter), and M. Bellare, J. Black, and T. Krovetz (auxiliary submitters). OCB Mode. Contribution to NIST (submitted along with the current submission). April 1, 2001. http://csrc.nist.gov/encryption/modes/
[27] US National Bureau of Standards. DES Modes of Operation. Federal Information Processing Standard (FIPS) Publication 81, December 1980. Available as http://www.itl.nist.gov/fipspubs/fip81.htm
[28] M. Wegman and L. Carter. New hash functions and their use in authentication and set equality. J. of Comp. and System Sciences. vol. 22, pp. 265-279, 1981.

```
Initialization
\(10 \quad L \stackrel{R}{\leftarrow}\{0,1\}^{n} ; \quad \pi\left(0^{n}\right) \leftarrow L\)
When \(A\) makes its \(r\)-th query, \(M_{r}=M_{r}[1] \cdots M_{r}\left[m_{r}\right]\), where \(r \in[1 . . q]\)
\(20 \quad\) for \(i \leftarrow 1\) to \(m_{r}-1\) do
\(21 \quad X_{r}[i] \leftarrow M_{r}[i] \oplus \gamma_{i} \cdot L\)
\(22 \quad\) if \(X_{r}[i] \in \operatorname{Domain}(\pi)\) then \(Y_{r}[i] \leftarrow \pi\left(X_{r}[i]\right)\)
\(23 \quad\) else \(Y_{r}[i] \stackrel{R}{\leftarrow}\{0,1\}^{n}\)
\(24 \quad\) if \(Y_{r}[i] \in \operatorname{Range}(\pi)\) then \(\left\{\operatorname{bad} \leftarrow\right.\) true; \(\left.Y_{r}[i] \stackrel{R}{\leftarrow}\{0,1\}^{n}\right\}\)
\(25 \quad \pi\left(X_{r}[i]\right) \leftarrow Y_{r}[i]\)
\(26 \quad \Sigma_{r} \leftarrow Y_{r}[1] \oplus Y_{r}[2] \oplus \cdots \oplus Y_{r}\left[m_{r}-1\right] \oplus \operatorname{pad}\left(M_{r}\left[m_{r}\right]\right)\)
\(27 \quad\) if \(\left|M_{r}\left[m_{r}\right]\right|=n\) then \(X_{r}\left[m_{r}\right] \leftarrow \Sigma_{r} \oplus\) huge \(\cdot L\) else \(X_{r}\left[m_{r}\right] \leftarrow \Sigma_{r}\)
28 if \(X_{r}\left[m_{r}\right] \in \operatorname{Domain}(\pi)\) then \(\left\{\right.\) bad \(\leftarrow\) true; \(\left.T A G_{r} \stackrel{R}{\leftarrow} \overline{\operatorname{Range}}(\pi)\right\}\)
29 else \(T A G_{r} \stackrel{R}{\leftarrow}\{0,1\}^{n}\)
        if \(T A G_{r} \in \operatorname{Range}(\pi)\) then \(\left\{\right.\) bad \(\leftarrow\) true; \(\left.T A G_{r} \stackrel{R}{\leftarrow} \overline{\operatorname{Range}}(\pi)\right\}\)
            \(\pi\left(X_{r}\left[m_{r}\right]\right) \leftarrow T A G_{r}\)
        \(T a g_{r} \leftarrow T A G_{r} \quad[\) first \(\tau\) bits]
        return \(\mathrm{Tag}_{r}\)
```

Figure 6: Game 1. This game accurately simulates $\operatorname{PMAC}[\operatorname{Perm}(n), \tau]$.

## A Proofs [DRAFT]

## A. 1 Proof of the Structure Lemma (Lemma 1)

Let $A$ be an adversary that attacks $\operatorname{PMAC}[\operatorname{Perm}(n), \tau]$. Since $A$ is computationally unbounded, there is no loss of generality to assume that $A$ is deterministic. One can imagine $A$ interacting with a $\operatorname{PMAC}[\operatorname{Perm}(n), \tau]$ oracle as $A$ playing a certain game, Game 1, as defined in Figure 6. This game perfectly simulates the behavior of $\operatorname{PMAC}[\operatorname{Perm}(n), \tau]$. It does so in a somewhat unusual way, sometimes setting a flag bad to true. We observe that if the flag bad is not set to true in an execution of the game, then the value $T a g_{r}$ returned by the game at line 33 is a random one - the first $\tau$ bits of the string randomly selected at line 29. It follows that $\mathbf{A d v}_{\operatorname{PMAC}[\operatorname{Perm}(n), \tau]}(A)$ is at most the probability that bad gets set to true in Game 1. The rest of the proof is devoted to bounding this probability.

We first consider the probability that bad gets set to true in line 24 or 30 . In both cases, we have just chosen a random $n$-bit string and then we are testing it for membership in a set. The size of this set starts at 1 (after executing line 10) and grows one element at a time until, by the time just before the last addition of a point, it has size $\sigma$. Thus we have that

$$
\begin{align*}
\mathrm{Pr}_{1}[\text { bad gets set in lines } 24 \text { or } 30] & \leq \frac{1+2+\ldots+\sigma}{2^{n}} \\
& \leq \frac{(\sigma+1)^{2}}{2^{n+1}} \tag{1}
\end{align*}
$$

Here the subscript of 1 in the probability reminds us that we are considering the behavior of Game 1.
We can now modify Game 1 by changing the behavior when and only when bad is set, and adding as a compensating factor the bound given by Equation (1). In particular, we may simply omit lines 24 and 30 , and the second statement in the compound statement of line 28 along with

```
Initialization
\(10 \quad L \stackrel{R}{\leftarrow}\{0,1\}^{n} ; \quad \pi\left(0^{n}\right) \leftarrow L\)
When \(A\) makes its \(r\)-th query, \(M_{r}=M_{r}[1] \cdots M_{r}\left[m_{r}\right]\), where \(r \in[1 . . q]\)
\(20 \quad\) for \(i \leftarrow 1\) to \(m_{r}-1\) do
\(X_{r}[i] \leftarrow M_{r}[i] \oplus \gamma_{i} \cdot L\)
22 if \(X_{r}[i] \in \operatorname{Domain}(\pi)\) then \(Y_{r}[i] \leftarrow \pi\left(X_{r}[i]\right)\)
else \(\left\{Y_{r}[i]^{R}{ }^{R}\{0,1\}^{n} ; \quad \pi\left(X_{r}[i]\right) \leftarrow Y_{r}[i]\right\}\)
\(24 \quad \Sigma_{r} \leftarrow Y_{r}[1] \oplus Y_{r}[2] \oplus \cdots \oplus Y_{r}\left[m_{r}-1\right] \oplus \operatorname{pad}\left(M_{r}\left[m_{r}\right]\right)\)
\(25 \quad\) if \(\left|M_{r}\left[m_{r}\right]\right|=n\) then \(X_{r}\left[m_{r}\right] \leftarrow \Sigma_{r} \oplus\) huge \(\cdot L\) else \(X_{r}\left[m_{r}\right] \leftarrow \Sigma_{r}\)
\(26 \quad\) if \(X_{r}\left[m_{r}\right] \in \operatorname{Domain}(\pi)\) then bad \(\leftarrow\) true
\(27 \quad T A G_{r}\left[m_{r}\right] \stackrel{R}{\leftarrow}\{0,1\}^{n}\)
\(28 \pi\left(X_{r}\left[m_{r}\right]\right) \leftarrow T A G_{r}\left[m_{r}\right]\)
\(29 \quad \mathrm{Tag}_{r} \leftarrow T A G_{r}\) [first \(\tau\) bits]
30 return \(\mathrm{Tag}_{r}\)
```

Figure 7: Game 2. A simplification of Game 1. Bounding the probability that bad gets set in this game, and then adding a correction factor, serves to bound $\mathbf{A d v} \mathbf{v}_{\mathrm{PMAC}[\operatorname{Perm}(n), \tau]}^{\mathrm{prf}}$.
the following else. The modified game is rewritten in Figure 7. At this point we know that

$$
\begin{equation*}
\operatorname{Adv}_{\mathrm{PMAC}[\operatorname{Perm}(n), \tau]}^{\mathrm{prf}}(A) \leq \operatorname{Pr}_{2}[\text { bad gets set }]+\frac{(\sigma+1)^{2}}{2^{n+1}} \tag{2}
\end{equation*}
$$

Notice in Game 2 that the value $\operatorname{Tag}_{r}$ returned in response to a query $M$ is always a random $\tau$-bit string. But of course the game does more than just return these strings: it also chooses $L$ at random, fills in $\pi$-values, and sets bad under certain conditions. We can defer doing all those things, and just return the random strings $\operatorname{Tag}_{1}, \ldots, \mathrm{Tag}_{q}$. This does not change the view of the adversary that interacts with the game, nor will it change the probability that bad is set to true. The modified game is called Game 3, and it is depicted in Figure 8.

We need to bound the probability that bad gets set to true in Game 3. This probability is over the random $T A G_{r}$-values selected at line 10 , the random value of $L$ selected at line 20 , and the random $Y_{r}[i]$-values selected at line 25 . We want to show that, over these random values, bad will rarely be set. In fact, we will claim something stronger: that even if one arbitrarily fixes the values of $T A G_{1}, \ldots, T A G_{q} \in\{0,1\}^{n}$ (and takes the probability over just the remaining values), still the probability that bad will be set to true is small. Now since the oracle responses have now been fixed, and since the adversary itself is deterministic, the queries $M_{1}, \ldots, M_{q}$ that the adversary will generate have likewise been fixed. Interaction and the adversary itself are essentially gone at this point, replaced by universal quantification. The new game is show in Figure 9. It depends on constants $\mathcal{C}=\left(q, T A G_{1}, \ldots, T A G_{q}, M_{1}, \ldots, M_{q}\right)$. At this point in the proof we have that

$$
\begin{equation*}
\operatorname{Adv}_{\operatorname{Perm}(n), \tau}^{\operatorname{prf}}(A) \leq \max _{\mathcal{C}}\left\{\operatorname{Pr}[\text { bad gets set to true in game } 4[\mathcal{C}]\}+\frac{(\sigma+1)^{2}}{2^{n+1}}\right. \tag{3}
\end{equation*}
$$

where, if $A$ is limited to $q$ queries of aggregate length $\sigma$, then $\mathcal{C}$ specifies $q$, strings $M_{1}, \ldots, M_{q}$ of aggregate block length $\sigma$, and $T A G_{1}, \ldots, T A G_{q} \in\{0,1\}^{n}$.

The next step is to modify Game 4 so that the new game, Game 5 , sets bad every time that Game 4 does, plus some additional times. Look at line 14 in Game 4 . The value $X_{r}[i]$ could have

```
When \(A\) makes its \(r\)-th query, \(M_{r}=M_{r}[1] \cdots M_{r}\left[m_{r}\right]\), where \(r \in[1 . . q]\)
\(10 \quad\) TAG \({ }_{r} \stackrel{R}{\leftarrow}\{0,1\}^{n}\)
11 return \(T A G_{r}\) [first \(\tau\) bits]
When \(A\) is done making its \(q\) queries
20
    \(L \stackrel{R}{\leftarrow}\{0,1\}^{n} ; \quad \pi\left(0^{n}\right) \leftarrow L\)
    for \(r \leftarrow 1\) to \(q\) do
        for \(i \leftarrow 1\) to \(m_{r}-1\) do
            \(X_{r}[i] \leftarrow M_{r}[i] \oplus \gamma_{i} \cdot L\)
            if \(X_{r}[i] \in \operatorname{Domain}(\pi)\) then \(Y_{r}[i] \leftarrow \pi\left(X_{r}[i]\right)\)
            else \(\left\{Y_{r}[i] \stackrel{R}{\leftarrow}\{0,1\}^{n} ; \quad \pi\left(X_{r}[i]\right) \leftarrow Y_{r}[i]\right\}\)
        \(\Sigma_{r} \leftarrow Y_{r}[1] \oplus Y_{r}[2] \oplus \cdots \oplus Y_{r}\left[m_{r}-1\right] \oplus \operatorname{pad}\left(M_{r}\left[m_{r}\right]\right)\)
        if \(\left|M_{r}\left[m_{r}\right]\right|=n\) then \(X_{r}\left[m_{r}\right] \leftarrow \Sigma_{r} \oplus\) huge \(\cdot L\) else \(X_{r}\left[m_{r}\right] \leftarrow \Sigma_{r}\)
        if \(X_{r}\left[m_{r}\right] \in \operatorname{Domain}(\pi)\) then \(\mathrm{bad} \leftarrow\) true
        \(\pi\left(X_{r}\left[m_{r}\right]\right) \leftarrow T A G_{r}\)
```

Figure 8: Game 3. Like Game 2, but we defer all but the selection of $T A G_{r}$-values.

```
\(L \stackrel{R}{\leftarrow}\{0,1\}^{n} ; \quad \pi\left(0^{n}\right) \leftarrow L\)
for \(r \leftarrow 1\) to \(q\) do
    for \(i \leftarrow 1\) to \(m_{r}-1\) do
                \(X_{r}[i] \leftarrow M_{r}[i] \oplus \gamma_{i} \cdot L\)
                if \(X_{r}[i] \in \operatorname{Domain}(\pi)\) then \(Y_{r}[i] \leftarrow \pi\left(X_{r}[i]\right)\)
                    else \(\left\{Y_{r}[i] \stackrel{R}{\left.\stackrel{R}{\leftarrow}\{0,1\}^{n} ; \quad \pi\left(X_{r}[i]\right) \leftarrow Y_{r}[i]\right\}}\right.\)
        \(\Sigma_{r} \leftarrow Y_{r}[1] \oplus Y_{r}[2] \oplus \cdots \oplus Y_{r}\left[m_{r}-1\right] \oplus \operatorname{pad}\left(M_{r}\left[m_{r}\right]\right)\)
        if \(\left|M_{r}\left[m_{r}\right]\right|=n\) then \(X_{r}\left[m_{r}\right] \leftarrow \Sigma_{r} \oplus\) huge \(\cdot L\) else \(X_{r}\left[m_{r}\right] \leftarrow \Sigma_{r}\)
        if \(X_{r}\left[m_{r}\right] \in \operatorname{Domain}(\pi)\) then \(\mathrm{bad} \leftarrow\) true
        \(\pi\left(X_{r}\left[m_{r}\right]\right) \leftarrow T A G_{r}\)
```

Figure 9: Game $4[\mathcal{C}]$. This game depends on constants $\mathcal{C}$ which specify: $q, \operatorname{TAG}_{1}, \ldots, \operatorname{TAG}_{q} \in\{0,1\}^{n}$, and $M_{1}=M_{1}[1] \cdots M_{q}\left[m_{q}\right], \ldots, M_{q}=M_{q}[1] \cdots M_{q}\left[m_{q}\right]$.

```
\(10 \quad L \stackrel{R}{\leftarrow}\{0,1\}^{n} ; \quad \pi\left(0^{n}\right) \leftarrow L\)
\(11 \quad\) for \(r \leftarrow 1\) to \(q\) do
        for \(i \leftarrow 1\) to \(m_{r}-1\) do
        \(X_{r}[i] \leftarrow M_{r}[i] \oplus \gamma_{i} \cdot L ; \quad Y_{r}[i] \stackrel{R}{\leftarrow}\{0,1\}^{n}\)
        if \(M_{r}[i]=M_{s}[i]\) for some \(s<r\) and \(i<m_{s}\) then \(Y_{r}[i] \leftarrow \pi\left(X_{s}[i]\right)\)
        else if \(X_{r}[i] \in \operatorname{Domain}(\pi)\) then \(\mathrm{bad} \leftarrow\) true
        \(\pi\left(X_{r}[i]\right) \leftarrow Y_{r}[i]\)
        \(\Sigma_{r} \leftarrow Y_{r}[1] \oplus Y_{r}[2] \oplus \cdots \oplus Y_{r}\left[m_{r}-1\right] \oplus \operatorname{pad}\left(M_{r}\left[m_{r}\right]\right)\)
        if \(\left|M_{r}\left[m_{r}\right]\right|=n\) then \(X_{r}\left[m_{r}\right] \leftarrow \Sigma_{r} \oplus\) huge \(\cdot L\) else \(X_{r}\left[m_{r}\right] \leftarrow \Sigma_{r}\)
        if \(X_{r}\left[m_{r}\right] \in \operatorname{Domain}(\pi)\) then bad \(\leftarrow\) true
        \(\pi\left(X_{r}\left[m_{r}\right]\right) \leftarrow 0^{n}\)
```

Figure 10: Game $\mathbf{5}[\mathcal{C}]$. This game set bad at least as often as Game $4[\mathcal{C}]$ does. It is this game that is related back to the one that defines the pairwise-collision probability.

```
\(L \stackrel{R}{\leftarrow}\{0,1\}^{n} ; \quad \pi\left(0^{n}\right) \leftarrow L\)
for \(i \leftarrow 1\) to \(m_{s}-1\) do
    \(X_{s}[i] \leftarrow M_{s}[i] \oplus \gamma_{i} \cdot L\)
    if \(X_{s}[i] \in \operatorname{Domain}(\pi)\) then \(\mathrm{bad} \leftarrow\) true
    \(Y_{s}[i] \stackrel{R}{\leftarrow}\{0,1\}^{n} ; \quad \pi\left(X_{s}[i]\right) \leftarrow Y_{s}[i]\)
\(\Sigma_{s} \leftarrow Y_{s}[1] \oplus Y_{s}[2] \oplus \cdots \oplus Y_{s}\left[m_{s}-1\right] \oplus \operatorname{pad}\left(M_{s}\left[m_{s}\right]\right)\)
if \(\left|M_{s}\left[m_{s}\right]\right|=n\) then \(X_{s}\left[m_{s}\right] \leftarrow \Sigma_{s} \oplus\) huge \(\cdot L\) else \(X_{s}\left[m_{s}\right] \leftarrow \Sigma_{s}\)
if \(X_{s}\left[m_{s}\right] \in \operatorname{Domain}(\pi)\) then bad \(\leftarrow\) true
\(\pi\left(X_{r}\left[m_{r}\right]\right) \leftarrow 0^{n}\)
for \(i \leftarrow 1\) to \(m_{u}-1\) do
    \(X_{u}[i] \leftarrow M_{u}[i] \oplus \gamma_{i} \cdot L ; \quad Y_{u}[i] \stackrel{R}{\leftarrow}\{0,1\}^{n}\)
    if \(i<m_{s}\) and \(M_{u}[i]=M_{s}[i]\) then \(Y_{u} \leftarrow \pi\left(X_{s}[i]\right)\) else
    if \(X_{u}[i] \in \operatorname{Domain}(\pi)\) then bad \(\leftarrow\) true
    \(\pi\left(X_{u}[i]\right) \leftarrow Y_{u}[i]\)
\(\Sigma_{u} \leftarrow Y_{u}[1] \oplus Y_{u}[2] \oplus \cdots \oplus Y_{u}\left[m_{u}-1\right] \oplus \operatorname{pad}\left(M_{u}\left[m_{u}\right]\right)\)
if \(\left|M_{u}\left[m_{u}\right]\right|=n\) then \(X_{u}\left[m_{u}\right] \leftarrow \Sigma_{u} \oplus\) huge \(\cdot L\) else \(X_{u}\left[m_{u}\right] \leftarrow \Sigma_{u}\)
if \(X_{u}\left[m_{u}\right] \in \operatorname{Domain}(\pi)\) then \(\mathrm{bad} \leftarrow\) true
```

Figure 11: Game 6[ $\left.M_{s}, M_{u}\right]$.
been in the domain of $\pi$ for the "trivial" reason that $M_{r}[i]=M_{s}[i]$ for some $s<r$ and where $i<m_{s}$, or for some other, "non-trivial" reason: $X_{r}[i]=0^{n}$, or $X_{r}[i]=X_{s}[j]$ for some $s<r$ and $j \neq i$, or $X_{r}[i]=X_{r}[j]$ for some $j<i$. If $X_{r}[i]$ was in the domain of $\pi$ for a non-trivial reason, we effectively give up, setting bad to true. Thus in Game 5 , line 15 , we set bad if $X_{r}[i]$ is already in the domain of $\pi$ and it is not due to the trivial cause (which is tested for in line 14).

Let us name coins used in Game 5 by $L$ and $R_{1}=R_{1}[1] \cdots R_{1}\left[m_{1}\right], \ldots, R_{q}=R_{q}[1] \cdots R_{1}\left[m_{q}\right]$. (We use coins $L$ at line 10, and then we use coins $R_{r}[i]$ at line 13.) When running Game 5 on this vector of coins, if the flag bad get set to true then it gets set to true because there is a smallest $s \in[1 . . q]$ such that, using the specified coins,

- $\quad X_{s}[i]=0^{n}$, or
- $\quad X_{s}[i]=X_{u}[j]$ where $u<s$ and $j \neq i$, or
- $\quad X_{s}[i]=X_{s}[j]$, where $j<i$.

In particular, if we execute game 5 but restrict line 11 to only use the particular values $r \in\{s, u\}$, and we use coins $L, R_{s}, R_{u}$, the flag bad will still be set to true in this execution. Let Game 6 denote this game. We have rewritten it, making some irrelevant simplifications, in Figure 11. (In particular, the image assigned to $\pi\left(X_{s}\left[m_{s}\right]\right)$ is irrelevant in this program-all that matters is that the domain point be defined.) At this point we know that if Game $5[\mathcal{C}]$ sets bad to true using coins $L, R_{1}, \ldots, R_{q}$, then there is some $s<u$ such that Game 6 sets bad to true using coins $L, R_{s}, R_{q}$. Letting $\operatorname{coll}_{n}^{\prime}\left(m_{s}, m_{u}\right)$ be the maximal probability that Game 6 will cause bad to be set to true when using messages $M_{s}$ and $M_{u}$ having $m_{s}$ and $m_{u}$ blocks, we have that the probability that Game $5\left[M_{1}, \ldots, M_{q}\right]$ gets set to true is at most $\sum_{m_{s}, m_{u}} \operatorname{coll}_{n}^{\prime}\left(m_{s}, m_{u}\right)$. But game $6\left[m_{s}, m_{u}\right]$ is equivalent to the game that defined the pairwise-collision probability in Figure 4. In particular, a pair of messages and a given set of coins causes a collision in one game if and only it causes a collision in the other. Thus $\operatorname{coll}_{n}^{\prime}(m, \bar{m})=\operatorname{coll}_{n}(m, \bar{m})$, and the proof is complete.

## A. 2 Proof of the Pairwise-Collision Bound (Lemma 2)

We show that for any two points of

$$
D=\left\{0^{n}, X[1], \ldots, X[m]\right\} \cup\{\bar{X}[i]: i \in \text { Unequal }\}
$$

the probability that they coincide is at most $2^{-n}$. The lemma will follow since there are at most $m+\bar{m}+1$ points listed above, so at most $\binom{m+\bar{m}+1}{2}$ pairs of these points.

Let Unequal ${ }^{\prime}=$ Unequal $\backslash\{X[\bar{m}]\}$ (multiset difference: remove one copy of $X[\bar{m}]$ ). The points of $D$ are the multiset union of the following multisets:

1. $D_{1}=\left\{0^{n}\right\}$
2. $D_{2}=\{X[1], \ldots, X[m-1]\}$
3. $D_{3}=\{X[m]\}$
4. $D_{4}=\left\{\bar{X}[j]: j \in\right.$ Unequal $\left.^{\prime}\right\}$
5. $D_{5}=\{X[\bar{m}]\}$

We now go through all of the $\binom{5}{2}+2=12$ types of pairings (that is, an element of $D_{i}$ with an element of $D_{j}$, for $i<j$; or two elements from $D_{2}$ or two elements from $D_{4}$ ). For each type of pairing we show that the probability of a collision is at most $2^{-n}$.

CASE 1: $\left(D_{1}, D_{2}\right): \quad \operatorname{Pr}\left[0^{n}=X[i]\right]=\operatorname{Pr}\left[M[i] \oplus \gamma_{i} \cdot L=0^{n}\right]=\operatorname{Pr}\left[L=\gamma_{i}^{-1} \cdot M[i]\right]=2^{-n}$. We have used the fact that $\gamma_{i}$ is nonzero and we are working in a field. (We will continue to use this without mention.)

CASE 2: $\left(D_{1}, D_{3}\right)$ : If $|M[m]|<n$ and $m \geq 2$ then $\Sigma$ is a random $n$-bit string and so $\operatorname{Pr}\left[0^{n}=\right.$ $X[m]]=\operatorname{Pr}\left[0^{n}=\Sigma\right]=\operatorname{Pr}\left[0^{n}=Y[1] \oplus \cdots Y[m-1] \oplus \operatorname{pad}(M[m])\right]=2^{-n}$. If $|M[m]|=n$ and $m \geq 2$ then $\Sigma$ is a random $n$-bit string that is independent of $L$ and so $\operatorname{Pr}\left[0^{n}=X[m]\right]=\operatorname{Pr}\left[0^{n}=\right.$ $\Sigma \oplus$ huge $\cdot L]=2^{-n}$. If $|M[m]|<n$ and $m=1$ then $\operatorname{Pr}\left[0^{n}=X[1]\right]=\operatorname{Pr}\left[0^{n}=\operatorname{pad}(M[m])\right]=0$. If $|M[m]|=n$ and $m=1$ then $\operatorname{Pr}\left[0^{n}=X[1]\right]=\operatorname{Pr}\left[0^{n}=\operatorname{pad}(M[m]) \oplus\right.$ huge $\left.\cdot L\right]=2^{-n}$.

CASE 3: $\left(D_{1}, D_{4}\right): \quad \operatorname{Pr}\left[0^{n}=\bar{X}[i]\right]=\operatorname{Pr}\left[\bar{M}[i] \oplus \gamma_{i} \cdot L=0^{n}\right]=2^{-n}$.
CASE 4: $\left(D_{1}, D_{5}\right): \quad$ If $M[\bar{m}]<n$ and $\bar{m} \geq 2$ then $\bar{\Sigma}$ is a random $n$-bit string and so $\operatorname{Pr}\left[0^{n}=\right.$ $\bar{X}[\bar{m}]]=\operatorname{Pr}\left[0^{n}=\bar{\Sigma}\right]=2^{-n}$. If $M[\bar{m}]=n$ and $\bar{m} \geq 2$ then $\bar{\Sigma}$ is a random $n$-bit string which is independent of $L$, and so $\operatorname{Pr}\left[0^{n}=\bar{X}[\bar{m}]\right]=\operatorname{Pr}\left[0^{n}=\bar{\Sigma} \oplus\right.$ huge $\left.\cdot L\right]=2^{-n}$. If $M[\bar{m}]<n$ and $\bar{m}=1$ then $\operatorname{Pr}\left[0^{n}=\bar{X}[1]\right]=\operatorname{Pr}\left[0^{n}=\operatorname{pad}(M[1])\right]=0$. If $M[\bar{m}]=n$ and $\bar{m}=1$ then $\operatorname{Pr}\left[0^{n}=\bar{X}[1]\right]=\operatorname{Pr}\left[0^{n}=\operatorname{pad}(M[1]) \oplus\right.$ huge $\left.\cdot L\right]=2^{-n}$.

Case 5: $\left(D_{2}, D_{2}\right): \quad$ For $i, j \in[1 . . m-1], i<j, \operatorname{Pr}[X[i]=X[j]]=\operatorname{Pr}\left[M[i] \oplus \gamma_{i} \cdot L=M[j] \oplus \gamma_{j} \cdot L\right]=$ $\operatorname{Pr}\left[M[i] \oplus M[j]=\left(\gamma_{i} \oplus \gamma_{j}\right) \cdot L\right]=2^{-n}$ because $\gamma_{i} \neq \gamma_{j}$ for $i \neq j$. (Here one assumes that $j<2^{n}$ because the lemma gives a non-result anyway if $j$ were larger.)

CASE 6: $\left(D_{2}, D_{3}\right)$ : Assume that $m \geq 2$, for otherwise there is nothing to show. Suppose first that $|M[m]|<n$. Then $\operatorname{Pr}[X[i]=X[m]]=\operatorname{Pr}\left[M[i] \oplus \gamma_{i} \cdot L=\Sigma\right]$. The value $\Sigma$ is uniformly random and independent of $L$, so this probability is $2^{-n}$. Suppose next that $|M[m]|=n$. Then $\operatorname{Pr}[X[i]=X[m]]=\operatorname{Pr}\left[M[i] \oplus \gamma_{i} \cdot L=\Sigma \oplus\right.$ huge $\left.\cdot L\right]=\operatorname{Pr}\left[M[i] \oplus \Sigma=\left(\gamma_{i} \oplus\right.\right.$ huge $\left.) \cdot L\right]$. This value is $2^{-n}$ since $\gamma_{i} \neq$ huge. Here we are assuming that $i<2^{n-1}$, which is without loss of generality since a larger value of $i$, and therefore $m$, would give a non-result in the theorem statement.

CASE 7: $\left(D_{2}, D_{4}\right):$ Let $i \in[1 . . m-1]$ and $j \in U n e q u a l^{\prime}$ and consider $\operatorname{Pr}[X[i]=\bar{X}[j]]=\operatorname{Pr}[M[i] \oplus$ $\left.\gamma_{i} \cdot L=\bar{M}[j] \oplus \gamma_{j} \cdot L\right]=\operatorname{Pr}\left[M[i] \oplus \bar{M}[j]=\left(\gamma_{i} \oplus \gamma_{j}\right) \cdot L\right.$. If $i \neq j$ then $\gamma_{i} \neq \gamma_{j}$ and this probability is $2^{-n}$. If $i=j$ then the probability is 0 since, necessarily, $M[i] \neq \bar{M}[j]$.

CASE 8: $\left(D_{2}, D_{5}\right)$ : Suppose that $|\bar{M}[\bar{m}]|<n$. Then $\operatorname{Pr}[X[i]=\bar{X}[\bar{m}]]=\operatorname{Pr}\left[M[i] \oplus \gamma_{i} \cdot L=\bar{\Sigma}\right]=2^{-n}$ because $\bar{\Sigma}$ is independent of $L$. Suppose that $|\bar{M}[\bar{m}]|=n$. Then $\operatorname{Pr}[X[i]=\bar{X}[\bar{m}]]=\operatorname{Pr}[M[i] \oplus$ $\gamma_{i} \cdot L=\bar{\Sigma} \oplus$ huge $\left.\cdot L\right]=\operatorname{Pr}\left[M[i] \oplus \bar{\Sigma}=\left(\gamma_{i} \oplus\right.\right.$ huge $\left.) \cdot L\right]=2^{-n}$ because $\bar{\Sigma}$ is independent of $L$ and $\gamma_{i} \neq$ huge.

CASE 9: $\left(D_{3}, D_{4}\right):$ Suppose that $|M[m]|<n$. Then $\operatorname{Pr}[X[m]=\bar{X}[j]]=\operatorname{Pr}\left[\Sigma=\bar{M}[j] \oplus \gamma_{j} \cdot L\right]=2^{-n}$ because $\Sigma$ is independent of $L$. Suppose that $|M[m]|=n$. Then $\operatorname{Pr}[X[m]=\bar{X}[j]]=\operatorname{Pr}[\Sigma \oplus$ huge. $\left.L=\bar{M}[j] \oplus \gamma_{j} \cdot L\right]=\operatorname{Pr}\left[\Sigma \oplus \bar{M}[j]=\left(\gamma_{j} \oplus\right.\right.$ huge $\left.) \cdot L\right]=2^{-n}$ because $\gamma_{j} \neq$ huge.

CaSE 10: $\left(D_{3}, D_{5}\right):$ Suppose that $|M[m]|<n$ and $|\bar{M}[\bar{m}]|<n$. If $m>\bar{m}$ then $\operatorname{Pr}[X[m]=$ $\bar{X}[\bar{m}]]=\operatorname{Pr}[\Sigma=\bar{\Sigma}]=2^{-n}$ because of the contribution of $Y[m-1]$ in $\Sigma —$ a random variable that is not used in the definition of $\bar{\Sigma}$. If $m<\bar{m}$ then $\operatorname{Pr}[X[m]=\bar{X}[\bar{m}]]=\operatorname{Pr}[\Sigma=\bar{\Sigma}]=2^{-n}$ because of the contribution of $\bar{Y}[\bar{m}-1]$ in $\bar{\Sigma}$-a random variable that is not used in the definition of $\Sigma$. If $m=\bar{m}$ and there is an $i<m$ such that $M[i] \neq \bar{M}[i]$ then $\operatorname{Pr}[X[m]=\bar{X}[\bar{m}]]=\operatorname{Pr}[\Sigma=\bar{\Sigma}]=2^{-n}$ because of the contribution of $\bar{Y}[i]$ in $\bar{\Sigma}$-a random variable that is not used in the definition of $\Sigma$. If $m=\bar{m}$ and for every $i<m$ we have that $M[i]=\bar{M}[i]$, then, necessarily, $M[m] \neq \bar{M}[m]$. In this case $\operatorname{Pr}[\Sigma=\bar{\Sigma}]=0$, as the two checksums differ by the nonzero value $\operatorname{pad}(M[m]) \oplus \operatorname{pad}(\bar{M}[m])$.

Suppose that $|M[m]|=n$ and $|\bar{M}[\bar{m}]|=n$. Then $X[m]$ and $\bar{X}[m]$ are being offset by the same amount, huge $\cdot L$, so this offset is irrelevant in computing $\operatorname{Pr}[X[m]=\bar{X}[\bar{m}]$ and one proceeds exactly as above.

Suppose that $|M[m]|<n$ and $|\bar{M}[\bar{m}]|=n$. Then $\operatorname{Pr}[X[m]=\bar{X}[m]]=\operatorname{Pr}[\Sigma=\bar{\Sigma} \oplus$ huge. $L]=2^{-n}$ since $\Sigma$ and $\bar{\Sigma}$ are independent of $L$. Similarly, if $|M[m]|=n$ and $|\bar{M}[\bar{m}]|<n$, then $\operatorname{Pr}[X[m]=\bar{X}[m]]=2^{-n}$.

Case 11: $\left(D_{4}, D_{4}\right)$ : This case mirrors Case 5 , so it is omitted.
Case 12: $\left(D_{4}, D_{5}\right)$ : This case mirrors Case 6 , so it is omitted.
This completes the proof.

## B Test Vectors

PMAC-AES test vectors and reference code, prepared by Ted Krovetz (tdk@acm.org), are available at http://www.cs.ucdavis.edu/~rogaway/

## C Document History

- October 16, 2000. First version of the PMAC document submitted to NIST [25].
- April 1, 2001. Second version of the PMAC document. Submitted to NIST.

