

XOR MACs: New Methods for Message Authentication Using Finite Pseudorandom Functions

Mihir Bellare¹ and Roch Guérin¹ and Phillip Rogaway²

¹ IBM T.J. Watson Research Center, PO Box 704, Yorktown Heights, NY 10598, USA. e-mail: {mihir, guerin}@watson.ibm.com

² Dept. of Computer Science, Eng. II Bldg., University of California, Davis, CA 95616, USA. e-mail: rogaway@cs.ucdavis.edu

Abstract. We describe a new approach for authenticating a message using a finite pseudorandom function (PRF). Our “XOR MACs” have several nice features, including parallelizability, incrementality, and provable security. The finite PRF can be “instantiated” via DES (yielding an alternative to the CBC MAC), via the compression function of MD5 (yielding an alternative to various “keyed MD5” constructions), or in a variety of other ways. The proven security is quantitative, expressing the adversary’s inability to forge in terms of her (presumed) inability to break the underlying finite PRF. This is backed by attacks showing the analysis is tight. Our proofs exploit linear algebraic techniques.

1 Introduction

A message authentication scheme enables two parties sharing a key a to authenticate their transmissions. This is one of the most widely used cryptographic primitives, and it may become even more so: as security concerns grow, it is reasonable to anticipate that virtually every transmitted message (or packet) will use cryptographic means to ensure authenticity. (For example, the ubiquitous use of message authentication is already being contemplated for the next generation of Internet Protocols.)

Message authentication is usually accomplished by including with each transmitted message M a short string, called its “message authentication code” (MAC) or “signature,” computed as a function of M and the shared key a . The most prevalent MAC is the “cipher block chaining message authentication code” (CBC MAC) specified in the International Standard ISO 9797 [ISO] and the U.S. Standard ANSI X9.9 [X9.9]. In recent years another type of MAC has started to become prevalent: these are constructed by somehow “keying” a cryptographic hash, as in $\text{MAC}_a(x) = \text{MD5}(x.a)$ (see, for example, [Ts]).

The goal of the present work is to provide new methods which have certain efficiency and security advantages. We call our methods “XOR schemes.” They are simple to describe and implement. They use as their underlying primitive any finite pseudorandom function (PRF). In particular, a finite PRF can be defined

from a block cipher (e.g. DES) or from the compression function of a cryptographic hash (e.g., MD5) yielding concrete alternatives to the above mentioned MACs.

What is an XOR MAC? At the highest level, the computation of an XOR MAC consists of three steps: (1) encode the message M as a collection of blocks (each block will depend on a small number of bits from the message); (2) apply the finite PRF to each of the blocks, thus creating a collection of PRF images (the MAC key a is the index for all of these PRF computations); and (3) XOR the set of PRF images together, building the MAC out of the result. Different ways of implement the encoding step (and different choices of the finite PRF) yield different XOR MACs. (Obviously not all encodings will result in secure MACs. We specify several simple ones which do, and also specify general conditions to determine which encodings work.)

This paper specifies, for every finite PRF family F and every value of a block size b , two XOR MACs—a stateless (and probabilistic) one called $\text{XMACR}_{F,b}$, and a stateful (and deterministic) one called $\text{XMACC}_{F,b}$. (In a stateful MAC the signer maintains information, in our case a counter, which he updates each time a message is signed.) The schemes are described concretely in Section 2, as are their main efficiency advantages, namely parallelizability and incrementality.

Security of our schemes. Our XOR schemes are *proven* secure— we show that if the F is a “secure” finite PRF family then the MAC schemes based on it are also “secure.” In formalizing this, security of a finite PRF family means indistinguishability from a family of random functions in the sense of [GGM], while security of a MAC means it resists chosen message attack. To make these results meaningful for practice, the security in both cases is made quantitative: we measure the success probabilities as a function of the resources (time and chosen message queries) available to the adversary, and specify exact reductions, enabling the protocol designer to compute, given some assumed security on the finite PRF, how many queries an XOR MAC based on it will withstand. This type of security analysis for a MAC, starting from a finite PRF, begins with [BKR].

Our XOR schemes are so simple that it is tempting to think one can easily find attacks. This is why we stress the importance of the proofs of security which show that no attacks short of breaking the underlying PRF will succeed.

An advantage of quantified security is that it allows one to compare the securities of different MACs based on the same finite PRF family. (Note that a concrete finite PRF family F , eg. a block cipher like DES, may possess strengths which are not reflected in the model of F being a finite PRF family, and these other strengths are potentially relevant in determining the strength of a MAC based on the block cipher. In making security determinations and comparisons we are treating the underlying primitives, eg. DES, as being known to only possess the properties which have been formally modeled, here the property of being a finite PRF.) We will see that our counter based MAC is “more secure” than our randomized one, and that both are “more secure” than the CBC MAC. In particular, the success probability of the adversary in the XOR schemes is

independent of the lengths of the messages in her chosen message attack (as long as they stay below a certain specified but very large length) while the attacks of [Kr, PV] show that the success probability of the adversary in the CBC scheme grows as a linear function of the message length. See Section 6.

We also describe the best attacks we know on the XOR schemes. They use birthday attacks (collisions) and indicate that the analysis from our proofs is tight.

2 The Schemes and Their Properties — Concretely

We begin by presenting concrete instantiations of our two main schemes using DES. (But we stress this is just an example. Other instantiations are possible, using other block ciphers, or even methods such as MD5, as discussed later.) We let $l = 64$ and $L = 48$. For any 56-bit key a and l -bit plaintext x we let $F_a(x)$ be the first L bits of $\text{DES}_a(x)$. (We stress that F_a outputs only 48 bits, and not the full 64-bit DES output. We have truncated the output because DES is a pseudorandom *permutation*, while what we want is a pseudorandom function.) Sender and receiver share a 56-bit DES key a which specifies F_a .

Message formatting and notation. We assume the length $|M|$ of M is a multiple of 32 bits, which can easily be achieved by a suitable padding. (For example, append a one and then append enough zeros to bring the length to a multiple of 32 bits.) The message is then viewed as a sequence of 32-bit blocks, $M = M[1] \dots M[n]$ with $|M[i]| = 32$ for $i = 1, \dots, n$. We assume that the number n of blocks is less than 2^{31} —equivalently $|M| \leq 32 * 2^{31} = 2^{36}$ bits— which would not normally be a significant restriction in practice.

Let (i) denote the binary representation of block index $i \in \{1, \dots, n\}$ as a string of exactly 31 bits. (This is why we assumed the bound on n .) Let $\alpha.\beta$ denote the concatenation of strings α and β . We give two schemes:

Scheme XMACR. The first scheme is called the randomized XOR scheme, XMACR. To authenticate the message $M = M[1] \dots M[n]$ do the following:

- Pick at random a 63-bit string r , hereafter called the *seed*
- Set $z = F_a(0.r) \oplus F_a(1.(1).M[1]) \oplus F_a(1.(2).M[2]) \oplus \dots \oplus F_a(1.(n).M[n])$
- Set the MAC of M to the pair $\mu = (r, z)$.

Thus the sender will transmit (M, μ) . The receiver, receiving (M', μ') , where $\mu' = (r', z')$, computes $z = F_a(0.r') \oplus F_a(1.(1).M'[1]) \oplus F_a(1.(2).M'[2]) \oplus \dots \oplus F_a(1.(n).M'[n])$. The receiver accepts M' if and only if $z = z'$.

We stress that new coins are flipped to determine the seed each time the sender wants to authenticate a message, and also that the seed is included in the signature.

Scheme XMACC. The second scheme is called the counter-based XOR scheme. Here it is required that the sender maintain a 63-bit counter C which is initially 0 and is incremented for each message. (Thus at most 2^{63} messages can be signed.) To authenticate message $M = M[1] \dots M[n]$ do the following:

- Increment the counter C by 1

- Set $z = F_a(0.C) \oplus F_a(1.\langle 1 \rangle.M[1]) \oplus F_a(1.\langle 2 \rangle.M[2]) \oplus \dots \oplus F_a(1.\langle n \rangle.M[n])$
- Set the MAC of M to the pair $\mu = (C, z)$.

Thus the sender will transmit (M, μ) . The receiver, receiving (M', μ') where $\mu' = (C', z')$, computes $F_a(0.C') \oplus F_a(1.\langle 1 \rangle.M'[1]) \oplus F_a(1.\langle 2 \rangle.M'[2]) \oplus \dots \oplus F_a(1.\langle n \rangle.M'[n])$ and accepts iff this value equals z' . Note the counter is included in the signature. Also the receiver maintains no state.

Stateful schemes are not necessarily “worse” than stateless ones; programmatically, a “static” variable is easy, but a good approximation to randomness is hard. We now discuss properties of XMACR and XMACC.

Parallelizability. The DES computations on different blocks can be made in *parallel*. In general, the throughput of an XOR MAC can be doubled by doubling the amount (and not speed) of the underlying hardware. An environment where this is crucial is message authentication over high speed networks (where packets will flow over optical links at rates of 1–10 GBit/second). In that setting one cannot realistically use the CBC MAC because of its sequential nature; an XOR scheme is a more appropriate choice. Note that even in the software setting parallelizability can be relevant: with an adequate degree of parallelism, multiple machine pipelines can all be kept busy doing useful work.

Incrementality. An XOR MAC is *incremental* [BGG1] with respect to block replacement. Suppose $M[i]$ is modified to a new 32-bit value m . Then, for a long message M , one can update the MAC much quicker than it would take to re-compute it. Let's illustrate for XMACR. Let $\mu = (r, z)$ be a MAC of M and let M' denote M with block i replaced by m . To compute a MAC for M' , pick r' at random and let $z' = z \oplus F_a(0.r) \oplus F_a(0.r') \oplus F_a(1.\langle i \rangle.M[i]) \oplus F_a(1.\langle i \rangle.m)$. Then (r', z') is a MAC for M' . Extensions of this scheme to support insertion and deletion of blocks (not just replacement) appear in [BGG2].

Out-of-order verification. Tag verification can proceed even if message blocks arrive out of order. Here it is only necessary that the each block be accompanied by its index. With other mechanisms MAC verification cannot proceed before the first block has been received, for example. Out-of-order MAC verification is useful since networks always have some degree of packet loss and re-transmission.

DES computations. The number of DES computations is twice that of the CBC MAC. (The overhead can be reduced as discussed in Section 4 by increasing the block size, currently set to 32, at the cost of reducing the maximum allowable message length.) So, in software, the above schemes are slower than the CBC MAC. But an XOR MAC based on DES is interesting for hardware efficiency, particularly for high-speed networks, or in settings where the incrementality compensates for the slower from-scratch MACing time. For a software-efficient XOR MAC use the MD5-based instantiation discussed later.

MD5-based instantiation. A software-efficient XOR MAC would start not with DES but with a software-efficient PRF. For example, from the compression function of a cryptographic hash function, say $md : \{0, 1\}^{640} \rightarrow \{0, 1\}^{128}$, one can define a finite PRF, say $F_a(x)$ equals the first 64 bits of $md(x.a)$, where

$|x| = 560$ and $|a| = 80$. Using 48-bits for the block index, we would get a MAC which uses one application of md for every 512 bits of message. This is as efficient as proposals like MD5($x.a$) or MD5($a.x.a$) which are currently being considered for the Internet, and has the advantages of parallelizability and incrementality.

Security. Observe that including the block indices in the argument to F_a is necessary— if they are omitted, permuting the message blocks would leave the MAC unchanged. One can also see that the block containing the random string r (resp. counter) of XMACR (resp. XMACC) cannot be omitted. In other words, the scheme in which the MAC is set to $F_a(1.(1).M[1]) \oplus F_a(1.(2).M[2]) \oplus \dots \oplus F_a(1.(n).M[n])$ is easily broken—e.g., set $M_1 = A.B$, $M_2 = A'.B$, $M_3 = A.B'$ and $M_4 = A'.B'$, and note that the MACs of M_1, M_2, M_3 sum to give the MAC of M_4 .

The idea behind the nonces is to prevent the attacker from forming new MACs via linear combinations of old ones. This is in fact the only attack short of breaking the PRF. This is not obvious, of course; indeed it is far from clear why XMACR and XMACC should be secure. That is why we have our proofs.

3 Definitions

In order to prove the security we first need to provide appropriate definitions of security for block ciphers and message authentication schemes.

We model block ciphers as finite pseudorandom functions in the manner of [BKR]. The underlying notion is the pseudorandom function notion of [GGM], appropriately tailored to take into account the fact that block ciphers have fixed input and output lengths that can't be treated asymptotically. This approach builds on a suggestion of [LuRa] that DES be viewed as a "pseudorandom in practice" function family.

Denote by $|x|$ the length of a string x . If $i \in \{1, \dots, 2^n\}$ is an integer then we denote by $\langle i \rangle_n$ the natural binary encoding of i as an n -bit string. (Thus the $\langle \cdot \rangle$ of Section 2 is $\langle \cdot \rangle_{31}$ in our current notation.) If S is a set (resp. probability space) then $x \stackrel{r}{\leftarrow} S$ denotes the operation of selecting an element uniformly at random from S (resp. at random according to the distribution specified by S).

Finite pseudorandom function families. A *function family* is a set of functions, and an associated set of strings called keys. Each key names a function in the family according to some fixed convention, and the function corresponding to key a is denoted F_a . (Note that two keys can name the same function.) To pick a function f at random from a family F means to pick a key a uniformly at random and let $f = F_a$; we write $f \stackrel{r}{\leftarrow} F$ for this operation. For example DES is a function family where the set of keys is the set of all 56-bit strings.

A family F has input length l and output length L if each $f \in F$ maps $\{0, 1\}^l$ to $\{0, 1\}^L$. (Eg. $l = L = 64$ for DES.) It has key length κ if the associated set of keys is the set of all strings of length κ . The family of random functions with input length l and output length L is the family R of all functions mapping $\{0, 1\}^l$ to $\{0, 1\}^L$. The key of a function f in this family is the string which describes its truth table. Note this is a very large family, consisting of 2^{L2^l} functions.

A finite function family F is “pseudorandom” if the input-output behavior of F_a “looks random” to someone who doesn’t know the key a . This is formalized via the notion of statistical tests [GGM]. Formally, such a test is an oracle algorithm A . Let F, G be finite function families. The advantage of A in distinguishing F from G is defined by

$$\text{Adv}_A(F, G) = \Pr_{g \leftarrow F} [A^g = 1] - \Pr_{g \leftarrow G} [A^g = 1].$$

The probability is over the indicated random choice of g and the coin tosses of A .

Let family F have input length l and output length L , and let R be the family of random functions with the same parameters. To discuss security quantitatively, we say that statistical test A $[t, q, \epsilon]$ -breaks F if A runs in at most t steps, makes at most q oracle queries, and achieves $\text{Adv}_A(F, R) \geq \epsilon$. (The running time here is measured in a standard RAM model of computation.) In informal discussion, the finite function family F is said to be $[t, q, \epsilon]$ -pseudorandom if there is no statistical test that $[t, q, \epsilon]$ -breaks F . (To be fully formal one ought to consider also other parameters such as the “code size”.) In other words, in time t and given q examples one cannot distinguish a random member of F from a random function with advantage more than ϵ .

Notice that the key size of the finite PRF family F does not need to be explicitly specified in the definition of security: its influence is captured in that it influences the values of t, q, ϵ for which the F is $[t, q, \epsilon]$ -pseudorandom.

Message authentication. We provide formal definitions of schemes and their security in the exact security setting. We begin with stateless schemes, in which no counters or other state information need be maintained. Then we briefly indicate how the definitions should be updated to take account of state.

A (*stateless*) message authentication scheme consists of a *signing* algorithm Sig and a *verifying* algorithm Vf . The signing algorithm may be probabilistic; the verifying one typically is not. Associated to the scheme are parameters κ and L_{sig} describing the key length and signature length, respectively. On input a κ -bit key a and a message M , algorithm Sig outputs an L_{sig} -bit string μ called the *signature*, or MAC, of M . On input a κ -bit key a , a message M and an L_{sig} -bit string μ , algorithm Vf outputs a bit, with 1 standing for accept and 0 for reject. We ask for a basic validity condition, namely that authentic signatures are accepted with probability one.

An adversary for a message authentication scheme is allowed a chosen message attack, and declared successful if, following this attack, she produces a forgery. Formally the adversary is a probabilistic algorithm E which is given oracle access to the signer and verifier—more precisely, to $\text{Sig}(a, \cdot)$ and $\text{Vf}(a, \cdot, \cdot)$ for a random but hidden choice of a . E can request a signature of a message of her choice; to do this, she writes M on a special query tape. She can also ask the verifier to verify that μ is a valid signature for M ; to do this she writes (M, μ) on a special query tape. E ’s *attack* on the scheme is described by the following experiment:

- (1) A random string a of length κ is selected as the shared secret. A random string r_E is selected as the coin tosses of E . E now starts computing.
- (2) Suppose E makes a signing query M . Then the oracle computes a signature $\mu \stackrel{R}{\leftarrow} \text{Sig}(a, M)$ and returns it to E . (Since Sig may be probabilistic, this step requires making the necessary underlying choice of a random string for Sig , anew for each signing query.)
- (3) Suppose E makes a verify query (M, μ) . The oracle computes the decision $d = \text{Vf}(a, M, \mu)$ and returns it to E .

The adversary is allowed an adaptive chosen message attack, as in the notion of [GMR], but we also allow verify queries because, unlike the setting in digital signatures, E cannot compute the verify predicate on her own (since the verify algorithm is not public). Note that E does not see a nor the coin tosses of Sig .

We say that E 's attack on \mathcal{M} is a (q_s, q_v) -attack if during the course of the attack she makes no more than q_s signing queries and no more than q_v verify queries. A (q_s, q_v) -attack is a (t, q_s, q_v) -attack if, in addition, E runs for no more than t steps, in the RAM model of computation we fixed above. It is useful to say that a verify query (M, μ) is *known-authentic* if a signing query M was made prior to this verify query and the signature returned was μ . Note validity implies that known-authentic verify queries are accepted. We thus assume of any adversary E that she never makes any known-authentic queries.

The outcome of running the protocol in the presence of an adversary is used to define security. We say that E is *successful* if she makes a verify query (M, μ) which is accepted but which is not known-authentic.³ (The verify query (M, μ) in question is called a *forgery*, and the definition reflects the notion of existential forgery [GMR].) We say that E $[q_s, q_v, \epsilon]$ -breaks the scheme if her attack is a (q_s, q_v, ϵ) -attack and her probability of success is at least ϵ . We say she $[t, q_s, q_v, \epsilon]$ -breaks the scheme if her attack is a (t, q_s, q_v, ϵ) -attack and her probability of success is at least ϵ . In informal discussion we'll say the scheme is $[t, q_s, q_v, \epsilon]$ -unforgeable if there is no adversary who can $[t, q_s, q_v, \epsilon]$ -break it. (To be fully formal we would have to consider also other parameters like the "code size.")

In a stateful message authentication scheme the signer maintains state across consecutive signing requests. (For example, in our counter-based scheme the signer maintains a message counter.) In such a case the signing algorithm can be thought of as taking an additional input —the "current" state C_s of the signer— and returning an additional output —the signer's next state. We must modify the experiment describing E 's attack: in Step (1) we also have that C_s is initialized to a value specified by the scheme; and in Step (2) we compute $(\mu, C'_s) \stackrel{R}{\leftarrow} \text{Sig}(a, M, C_s)$, then return μ to the adversary and replace C_s by C'_s . Note the adversary doesn't see the revised state (though in the stateful scheme

³ This is slightly stronger than the more standard definition in which one would only ask that the message M was not a previous signing query. We make this stronger requirement because we achieve it and because it is useful in contexts like entity authentication.

of this paper this wouldn't matter). Also note that we allow the signer a state, but not the verifier.

4 The Randomized XOR Scheme

We first present the general scheme, of which the scheme XMACR in Section 2 is a special case, and then proceed to the security analysis.

SPECIFICATION. Let F be a family of functions with key length κ , input length l , and output length L . We fix in addition a parameter $b \leq l - 1$ which will be the block size. We will assume that any message M to be authenticated has length at most $|M| \leq b2^{l-b-1}$. By standard padding arguments we may assume wlog that the message length is a multiple of b . We then regard M as a sequence of b -bit blocks. The number of blocks is denoted $\|M\|_b$, and with b understood the i -th block is denoted $M[i]$, for $i = 1, \dots, \|M\|_b$. Let $r \in \{0, 1\}^{l-1}$, and let $a \in \{0, 1\}^\kappa$ be the shared key. We define $\text{tag}_{F,b}(a, M, r)$ by

$$F_a(0.r) \oplus F_a(1.\langle 1 \rangle_{l-b-1}.M[1]) \oplus \dots \oplus F_a(1.\langle n \rangle_{l-b-1}.M[n]). \quad (1)$$

We use this function in both the randomized and the counter-based schemes. We call r the *seed*. The (stateless) message authentication scheme is:

<p>function $\text{SigR}_{F,b}(a, M)$</p> <p>$r \xleftarrow{\\$} \{0, 1\}^{l-1};$</p> <p>$z \leftarrow \text{tag}_{F,b}(a, M, r)$</p> <p>return (r, z)</p>	<p>function $\text{VfR}_{F,b}(a, M', (r', z'))$</p> <p>$z \leftarrow \text{tag}_{F,b}(a, M', r')$</p> <p>if $z = z'$ then return 1</p> <p>else return 0</p>
--	--

We call $\text{XMACR}_{F,b}$ the randomized XOR scheme based on function family F and using block size b . The validity condition is easy to verify. Note that the XMACR scheme of Section 2 is, in the current terminology, $\text{XMACR}_{F,32}$ with F being the family specified by $F_a(\cdot) = \text{first 48 bits of the output of } \text{DES}_a(\cdot)$.

TRADING EFFICIENCY FOR MESSAGE LENGTH. Note that choosing different values of b will tradeoff the number of F_a computations with the allowable length of messages that can be signed. Namely, the scheme calls for $1 + \|M\|_b = 1 + (|M|/b)$ evaluations of F_a and allows $|M|$ to be $b2^{l-b-1}$ so that increasing b reduces the number of F_a evaluations at the cost of restricting the scheme to shorter messages. For example, the XMACR scheme of Section 2, with $b = 32$, currently has twice the DES operations of the CBC MAC, and allows $|M|$ up to 2^{36} . But we could set $b = 48$ and have only 33% more DES operations than the CBC MAC, now with $|M| \leq 48 * 2^{15} = 3 * 2^{19}$.

SECURITY: INFORMATION THEORETIC CASE. Begin by thinking of F as ideal (i.e., truly random). Namely, we consider $\text{XMACR}_{R,b}$, which we call the information theoretic case. The following theorem provides an absolute bound on the success of the adversary in terms of the number of sign and verify queries she makes.

Theorem 1. *Let R be the family of random functions with input length l and output length L , let b be at most $l - 1$, and let E be any adversary making a (q_s, q_v) -attack on $\text{XMACR}_{R,b}$. Then the probability that E is successful is at most $\delta_R \stackrel{\text{def}}{=} 2q_s^2 \cdot 2^{-l} + q_v \cdot 2^{-L}$.*

Proof. Due to page limits we provide a very brief sketch; for a full proof see [BGR]. The proof has two parts: first we relate the security of the scheme to the probability that a certain matrix has full rank; then we bound this probability.

Since E is computationally unbounded she is wlog deterministic. The probabilistic choices in E 's attack on the scheme are thus the initial choice a of key (naming a random function $R_a \in R$), and the choices of seeds made by the signer in the course of signing. Let M_i denote the random variable whose value is the i -th message whose signature E requests. Let R_i be the random seed chosen by the signer to sign M_i and let $Z_i = \text{tag}_{R,b}(a, M_i, R_i)$ denote its tag, $i = 1, \dots, q_s$. Let Distinct be the event that R_1, \dots, R_{q_s} are all distinct and Succ the event that E is successful. By birthday bounds we can show that $\Pr[\neg \text{Distinct}] \leq q_s^2 \cdot 2^{l-1}$. Now we want to show that $\Pr[\text{Succ} \mid \text{Distinct}] \leq q_v \cdot 2^{-L}$ whence the theorem follows.

Fix a particular sequence of messages M_1, \dots, M_{q_s} , a particular choice $r_1, \dots, r_{q_s} \in \{0, 1\}^{l-1}$ of *distinct* seeds and a particular choice z_1, \dots, z_{q_s} of L -bit strings, for which $\Pr[M_i = M_i \text{ and } R_i = r_i \text{ and } Z_i = z_i \text{ for } i = 1, \dots, q_s] > 0$. We let

$$\Pr_1[\cdot] = \Pr[\cdot \mid M_i = M_i \text{ and } R_i = r_i \text{ and } Z_i = z_i \text{ for } i = 1, \dots, q_s]$$

denote the indicated conditional probability measure. (The probability is effectively over only the random choice of the shared key a , since everything else is fixed.) Fix a message M_{q_s+1} distinct from M_1, \dots, M_{q_s} , a seed $r_{q_s+1} \in \{0, 1\}^{l-1}$ and an L -bit string z_{q_s+1} . These are intended to stand for a possible forgery $(M_{q_s+1}, (r_{q_s+1}, z_{q_s+1}))$. (Notice that although M_{q_s+1} is distinct from previous messages, r_{q_s+1} is not assumed distinct from previous seeds—indeed, since the adversary may choose it, we cannot make such an assumption.) Below we will show that

$$\Pr_1[\text{tag}_{R,b}(a, M_{q_s+1}, r_{q_s+1}) = z_{q_s+1}] \leq 2^{-L}. \quad (2)$$

Given this, standard conditioning arguments can then be used to show that $\Pr[\text{Succ} \mid \text{Distinct}] \leq q_v \cdot 2^{-L}$. In what follows, we make the wlog assumption that E first makes its q_s signing queries, and then makes its q_v verify queries.

Recall $M_i[j] \in \{0, 1\}^b$ denotes the j -th block of M_i . We define a $q_s + 1$ by 2^l matrix B over $\text{GF}(2)$. Its rows are indexed $1, \dots, q_s + 1$ and its columns are indexed by the l -bit strings in lexicographic order. The entry in row i , column x is denoted $B[i, x]$, and is defined as follows. First consider the case where the first bit of x is 0, so that $x = 0.y$. Then we set $B[i, x] = 1$ if $y = r_i$ and 0 otherwise. Now suppose the first bit of x is 1, and write it as $x = 1.\langle j \rangle_{l-b-1}.y$, where $|y| = b$. Then we set $B[i, x] = 1$ if $M_i[j] = y$ and 0 otherwise. (In particular, $B[i, x] = 0$ if $j > \|M_i\|_b$.) Note the matrix is not a random variable—it is fixed given that M_1, \dots, M_{q_s+1} and r_1, \dots, r_{q_s+1} are fixed.

LEMMA. The matrix B has full rank.

PROOF. Transform B by row and column operations until it has a $q_s + 1$ by $q_s + 1$ identity matrix in its upper left corner. At any time, the left half of B means the first 2^{l-1} columns and the right half means the rest. Initially the left (resp. right) half consists of those columns whose index has first bit 0 (resp. 1). Since r_1, \dots, r_{q_s} are distinct, we can permute columns until the first q_s rows of the left half of B consist of a q_s by q_s identity matrix followed by a q_s by $2^{l-1} - q_s$ matrix of zeroes. We now consider two cases. First if r_{q_s+1} is distinct from r_1, \dots, r_{q_s} , then a single column swap suffices. The case when $r_{q_s+1} = r_\alpha$ for some $\alpha \in \{1, \dots, q_s\}$ uses that M is by assumption different from M_α and is more complex, requiring a few operations. \square

Equation (2) is established via standard relations of linear to probabilistic independence that have been used in several places (eg. [ABI, BeRo]). \square

Note the bound is independent of b : the latter figures only in our assumption that any query M made by E above satisfies $\|M\|_b \leq 2^{l-b-1}$. We stress that δ_R grows with the *square* of the number of signing queries: a “birthday” type behavior. Attacks we present later will show that this analysis and behavior is essentially the “best possible.”

SECURITY: COMPUTATIONAL CASE. We now assume we are given a family F which is not truly random, but $[t', q', \epsilon']$ -pseudorandom. In that case, how secure is $\text{XMACR}_{F,b}$? This is what the following tells us. It is the result of more direct interest in practice (although Theorem 1 is in some ways more basic). The constant c below depends only on details of the computational model. The proof is not too hard and can be found in [BGR].

Theorem 2. *There is an oracle machine U and a constant c such that the following is true. Let F be a family of functions with input length l and output length L and let b be at most $l - 1$. Let E be an adversary who $[t, q_s, q_v, \epsilon]$ -breaks $\text{XMACR}_{F,b}$ and suppose any message M in a query of E has a number $\|M\|_b$ of blocks which is bounded by n . Let $\delta_R = 2q_s^2 \cdot 2^{-l} + q_v \cdot 2^{-L}$. Then U^E $[t', q', \epsilon']$ -breaks F , where*

$$t' = t + c(l + L)q' \quad ; \quad q' = (q_s + q_v) \cdot (n + 1) \quad ; \quad \epsilon' = \epsilon - \delta_R .$$

In other words if F is $[t', q', \epsilon']$ -pseudorandom (the values t', q', ϵ' depending on the key size and cryptanalytic strength of the finite PRF F) then $\text{XMACR}_{F,b}$ is $[t, q_s, q_v, \epsilon]$ -unforgeable, where $t = t' - c(l + L)q'$, $q_s + q_v = q'/(n + 1)$ and $\epsilon = \epsilon' + \delta_R$. Thus a success probability of δ_R for the adversary is unavoidable, even if the PRF is “ideal;” beyond that, the success of the adversary is bounded in terms of the parameters of the block cipher.

ATTACKS. We present the best attacks we know. Since we think of F as pseudorandom, we will do the attack assuming it is in fact random; that is, we look at $\text{XMACR}_{R,b}$ where R is the family of random functions with input length l and output length L . Given q_s, q_v we specify a particular adversary E who makes q_s sign queries and q_v verify queries, and then outputs a forgery with probability

$\epsilon = \Omega(\delta_R)$, where $\delta_R = 2q_s^2 \cdot 2^{-l} + q_v \cdot 2^{-L}$. The attack is based on birthday attacks, and finds enough collisions that linearity can be exploited.

Proposition 3. *Let R be the family of random functions with input length l and output length L , and let b be at most $l - 1$. Then there is a constant $c > 0$ such that for any q_s, q_v satisfying $q_s^2 \leq 2^l$ and $q_v \leq 2^L$, there is an adversary E who $[t, q_s, q_v, \epsilon]$ -breaks $\text{XMACR}_{R,b}$, where*

$$t = c(l + L)(q_s + q_v) \quad ; \quad \epsilon = \max \left\{ \left(1 - \frac{1}{e} \right) \cdot \frac{q_s^2 - 3q_s}{4 \cdot 2^l}, \frac{q_v}{2^L} \right\}.$$

Proof. We provide a very brief sketch; for a full proof see [BGR]. Given distinct b -bit strings A', B' we show how to forge the signature of the message $M_4 = A' \cdot B'$. E chooses a b -bit string $A \notin \{A', B'\}$ and a b bit string $B \notin \{A', B', A\}$. She sets $M_1 = A \cdot B$, $M_2 = A' \cdot B$ and $M_3 = A \cdot B'$. She sets $q = \lfloor (q_s - 1)/2 \rfloor$. Now she mounts the following attack—

- (1) For $i = 1, 2$, she makes the signing query M_i a total of q times. Let $(r_{i,j}, z_{i,j})$ denote the answers, $i = 1, 2$ and $j = 1, \dots, q$.
- (2) She makes the signing query M_3 . Let (r, z_3) denote the answer.

Notice that the total number of signing queries made is $2q + 1 \leq q_s$. Now let Coll be the event that there exist j_1, j_2 such that $r_{1,j_1} = r_{2,j_2}$. Then:

- (3) If Coll is true then E sets $\mu = (r, z)$ where $z = z_{1,j_1} \oplus z_{2,j_2} \oplus z_3$. She then makes the verify query (M_4, μ) .
- (4) Else, she picks a random $r' \in \{0, 1\}^{l-1}$ and lets z'_1, \dots, z'_{q_v} be distinct L -bit strings, for example the q_v lexicographically least L -bit strings. She makes the q_v verify queries $(M_4, (r', z'_j))$ for $j = 1, \dots, q_v$.

Note the number of verify queries made is at most q_v . For the analysis first check that if Coll is true then E is successful in forgery; otherwise via step (4) she is successful with probability $q_v \cdot 2^{-L}$. To conclude we derive some birthday lower bounds to show that $\Pr[\text{Coll}]$ is at least $(1 - e^{-1}) \cdot (q_s^2 - 3q_s)/(4 \cdot 2^l)$. \square

The above indicates our analysis in Theorem 1 is tight up to small constant factors. Thus we have been able to give a pretty complete picture of the security. Namely, because of Theorem 1, the attack in the proof of Proposition 3 represents the best possible one (up to constant factors) short of breaking the PRF. We remark that the proof shows that the adversary forges the signature of essentially any message of her choice. This makes the attack all the more relevant.

5 The Counter-Based XOR Scheme

Here we present another scheme, one which allows the signer to maintain state in the form of a counter. The gain is greater security: the success probability of the adversary in the analogue of Theorem 1 does not depend on the number q_s of signing queries at all (as long as the latter is bounded by a certain exponential function of l)!

SPECIFICATION. Let F, l, L, κ be as before, and $a \in \{0, 1\}^\kappa$ the shared key. The idea is to use the same tagging function as above, but use the current counter value as the seed. Formally the scheme $\text{XMACC}_{F,b}$ is specified by functions $\text{SigC}_{F,b}, \text{VfC}_{F,b}$. The signing function depends on a counter C maintained by the signer; it is initially 0 and then incremented by the signing function itself. (In Section 2 we loosely identified C with its 63 bit representation. Now we are more precise, viewing it as an integer and writing $\langle C \rangle_{l-1}$ for the corresponding string.) The verifying function has no state. Below $\text{tag}_{F,b}$ is the function specified in Equation (1) of Section 4.

function $\text{SigC}_{F,b}(a, M, C)$ $C \leftarrow C + 1;$ $z \leftarrow \text{tag}_{F,b}(a, M, \langle C \rangle_{l-1})$ return $((C, z), C)$	function $\text{VfC}_{F,b}(a, M', \langle C' \rangle_{l-1}, z')$ $z \leftarrow \text{tag}_{F,b}(a, M', \langle C' \rangle_{l-1})$ if $z = z'$ then return 1 else return 0
--	--

We call $\text{XMACC}_{F,b}$ the counter-based XOR scheme based on function family F and using block size b . The validity of the counter-based XOR scheme is easy to verify. Note that the XMACC scheme of Section 2 is, in the current terminology, $\text{XMACC}_{F,32}$ with F being the family specified by $F_a(\cdot) = \text{first 48 bits of the output of } \text{DES}_a(\cdot)$.

As before the length of any message whose signature the adversary requests is assumed bounded by $b2^{l-b-1}$. But also we will now assume that the total number of signing requests is bounded by 2^{l-1} . That is, we require that C not exceed 2^{l-1} . (Typically, this is not a significant restriction.) These assumptions are made in the theorems that follow.

SECURITY: INFORMATION THEORETIC CASE. There is a dramatic increase in security: the success probability of the adversary now depends only on the number q_v of its verify queries, rather than this plus $2q_s^2 \cdot 2^{-l}$. The proof is like that of Theorem 1 and can be found in [BGR].

Theorem 4. *Let R be the family of random functions with input length l and output length L , let b be at most $l - 1$, and let E be any adversary making a (q_s, q_v) -attack on $\text{XMACC}_{R,b}$, where $q_s < 2^{l-1}$. Then the probability that E is successful is at most $\delta_C \stackrel{\text{def}}{=} q_v \cdot 2^{-L}$.*

To see concretely what this improvement means, think of $F_a = \text{first 48 bits of } \text{DES}_a$, where we have $l = 64$ and $L = 48$. If $q_s = 2^{20}$ and $q_v = 1$, then the success probability is a marginal 2^{-23} in the randomized scheme, but it is 2^{-48} in the counter-based one.

SECURITY: COMPUTATIONAL CASE. We get a corresponding improvement. The proof is like that of Theorem 2 and can be found in [BGR].

Theorem 5. *There is an oracle machine U and a constant c such that the following is true. Let F be a family of functions with input length l and output length L and let b be at most $l - 1$. For $q_s < 2^{l-1}$ let E be an adversary who*

$[t, q_s, q_v, \epsilon]$ -breaks $\text{XMACC}_{F,b}$, and suppose any message M in a query of E has a number $\|M\|_b$ of blocks which is bounded by n . Let $\delta_C = q_v \cdot 2^{-L}$. Then U^E $[t', q', \epsilon']$ -breaks F , where

$$t' = t + c(l + L)q' \quad ; \quad q' = (q_s + q_v) \cdot (n + 1) \quad ; \quad \epsilon' = \epsilon - \delta_C .$$

Again, what this means is that if F is $[t, q', \epsilon']$ -pseudorandom then $\text{XMACC}_{F,b}$ is $[t, q_s, q_v, \epsilon]$ -unforgeable, where $t = t' - c(l + L)q'$, $q_s + q_v = q'/(n + 1)$, and, most importantly, $\epsilon = \epsilon' + \delta_C$.

ATTACKS. The best attack is just to guess signatures! Furthermore one does not expect better than the following (short of breaking the PRF) by virtue of the above theorems.

Proposition 6. *Let R be the family of random functions with input length l and output length L , and let b be at most $l - 1$. There is a constant $c > 0$ such that for any $q_v \leq 2^L$, there is an adversary E who $[t, 0, q_v, \epsilon]$ -breaks $\text{XMACC}_{R,b}$, where $t = c(l + L)q_v$ and $\epsilon = q_v \cdot 2^{-L}$.*

6 Comparison with the CBC MAC

We compare the security of our schemes to that of the CBC MAC. First, let us recall that scheme. Let F be a family of functions with input and output length l . A message $M = M[1] \dots M[n]$ is viewed as a sequence of l -bit blocks. The (full) CBC scheme is specified by the following:

<p>function $\text{SigCBC}_{F,n}(a, M[1] \dots M[n])$</p> <p style="padding-left: 20px;">$y_0 \leftarrow 0^l$</p> <p style="padding-left: 20px;">for $i \leftarrow 1$ to n do $y_i \leftarrow F_a(y_{i-1} \oplus M[i])$</p> <p style="padding-left: 20px;">return y_n</p>	<p>function $\text{VfCBC}_{F,n}(a, M', \mu')$</p> <p style="padding-left: 20px;">$\mu \leftarrow \text{SigCBC}_{F,n}(a, M')$</p> <p style="padding-left: 20px;">if $\mu = \mu'$ then return 1</p> <p style="padding-left: 20px;">else return 0</p>
---	---

The scheme is denoted $\text{CBC-MAC}_{F,n}$. We will consider the information theoretic case. The following was proved in [BKR]. Let R be the family of random functions of input and output length l , and let E be any adversary. Then the probability that E $[q_s, q_v, \epsilon]$ -breaks $\text{CBC-MAC}_{R,n}$ is at most $\delta_{\text{CBC}} \stackrel{\text{def}}{=} 3(n^2 + 1) \cdot (q_s + q_v)^2 \cdot 2^{-l}$. To compare this to our schemes set $L = l$ in Theorems 1 and 4. Clearly, δ_R is smaller than δ_{CBC} , and δ_C is considerably smaller than δ_{CBC} ; in particular, δ_R and δ_C don't depend on n while δ_{CBC} does, a significant difference.

Yet this by itself is not proof that our schemes are more secure, because it may be that the analysis of [BKR] is not tight. In fact, however, there are attacks (lower bounds) which indicate that the best improvement one could hope for in their analysis would be that $\delta_{\text{CBC}} = \Omega(nq_s^2 + q_v)2^{-l}$. This result is due independently to Krawczyk [Kr] and Preneel and Van Oorschot [PV]—what they show is an attack on the CBC MAC which succeeds in forging the signature of a new message with probability $\Omega(nq_s^2) \cdot 2^{-l}$, after having made q_s signing queries on n -block messages. Thus the dependence on n in δ_{CBC} is unavoidable.

We comment that the CBC-MAC $_{F,n}$ is only secure for fixed n ; the scheme must be modified to accommodate n 's of varying length. In contrast, both XMACR $_{F,b}$ and XMACC $_{F,b}$ operate on inputs of varying lengths (with the security bounds given by our theorems).

Acknowledgments

We thank Mike Luby for pointing out an oversight in a previous version of the proof of Theorem 1. We thank Hugo Krawczyk for a great deal of useful advice and information.

References

- [ABJ] N. ALON, L. BABAI AND A. ITAI. A fast and simple randomized parallel algorithm for the maximal independent set problem. *J. of Algorithms*, Vol.7, 567-583, 1986.
- [BGR] M. BELLARE, R. GUÉRIN AND P. ROGAWAY. XOR MACs: New methods for message authentication using finite pseudorandom functions. Available from the authors or out of <http://www.cs.ucdavis.edu/~rogaway/>
- [BKR] M. BELLARE, J. KILIAN AND P. ROGAWAY. On the security of cipher block chaining. *Advances in Cryptology - Crypto 94 Proceedings*.
- [BGG1] M. BELLARE, O. GOLDBREICH AND S. GOLDWASSER. Incremental cryptography: The case of hashing and signing. *Advances in Cryptology - Crypto 94 Proceedings*.
- [BGG2] M. BELLARE, O. GOLDBREICH AND S. GOLDWASSER. Incremental cryptography and application to virus protection. *Proceedings of the Twenty Seventh Annual Symposium on the Theory of Computing*, ACM, 1995.
- [BeRo] B. BERGER AND J. ROMPEL, "Simulating $(\log^c n)$ -wise independence in NC," *Proceedings of the Thirtieth Annual Symposium on the Foundations of Computer Science*, IEEE, 1989.
- [GGM] O. GOLDBREICH, S. GOLDWASSER, AND S. MICALI. How to construct random functions. *Journal of the ACM*, Vol. 33, No. 4, 210-217, 1986.
- [GMR] S. GOLDWASSER, S. MICALI, AND R. RIVEST. A digital signature scheme secure against adaptive chosen-message attacks. *SIAM Journal of Computing*, 17(2):281-308, April 1988.
- [ISO] ISO/IEC 9797. Data cryptographic techniques - Data integrity mechanism using a cryptographic check function employing a block cipher algorithm, 1989.
- [Kr] H. KRAWCZYK. Personal communication, September 1994.
- [LuRa] M. LUBY AND C. RACKOFF, "How to construct pseudorandom permutations from pseudorandom functions," *SIAM J. Comput*, Vol. 17, No. 2, April 1988.
- [PV] B. PRENEEL AND P. VAN OORSCHOT. A new generic attack on message authentication codes. *Advances in Cryptology - Crypto 95 Proceedings*.
- [Ri] R. RIVEST, "The MD5 message digest algorithm." IETF RFC-1321, 1992.
- [Ts] G. TSUDIK, "Message authentication with one-way hash functions." *Proceedings of Infocom 92*, IEEE Press, 1992.
- [X9.9] ANSI X9.9, American National Standard for Financial Institution Message Authentication (Wholesale), American Bankers Association, 1981. Revised 1986.