

# On the Impossibility of Highly-Efficient Blockcipher-Based Hash Functions

John Black<sup>1</sup>, Martin Cochran<sup>1</sup>, and Thomas Shrimpton<sup>2</sup>

<sup>1</sup> Dept. of Computer Science, University of Colorado, Boulder CO 80309, USA  
jrblack@cs.colorado.edu, Martin.Cochran@colorado.edu

www.cs.colorado.edu/~jrblack, ucsu.colorado.edu/~cochranm

<sup>2</sup> Dept. of Computer Science, Portland State University, Portland OR, 97207, USA  
teshrim@cs.pdx.edu  
www.cs.pdx.edu/~teshrim

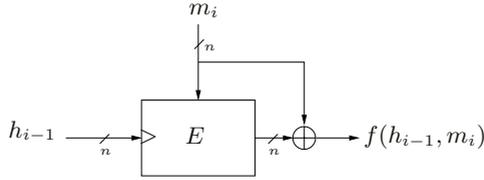
**Abstract.** Fix a small, non-empty set of blockcipher keys  $\mathcal{K}$ . We say a blockcipher-based hash function is *highly-efficient* if it makes exactly one blockcipher call for each message block hashed, and all blockcipher calls use a key from  $\mathcal{K}$ . Although a few highly-efficient constructions have been proposed, no one has been able to prove their security. In this paper we prove, in the ideal-cipher model, that it is *impossible* to construct a highly-efficient iterated blockcipher-based hash function that is provably secure. Our result implies, in particular, that the Tweakable Chain Hash (TCH) construction suggested by Liskov, Rivest, and Wagner [7] is *not* correct under an instantiation suggested for this construction, nor can TCH be correctly instantiated by any other efficient means.

**Keywords:** Collision-resistant hash functions, tweakable blockciphers, provable security.

## 1 Introduction

BACKGROUND. Essentially all modern hash functions are built by iterating a compression function according to the Merkle-Damgård paradigm [4, 10]. Moreover, these compression functions are almost always built from a blockcipher. Constructions like the Matyas-Meyer-Oseas (MMO) compression function [8] are explicit about their use of a blockcipher, but even so-called “dedicated” hashing primitives like MD5 and SHA-1 are in fact blockcipher-based. The SHA-1 compression function, for example, uses a 160-bit blockcipher that takes a 512-bit key; this blockcipher has been named SHACAL-1 [6].

This idea of building hash functions from blockciphers goes back more than 25 years. The earliest construction by Rabin [12] proposed to hash a message  $M = m_1 m_2 \cdots m_\ell$  by fixing an initial value  $h_0$  and computing  $H(M) = \text{DES}_{m_\ell}(\text{DES}_{m_{\ell-1}}(\cdots(\text{DES}_{m_1}(h_0))))$ ; effectively, this is a Merkle-Damgård construction with blockcipher-based compression function  $f(h_{i-1}, m_i) = E_{m_i}(h_{i-1})$ . Other constructions like Davies-Meyers [9] and MMO followed, but it was Preneel, Govaerts, and Vandewalle [11] who conducted the first systematic study



**Fig. 1.** The Matyas-Meyer-Oseas (MMO) compression function [8].  $E: \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}^n$  is a block cipher; the hatch mark denotes the location of the key. Iterating this compression function results in a provably-secure hash function [2], however notice that the above compression function will be rekeyed each round

of blockcipher-based hash functions. They considered the security of 64 iterated constructions with compression functions of the form  $f(h_{i-1}, m_i) = E_a(b) \oplus c$  where  $a, b, c \in \{h_{i-1}, m_i, h_{i-1} \oplus m_i, v\}$  for some fixed constant  $v$ . The analysis of PGV was attack-based, and schemes not broken by their attacks were deemed secure. Subsequently, Black, Rogaway and Shrimpton [2] considered these same 64 constructions using a proof-based approach. They showed that, in the ideal-cipher model, 20 of the 64 schemes are collision-resistant up to the birthday bound; Figure 1 gives one example.

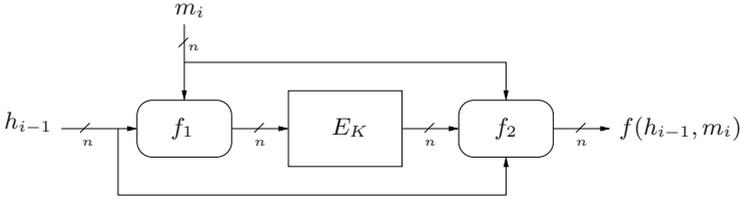
Although provably secure, these 20 schemes could be viewed as inefficient in the following sense: in each, the blockcipher key is changed every round. For all conventional blockciphers, changing the key each round is undesirable since scheduling a new key entails a significant computational cost. It is natural to ask then if it is possible to achieve provable security without incurring this cost, and this question is the focus of our work.

**MAIN RESULT.** Fix a small, non-empty set of blockcipher keys  $\mathcal{K}$ . We term a blockcipher-based hash function *highly-efficient* if its compression function uses exactly one call to a blockcipher (ie, it is rate-1), and if the blockcipher uses only keys from  $\mathcal{K}$ . Since we can preschedule each key in  $\mathcal{K}$ , we enjoy a significant performance gain: key scheduling reduces to looking up a precomputed permutation. It is possible that those researchers who have worked on blockcipher-based hash functions over the past 25 years have considered highly-efficient constructions, but found attacks that broke these constructions, or were unable to prove their security. Indeed, the present authors also spent some time trying to find highly-efficient constructions without success. We now explain why.

One would like to construct a highly-efficient hash function that is provably collision resistant. If such a construction did exist, its underlying compression function could be constructed as follows (see Figure 2): let  $f_1: \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}^n$  and  $f_2: \{0, 1\}^n \times \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}^n$  be arbitrary functions. We define  $f: \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}^n$  as

$$f(h_{i-1}, m_i) = f_2(h_{i-1}, m_i, E_K(f_1(h_{i-1}, m_i))),$$

where  $E_K$  is an  $n$ -bit blockcipher with key  $K \in \mathcal{K}$  that is the output of a deterministic key-selection function  $g$ .

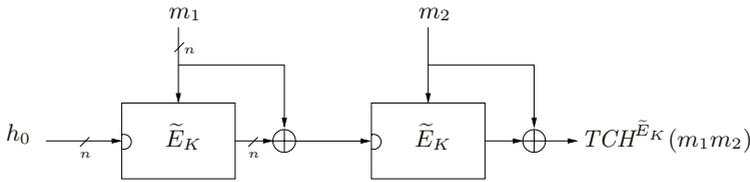


**Fig. 2.** The general compression function built from a blockcipher keyed from  $\mathcal{K}$ . Functions  $f_1: \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}^n$  and  $f_2: \{0, 1\}^n \times \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}^n$  are arbitrary;  $E_K$  is some  $n$ -bit blockcipher with key  $K \in \mathcal{K}$  selected by a deterministic key-selection function  $g$

It isn't hard to see that this construction captures all possible rate-1 compression functions built from a blockcipher used in the forward direction and keyed from  $\mathcal{K}$ : both  $f_1$  and  $f_2$  may process every bit of the input to  $f$  in any arbitrary way. Notice that it isn't necessary to feed forward the output of  $f_1$  to  $f_2$ , since  $f_2$  can compute  $f_1(h_{i-1}, m_i)$  itself. The key-selection function  $g$  must be deterministic and well-defined, but beyond this may depend on other message blocks, chaining values, the round number, or other parameters, provided it always returns a value from  $\mathcal{K}$ . These notions are made precise in Section 2. (Note that this approach does not capture *all* blockcipher-based hash functions with a fixed key-set, only those that are iterated via Merkle-Damgård and that use the blockcipher in the forward direction.)

In this paper we prove that any compression function constructed as just described *cannot* produce a provably collision-resistant hash function when iterated. Specifically, we show—in the ideal-cipher model—that for any functions  $f_1, f_2$ , there exists an information-theoretic adversary that finds a collision in the iterated function  $H^f$  in at most  $|\mathcal{K}|(n + \lceil \lg(n) \rceil)$  blockcipher invocations. In fact, for many natural functions  $f_1, f_2$  (like XOR) we find collisions in just  $2|\mathcal{K}|$  blockcipher invocations. This is in stark contrast to the  $\Theta(2^{n/2})$  expected invocations needed to produce a collision in the 20 rekeying constructions proven secure in [2]. Our impossibility proof uses a greedy algorithm that builds large numbers of messages along with their associated hash outputs. We prove that this algorithm builds a tree with height at most  $n + \lceil \lg(n) \rceil$  containing at least  $2^n(n + \lceil \lg(n) \rceil) + 1$  hash values, thereby yielding a collision on some level of the tree. We stress that our results should *not* be interpreted as giving practical attacks on all highly-efficient hash functions: our attacks have exponential running time. Instead we are exhibiting a proof that no highly-efficient iterated blockcipher-based hash functions can exist, in the model we have described.

**SECURITY OF TWEAK CHAIN HASH.** Tweakable blockciphers were introduced by Liskov, Rivest, and Wagner [7]. They define a tweakable blockcipher as a map  $\tilde{E}: \{0, 1\}^k \times \{0, 1\}^t \times \{0, 1\}^n \rightarrow \{0, 1\}^n$  where the inputs are called the *key*, the *tweak* and the *message*. We sometimes write  $\tilde{E}_K(T, M)$  instead of  $\tilde{E}(K, T, M)$ . For any fixed  $K \in \{0, 1\}^k$  and  $T \in \{0, 1\}^t$ , we require that  $\tilde{E}(K, T, \cdot)$  is a



**Fig. 3.** Two rounds of the Tweak Chain Hash compression function.  $\tilde{E}_K$  is the tweakable blockcipher,  $K$  is a fixed public key; the arc denotes the location of the tweak

permutation on  $n$  bits. The idea is for the tweakable blockcipher to act like a normal blockcipher but with an extra (public) input, the tweak, which adds variability. The key may be expensive to schedule and to change, but changes to the tweak should be inexpensive. Security is defined as indistinguishability of a family of random permutations from  $\tilde{E}_K(\cdot, \cdot)$  with random key  $K$ , where the adversary controls the tweak and the message. See Section 2 for a formal definition.

Along with several other constructions, Liskov, Rivest, and Wagner suggest a new iterated hash-function construction built on tweakable blockciphers called the “Tweak Chain Hash” (TCH). This is a straightforward adaptation of the MMO construction into the tweakable setting: let  $\tilde{E}$  be a tweakable blockcipher with  $t = n$ , and fix a key  $K \in \{0, 1\}^k$ . For any  $M \in (\{0, 1\}^n)^+$  write  $M = m_1 \cdots m_\ell$  where each  $|m_i| = n$ , define  $TCH^{\tilde{E}_K}(M)$  as

```

function  $TCH^{\tilde{E}_K}(m_1 \cdots m_\ell)$ 
  for  $i \leftarrow 1$  to  $\ell$  do  $h_i \leftarrow \tilde{E}_K(h_{i-1}, m_i) \oplus m_i$ 
  return  $h_\ell$ 
    
```

where  $h_0$  is a fixed constant, say  $0^n$ . See Figure 3. A main motivation for TCH is efficiency: in each round the (expensive to change) key  $K$  remains fixed while the (cheap to change) tweak and message vary. One might therefore expect TCH to be substantially faster than MMO.

However the security of TCH is left as an open question. In the same paper, the authors propose two ways to create tweakable blockciphers from conventional blockciphers: one construction based on the CBC-MAC and one using universal hash families. So it is natural to wonder whether TCH is secure when built using either of these constructions. But inserting the second construction into TCH yields a fixed-key rate-1 hash function constructed from a *conventional* blockcipher and given our results above, we would expect that any such construction should be insecure. In fact we show something even stronger.

We demonstrate that using *either* tweakable-blockcipher construction from Liskov et al., the resulting TCH construction admits a simple attack. These attacks produce an infinite number of same-length colliding message pairs under TCH, regardless of the parameters chosen for the underlying tweakable blockcipher and regardless of the security model. Appealing to our main result, we further show in the ideal-cipher model that *any* tweakable blockcipher—built using one call to a conventional blockcipher—will yield an insecure TCH construction.

Our result does not, however, rule out TCH being secure when constructed from a tweakable blockcipher *primitive*, such as the Hasty Pudding cipher [13]. This is discussed further in Section 4.

**SECURITY MODEL.** The standard-model assumption for blockciphers is that they are good pseudo-random permutations (PRPs) [9]. However this assumption is insufficient for proving the security of hash functions based on blockciphers; indeed, Simon has shown [15] that the PRP assumption alone is insufficient for secure blockcipher-based hashing. For this reason, all proofs of security for blockcipher-based hash functions have been done in the ideal-cipher model [2, 9, 10, 15, 16]. This model, which dates back to Shannon [14], treats a blockcipher as a random and independent permutation for each key. Some believe that modeling blockciphers in this way is not realistic: often we find correlations and biases in real blockciphers that one would not expect to see if the object were drawn uniformly from the family of all blockciphers with the same block and key size. Nonetheless, proofs of security in this model do have meaning: security is guaranteed against adversaries that ignore the structure of the underlying blockcipher.

Except for the two simple attacks given for TCH, all attacks in this paper are in the ideal-cipher model. We do not make any probabilistic assumptions nor do we depend on the permutivity of the blockciphers.

The normal measure of security for blockcipher-based hash functions is that they are secure against information-theoretic adversaries in the ideal-cipher model [2]. Our adversaries are therefore information-theoretic. While this may seem to be giving too much power to an adversary when thinking of real-world attacks, we once again stress that our goal is to demonstrate the nonexistence of *provably* collision-resistant schemes, not to give practically instantiable attacks. Indeed, it is clearly impossible to exhibit a practical attack given the generality of our setting:  $f_1$  could itself be a collision-resistant hash function and a practical attack on the resulting hash function would imply a practical attack on  $f_1$ .

**MESSAGE LENGTHS.** Our definition for collision resistance will count as valid *any* pair of messages that produce the same hash value. Finding collisions in practice is often much harder than this due to techniques such as Merkle-Damgård strengthening [9]. In view of this, all attacks in this paper produce colliding messages of the same length, and therefore still apply even in the presence of such techniques.

## 2 Security Definitions

**BASIC NOTIONS.** Let  $k$  and  $n$  be positive integers. A *blockcipher* is a function  $E: \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$  where for each  $K \in \{0, 1\}^k$  we require that  $E_K(\cdot) = E(K, \cdot)$  is a permutation on  $\{0, 1\}^n$ . Let  $\text{Perm}(n)$  be the set of all permutations on  $\{0, 1\}^n$ , and let  $\text{Bloc}(k, n)$  be the set of all blockciphers  $E: \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ . A function  $f: \text{Perm}(n) \times \mathcal{D} \rightarrow \{0, 1\}^c$  is a (permutation-based) *compression function* if  $\mathcal{D} = \{0, 1\}^a \times \{0, 1\}^n$  for some

$a \geq 1$  where  $a + n \geq c$ . If the program for computing  $f$  uses a single query  $\pi(x)$  to compute  $f^\pi(h, m) = f(\pi, h, m)$  then  $f$  is *rate-1*.

For non-empty sets  $\mathcal{K} \subset \{0, 1\}^k$  and  $\mathcal{S} \subseteq \{0, 1\}^*$ , fix a deterministic *key-selection function*  $g: \mathcal{S} \times \{0, 1\}^n \times \{0, 1\}^n \rightarrow \mathcal{S} \times \mathcal{K}$ . Fix a constant  $h_0$  and let  $\sigma_0 = \varepsilon$ . We define a (highly-efficient, blockcipher-based) *hash function* by the following program:

```

function  $H[g, \mathcal{K}]^E(m_1 \cdots m_\ell)$ 
  for  $i \leftarrow 1$  to  $\ell$  do
     $(\sigma_i, K) \leftarrow g(\sigma_{i-1}, h_{i-1}, m_i)$ 
     $\pi \leftarrow E_K$ 
     $h_i \leftarrow f^\pi(h_{i-1}, m_i)$ 
  return  $h_\ell$ 
    
```

where  $f: \text{Perm}(n) \times \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}^n$  is a rate-1 compression function. This program computes a map  $H[g, \mathcal{K}]^E: \text{Bloc}(k, n) \times (\{0, 1\}^n)^+ \rightarrow \{0, 1\}^n$ , and we say that  $H[g, \mathcal{K}]^E$  is rate-1 because  $f$  is. Sometimes we will call  $H[g, \mathcal{K}]^E$  the *iterated hash* of  $f$ , for obvious reasons.

When it is understood from context, we will omit the superscript  $\pi$  to  $f$ , and  $E$  to  $H$ . We will also often omit explicit reference to  $g$  and  $\mathcal{K}$ , simply writing  $H$  for  $H^E[g, \mathcal{K}]$ .

We write  $x \stackrel{\$}{\leftarrow} S$  for the experiment of choosing a random element from the finite set  $S$  and calling it  $x$ . An *adversary* is an algorithm with access to one or more oracles, which we write as superscripts.

**COLLISION RESISTANCE.** To quantify the collision resistance of a highly-efficient, blockcipher-based hash function, we model the blockcipher as a randomly chosen  $E \in \text{Bloc}(k, n)$ . An adversary  $A$  is given oracles for  $E(\cdot, \cdot)$  and its inverse  $E^{-1}(\cdot, \cdot)$ , and wants to find a *collision* for  $H$ —that is,  $M, M' \in \mathcal{D}$  where  $M \neq M'$  but  $H(M) = H(M')$ . We look at the number of queries that the adversary makes and compare this with the probability of finding a collision.

**Definition 1. [Collision resistance of a hash function]** Fix  $k, n > 0$ , and let  $E: \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$  be a blockcipher. Let  $\mathcal{K} \subset \{0, 1\}^k$  and  $\mathcal{S} \subseteq \{0, 1\}^*$  be non-empty sets, and let  $g: \mathcal{S} \times \{0, 1\}^n \times \{0, 1\}^n \rightarrow \mathcal{S} \times \mathcal{K}$  be a key-selection function. Let  $H[g, \mathcal{K}]$  be a highly-efficient, blockcipher-based hash function,  $H[g, \mathcal{K}]: \text{Bloc}(k, n) \times \mathcal{D} \rightarrow \{0, 1\}^n$ , and let  $A$  be an adversary. Then the advantage of  $A$  in finding collisions in  $H$  is the real number

$$\text{Adv}_H^{\text{coll}}(A) = \Pr \left[ E \stackrel{\$}{\leftarrow} \text{Bloc}(k, n); (M, M') \stackrel{\$}{\leftarrow} A^{E(\cdot, \cdot), E^{-1}(\cdot, \cdot)} : M \neq M' \wedge H^E(M) = H^E(M') \right]$$

For  $q \geq 1$  we write  $\text{Adv}_H^{\text{coll}}(q) = \max_A \{ \text{Adv}_H^{\text{coll}}(A) \}$  where the maximum is taken over all adversaries that ask at most  $q$  oracle queries.

TWEAKABLE BLOCKCIPHERS. Fix  $k, t, n > 0$ . A *tweakable blockcipher* is a function  $\tilde{E}: \{0, 1\}^k \times \{0, 1\}^t \times \{0, 1\}^n \rightarrow \{0, 1\}^n$  such that for any  $K \in \{0, 1\}^k$  and any  $T \in \{0, 1\}^t$  we are guaranteed that  $\tilde{E}(K, T, \cdot) = \tilde{E}_K(T, \cdot)$  is a permutation on  $\{0, 1\}^n$ . If we write  $\tilde{\pi} \stackrel{\$}{\leftarrow} \{0, 1\}^t \times \text{Perm}(n)$  we are choosing  $2^t$  random permutations on  $\{0, 1\}^n$ , one for each  $T \in \{0, 1\}^t$ . The permutation associated to  $T$  is  $\tilde{\pi}(T, \cdot)$ .

**Definition 2.** [Security of Conventional and Tweakable Blockciphers] Let  $\tilde{E}: \{0, 1\}^k \times \{0, 1\}^t \times \{0, 1\}^n \rightarrow \{0, 1\}^n$  be a tweakable blockcipher, and let  $A$  be an adversary. Then

$$\begin{aligned} \text{Adv}_E^{\text{PRP}}(A) &= \Pr[K \stackrel{\$}{\leftarrow} \mathcal{K}: A^{E_K(\cdot)} = 1] - \Pr[\pi \stackrel{\$}{\leftarrow} \text{Perm}(n): A^{\pi(\cdot)} = 1] \\ \text{Adv}_{\tilde{E}}^{\text{tPRP}}(A) &= \Pr[K \stackrel{\$}{\leftarrow} \{0, 1\}^k: A^{\tilde{E}_K(\cdot, \cdot)} = 1] - \\ &\quad \Pr[\tilde{\pi} \stackrel{\$}{\leftarrow} \{0, 1\}^t \times \text{Perm}(n): A^{\tilde{\pi}(\cdot, \cdot)} = 1] \end{aligned}$$

Write  $\text{Adv}_E^{\text{PRP}}(q) = \max_A \{ \text{Adv}_E^{\text{PRP}}(A) \}$  and  $\text{Adv}_{\tilde{E}}^{\text{tPRP}}(q) = \max_A \{ \text{Adv}_{\tilde{E}}^{\text{tPRP}}(A) \}$  for  $q \geq 1$  and where the maxima are taken over all adversaries that ask at most  $q$  oracle queries.

### 3 Hash Function Constructions and Attacks

We begin this section with a more detailed discussion of the generalized rate-1 blockcipher-based compression function shown in Figure 2, and of certain assumptions we might make in practice (though our later proofs will make no such assumptions). Next we consider attacks on the iterated hash of this compression function. The first attack is particularly efficient (it requires only  $2|\mathcal{K}|$  blockcipher invocations) and we argue that it probably applies to many “reasonable” constructions for the compression function. The second is more general: it shows there cannot exist an iterated hash function based on this type of compression function that is provably collision resistant in our model.

GENERALIZED RATE-1 BLOCKCIPHER-BASED COMPRESSION FUNCTION. We consider any compression function  $f$  that is built in the following way. Let  $f_1: \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}^n$  and  $f_2: \{0, 1\}^n \times \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}^n$  be arbitrary functions. We define  $f: \text{Perm}(n) \times (\{0, 1\}^n \times \{0, 1\}^n) \rightarrow \{0, 1\}^n$  as  $f^\pi(h, m) = f_2(h, m, \pi(f_1(h, m)))$ . See Figure 2 with  $\pi = E_K$ . We will not formally argue that this construction covers all possible  $2n$  to  $n$  bit functions that call  $\pi$  at most once; this would take us quite far afield. Instead we give the following informal justification.

The function  $f$  takes two  $n$ -bit inputs,  $h$  and  $m$ . We make both of these inputs available to the “preprocessing” function  $f_1$  and to the “postprocessing” function  $f_2$ . Additionally,  $f_2$  has access to the output of  $\pi$ . We do not feed the output of  $f_1$  to  $f_2$  since  $f_2$  is capable of recomputing  $f_1$  itself. Similarly, we do not feed the output of  $f_2$  back to  $f_1$  since any computation performed by  $f_2$

only on inputs  $h$  and  $m$  could have been computed by  $f_1$ ; if the output of  $f_2$  depends also on  $\pi$  then it cannot be fed back into  $f_1$  (and thus  $\pi$ ) because we are requiring  $f$  be rate-1.

Although  $f_1$  and  $f_2$  are fully arbitrary, we imagine that in practice they will be simple and fast-to-compute functions. In PGV [11], for example, these functions are never more complex than XOR. It would make little sense to have  $f_1$  be, say, SHA-1 since our overall construction is itself aiming to be a cryptographic hash function. Nonetheless, our results continue to hold even for such far-fetched constructions: since our adversary is information-theoretic, it is able to find all  $2n$ -bit inputs that yield some particular  $n$ -bit output for  $f_1$  in constant time.

**THE TWO-FIBER ATTACK.** We begin by describing a simple collision-finding attack called the “two-fiber attack” which works on many natural highly-efficient constructions, including all of the fixed-key constructions from [11]. In this attack the function  $f_1$  has a certain property, which we call the “two-fiber property.” We now explain.

Let  $f_1^{-1}(i)$  represent the set  $\{(h, m) : f_1(h, m) = i\}$ . This is commonly called the *fiber* of  $f_1$  under  $i$ , or the  $i$ -fiber. We now define the notion of a well-balanced fiber or function.

**Definition 3. [Well-Balanced Fibers]** Fix integer  $n > 0$ , and let  $f : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}^n$  be a function. Then the fiber  $f^{-1}(i)$  is *well-balanced* if each  $h \in \{0, 1\}^n$  appears exactly once as a first coordinate of some ordered pair of  $f^{-1}(i)$ . If every fiber of  $f$  is well-balanced, then we say that  $f$  is well-balanced.

An example of a well-balanced function is  $f_1(h, m) = h \oplus m$ . In fact, it’s not hard to see that if  $f_1(h, \cdot)$  is a permutation on  $\{0, 1\}^n$  for each  $h \in \{0, 1\}^n$  then  $f_1$  will be well-balanced.

For the purposes of the present attack, we require only that there exist distinct  $i_1, i_2 \in \{0, 1\}^n$  such that  $f_1^{-1}(i_1)$  and  $f_1^{-1}(i_2)$  are well-balanced. If  $f_1$  has two such fibers, we say that  $f_1$  has the *two-fiber property* and the resulting attack is called the *two-fiber attack*. Notice that this property can be determined by the adversary without requiring any  $E$ -queries.

The attack is given in the theorem below. The idea behind the attack is that by doing just  $2|\mathcal{K}|$  queries to  $E$  on points  $i_1$  and  $i_2$ , an adversary can produce arbitrarily many same-length messages along with their hash values.

**Theorem 1. [Two-Fiber Attack]** Fix  $k, n > 0$  and let  $E : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$  be a blockcipher. Fix  $\mathcal{K} \subset \{0, 1\}^k$ ,  $\mathcal{S} \subseteq \{0, 1\}^*$ , and let  $g : \mathcal{S} \times \{0, 1\}^n \times \{0, 1\}^n \rightarrow \mathcal{S} \times \mathcal{K}$  be a key-selection function. Let the compression function  $f : \text{Perm}(n) \times (\{0, 1\}^n \times \{0, 1\}^n) \rightarrow \{0, 1\}^n$  be defined as usual by  $f^\pi(h, m) = f_2(h, m, \pi(f_1(h, m)))$ , where  $f_1$  has the two-fiber property. Finally, let hash function  $H[g, \mathcal{K}] : \text{Bloc}(k, n) \times \mathcal{D} \rightarrow \{0, 1\}^n$  be the iterated hash of  $f$ . Then  $\text{Adv}_H^{\text{coll}}(2|\mathcal{K}|) = 1$ .

*Proof.* Let  $i_1, i_2 \in \{0, 1\}^n$ ,  $i_1 \neq i_2$ , be chosen such that  $f_1^{-1}(i_1)$  and  $f_1^{-1}(i_2)$  are well-balanced fibers. We now define  $A^{E, E^{-1}}$ , a collision-finding adversary for  $H$ .

First  $A$  makes  $2|\mathcal{K}|$  queries to its left oracle at  $(K, i_1)$  and  $(K, i_2)$  for each  $K \in \mathcal{K}$ . With the resulting values,  $A$  grows a rooted tree  $T$ . Tree  $T$  will be annotated with node-labels and edge-labels; the edge-labels will represent message blocks and the node-labels will contain the intermediate hash values obtained by traversing  $T$  from the root to that node. Edges are added by specifying an ordered pair  $(u, v)$  of node labels where  $u$  is already in  $T$  and  $v$  is a new node with label  $v$ . Thus each edge-addition always creates a leaf. Each time we add an edge to  $T$  we will also specify the label for that edge.

Let the root of  $T$  be labeled  $h_0$ . Since  $h_0 \in \{0, 1\}^n$  and  $f_1^{-1}(i_1)$  and  $f_1^{-1}(i_2)$  are well-balanced fibers, then there exist distinct  $m_1, m_2$  such that  $i_1 = f_1(h_0, m_1)$  and  $i_2 = f_1(h_0, m_2)$ . Therefore  $A$  can now compute  $x_1 = f(h_0, m_1)$  and  $x_2 = f(h_0, m_2)$  without any oracle queries since  $E_K(i_1)$  and  $E_K(i_2)$  have been pre-computed for any  $K$  that was output by  $g$ . If  $x_1 = x_2$ , then  $A$  halts returning the collision  $m_1, m_2$ . If not,  $A$  adds an edge  $(h_0, x_1)$  labeled  $m_1$  and an edge  $(h_0, x_2)$  labeled  $m_2$ . Then  $A$  continues at the leaves of  $T$ , doubling their number using the same technique as above; no additional oracle queries are required. This process is continued by  $A$  until a collision occurs. Since there are only  $2^n$  possible output values for  $f$  and because the number of leaves doubles at each step, we are guaranteed that  $A$  will find a collision among the leaves within  $n+1$  iterations of this process.  $\blacksquare$

Note that the proof holds even if  $E_K$  is not a permutation; we require only that  $E_K$  be a map from  $n$  bits to  $n$  bits. Also notice that the colliding messages produced by  $A$  are the same length; this means that a length-encoding scheme like MD-strengthening does not help avert the attack.

There are several obvious extensions to the two-fiber attack: for example, had we not insisted that messages be of the same length, a single well-balanced fiber would have sufficed. Also, if  $f_1$  did not have the two-fiber property, perhaps it had  $\ell$  fibers in which every  $h \in \{0, 1\}^n$  occurred at least twice among the first coordinates in those  $\ell$  fibers. This would admit an analogous attack using  $\ell|\mathcal{K}|$  oracle queries. Rather than pursue these ideas further, we instead proceed to the generalized attack that shows that  $|\mathcal{K}|(n + \lceil \lg(n) \rceil)$  oracle queries are sufficient to find distinct same-length messages that collide for *any* generalized compression function.

**MAIN RESULT.** The central result of this paper is to show that *no* rate-1 compression function using blockcipher keys from a small fixed set  $\mathcal{K}$  can give rise to a provably collision-resistant hash function when iterated. We show this by using at most  $|\mathcal{K}|(n + \lceil \lg(n) \rceil)$  oracle queries to produce an overwhelming number of hash outputs that correspond to distinct messages. More specifically, our attack implements an algorithm to grow a tree of messages where the number of nodes in the tree at least doubles with each level added to it. We then show that the tree will have height at most  $n + \lceil \lg(n) \rceil$  but with more than  $2^n(n + \lceil \lg(n) \rceil)$  nodes which means there must exist a collision at some level of the tree.

Although the theorems below hold for all  $n > 0$ , we restrict our statements to  $n \geq 8$  since small values are of no interest and addressing them would introduce

special cases into the proofs. Due to space constraints, we present here a proof sketch. The complete proof can be found in the full version of the paper [1].

**Theorem 2. [General Attack]** Fix  $k > 0$ ,  $n \geq 8$  and let  $E: \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^n$  be a blockcipher. Fix  $\mathcal{K} \subset \{0, 1\}^k$ ,  $\mathcal{S} \subseteq \{0, 1\}^*$ , and let  $g: \mathcal{S} \times \{0, 1\}^n \times \{0, 1\}^n \rightarrow \mathcal{S} \times \mathcal{K}$  be a key-selection function. Define function  $f: \text{Perm}(n) \times (\{0, 1\}^n \times \{0, 1\}^n) \rightarrow \{0, 1\}^n$  as usual by  $f^\pi(h, m) = f_2(h, m, \pi(f_1(h, m)))$ , with  $f_1$  and  $f_2$  arbitrary. Let  $H[g, \mathcal{K}]: \text{Bloc}(k, n) \times \mathcal{D} \rightarrow \{0, 1\}^n$  be the iterated hash of  $f$ . Then  $\text{Adv}_H^{\text{coll}}(|\mathcal{K}|(n + \lceil \lg(n) \rceil)) = 1$ .

*Proof.* We sketch the main ideas. Our goal is to grow a (rooted, directed) tree  $T$  where each edge and each node has a label. A path from the root to any node represents a message and its hash value; each edge along a given path is labeled with a message block, and the message associated to a given path is the concatenation of the edge-labels along that path. Each node is labeled with an  $n$ -bit chaining value; for any given path, the message associated to that path results in the hash value indicated at the terminal node of the path. We initialize  $T$  with a single node: the root, labeled by  $h_0$ . Thus far this is the same as it was for the two-fiber attack above.

For this sketch, we fix  $|\mathcal{K}| = 1$ , so  $g$  is trivial and  $K$  is fixed for all oracle queries. It is straightforward to generalize for larger  $\mathcal{K}$ .

Now, let  $N = 2^n$  and for all  $i \in [0..N - 1]$  define  $\mathcal{R}_i = \{(h, m, s) : h, m \in \{0, 1\}^n \wedge f_1(h, m) = i\}$ . Here  $s \in \{0, 1\}^n \cup \{?\}$ , and initially we set  $s = ?$  for all triples in the  $\mathcal{R}_i$  above. The  $?$  is a symbol indicating that the third coordinate of the triple is unknown; our goal is to have the third coordinate set to  $f(h, m)$ , but computing this requires an  $E$ -oracle query of  $(K, i)$  by  $A$  and we will use our oracle queries sparingly.

Note that once  $A$  performs the  $E$ -oracle query  $(K, i)$  he is able to fill in all third coordinates for triples in  $\mathcal{R}_i$ . And the goal is to grow  $T$  by as much as possible using as few queries as possible. Therefore  $A$  uses a greedy strategy: when choosing which  $(K, i)$  query to make next, he chooses the one that allows the largest number of new nodes to be added to  $T$ . When  $A$  makes this optimal query  $(K, i)$  he fills in the third coordinates of all triples in  $\mathcal{R}_i$  and then adds edges to  $T$  with the corresponding edge-labels and node-labels.

As  $A$  proceeds, he adds at most one level to  $T$  for each query performed. The crux of the proof is to show that the number of nodes in  $T$  at least doubles each time  $A$  makes a query and extends  $T$ . In the full proof, we show that (except for the first query) a tree  $T$  with  $t$  nodes receives at least  $t + 1$  new nodes when  $A$  makes a query according to the greedy strategy above. The idea is a counting argument over the triples for those  $\mathcal{R}_i$  that still contain a  $?$ -symbol in their third coordinates. A simple inductive argument then shows that after  $\ell$  queries we will have at least  $2^\ell + 2^{\ell-1} - 1$  nodes in  $T$ .

Now letting  $m = \lceil \lg(n) \rceil$ , we see that after  $n + m$  queries we are guaranteed at least  $nN + 2^{m-1}N - 1$  nodes in  $n + m + 1$  levels of  $T$ . Ignoring the root node, this is  $nN + 2^{m-1}N - 2$  nodes in  $n + m$  levels. Since  $n \geq 8$  then  $2^{m-1} \geq m + 1$  and

$$nN + 2^{m-1}N - 2 \geq (n + m)N + N - 2 > (n + m)N.$$

Thus there are more than  $(n + \lceil \lg(n) \rceil)N$  nodes on  $n + \lceil \lg(n) \rceil$  levels of  $T$  yielding a collision on some level. The same-length messages  $M$  and  $M'$  that collide under hash function  $H$  are extracted from  $T$  by traversing  $T$  from the root to each colliding node and reading off the edge labels that form the message blocks of  $M$  and  $M'$ . ▀

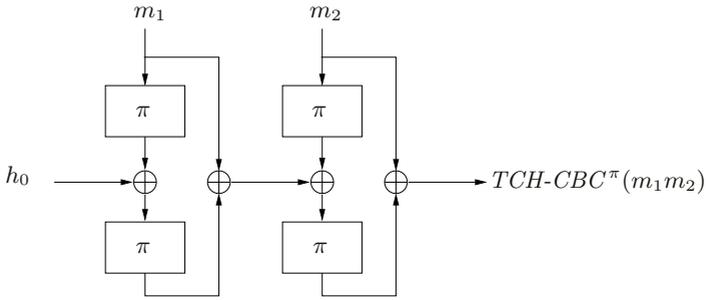
INTERPRETING THE RESULT. We have used the ideal-cipher model for the block-cipher and endowed the adversary with limitless computational abilities. In this setting we were able to find an attack far more efficient than we can for known-secure constructions like MMO. However, we must realize that this model is not realistic in two ways: (1) When we plug a real blockcipher in for  $E$ , say 256-bit Rijndael, and fix a key-selection algorithm  $g$  and key-set  $\mathcal{K}$ , we do not then have a random object. We have a fixed public object that can be attacked via directed cryptanalysis. (2) If we attempt to mount the attacks described here, we will be using real computers with real computational limitations. Building a tree with  $\Omega(2^n)$  nodes is not feasible for typical values of  $n$ . Of course collisions will appear long before the tree reaches this size, under reasonable probabilistic assumptions, but even a tree containing  $\Omega(2^{n/2})$  nodes is impractical to store when  $n$  is (say) 160 or 256.

So one might reasonably ask if the attacks just shown are really of any concern at all. Perhaps we can use 256-bit Rijndael, fix a single key  $0^{256}$ , and find some fast and simple functions  $f_1$  and  $f_2$  that do not admit any “obvious” attacks on the resulting iterated hash function. This may very well produce a collision-resistant hash function in the same sense that SHA-1 or RIPEMD-160 is thought to be collision resistant: no one has yet found collisions. However, we are taking a step backwards in this way of thinking because we are once again relying on the lack of effective attacks to give evidence of security. In a sense, we would be designing yet another primitive when we already have several primitives without any known attacks (at the time of this writing) and a longer established presence. But one thing we *can* guarantee about such an object is this: it will never admit a proof of security in the established model.

## 4 The Tweak Chain Hash

Tweakable blockciphers [7] are a map  $\tilde{E}: \{0, 1\}^k \times \{0, 1\}^t \times \{0, 1\}^n \rightarrow \{0, 1\}^n$  where the inputs are called the “key,” the “tweak,” and the “message,” respectively. We sometimes write  $\tilde{E}_K(T, M)$  instead of  $\tilde{E}(K, T, M)$ . For any fixed  $K \in \{0, 1\}^k$  and  $T \in \{0, 1\}^t$ , we require that  $\tilde{E}(K, T, \cdot)$  is a permutation on  $n$  bits. The idea is for the tweakable blockcipher to act like a normal blockcipher but with an extra (public) input, the tweak, which adds variability. The key may be expensive to schedule and to change, but changes to the tweak should be cheap. Security for a tweakable blockcipher was defined in Section 2.

In their paper, Liskov et al. give (among other things) two proposals for constructing tweakable blockciphers from conventional blockciphers, along with several other constructions for using tweakable blockciphers. Their paper suggests a new hash-function construction built on tweakable blockciphers called



**Fig. 4.** Two rounds of the  $TCH-CBC^\pi$  hash function. Function  $\pi: \{0, 1\}^n \rightarrow \{0, 1\}^n$  is a fixed permutation. We can easily generate an infinite number of same-length message pairs that collide using this construction

the “Tweak Chain Hash” (TCH), defined as follows: for any  $m \in (\{0, 1\}^n)^+$  write  $M = m_1 \cdots m_\ell$  where each  $|m_i| = n$ , define  $TCH^{\tilde{E}_K}(M)$  as

```

function  $TCH^{\tilde{E}_K}(m_1 \cdots m_\ell)$ 
  for  $i \leftarrow 1$  to  $\ell$  do  $h_i \leftarrow \tilde{E}_K(h_{i-1}, m_i) \oplus m_i$ 
  return  $h_\ell$ 
    
```

where  $h_0$  is a fixed constant, say  $0^n$ , and  $\tilde{E}$  is a tweakable blockcipher with  $n = t$  and key  $K$  a constant; see Figure 3. Their idea is that this construction should be faster than blockcipher-based constructions that rekey: the key  $K$  is fixed and only the tweak and message change for each message block digested. Since changing these two inputs should be cheap (ie, nothing equivalent to rescheduling a key should be required), each round of TCH should be faster than a round of, say, MMO. The authors leave the security of TCH as an open question. This is a question we aim to address in this section.

**THE FIRST ATTACK: TCH-CBC.** Liskov et al. give two provably-secure constructions of tweakable blockciphers from conventional blockciphers. The first construction is the CBC MAC of the two-block message  $M \parallel T$ . In other words, for a given blockcipher  $E$  they define  $\tilde{E}_K(T, M) = E_K(T \oplus E_K(M))$ . They show that this construction is birthday-close to the underlying blockcipher  $E$ . That is,  $\text{Adv}_{\tilde{E}}^{\text{tprp}}(q) < \text{Adv}_E^{\text{prp}}(q) + q^2/2^n$ . We call this the “CBC construction.”

Taking the CBC construction and inserting it into the TCH construction seems like a natural try at building a collision-resistant hash function from a blockcipher. However, we immediately notice that the resulting TCH-CBC scheme is rate-1/2; that is, two blockcipher calls are required for each message block digested. (This means that our analysis from Section 3 does not apply because the compression function here is not rate-1.) This may in fact be more expensive than a rate-1 scheme that rekeys (like MMO). But TCH-CBC would be an interesting scheme nonetheless because it fixes the blockcipher key; no secure scheme has ever been exhibited that does this.

Unfortunately TCH-CBC is not collision resistant, as we now show. Fix a key  $K$ : this induces a fixed permutation that, for notational convenience, we name  $\pi = E_K$ . For any  $M \in (\{0, 1\}^n)^+$  write  $M = m_1 \cdots m_\ell$  where each  $|m_i| = n$ , define  $TCH\text{-}CBC^\pi(M)$  as

```

function  $TCH\text{-}CBC^\pi(m_1 \cdots m_\ell)$ 
  for  $i \leftarrow 1$  to  $\ell$  do  $h_i \leftarrow \pi(h_{i-1} \oplus \pi(m_i)) \oplus m_i$ 
  return  $h_\ell$ 

```

where as usual,  $h_0$  is some fixed constant. See Figure 4. Now for a two-block message  $M = m_1 m_2$  we have

$$TCH\text{-}CBC^\pi(M) = \pi(\pi(h_0 \oplus \pi(m_1)) \oplus m_1 \oplus \pi(m_2)) \oplus m_2.$$

Let  $M_c^* = \pi^{-1}(c \oplus h_0) \parallel c$  and notice that  $h(M_c^*) = h_0$  for any  $c \in \{0, 1\}^n$ , yielding a large number of 2-block collisions. This idea can easily be generalized to generate collisions for messages of any even number of blocks  $> 2$  as well.

**THE SECOND ATTACK: TCH-AXU.** The second tweakable blockcipher construction proposed by Liskov et al. is based on the use of a universal hash family [3]. The flavor they used are known as  $\epsilon$ -AXU<sub>2</sub> hash families. This is the preferred flavor because it leads to an efficient tweakable-blockcipher construction with good security. However, as we will see, plugging their construction into TCH allows a simple attack, and this attack does not depend on the  $\epsilon$ -AXU<sub>2</sub> property.

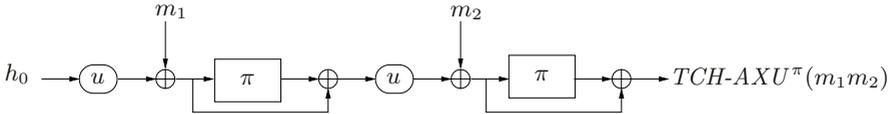
**Definition 4.** [ $\epsilon$ -AXU<sub>2</sub> Hash Families] Fix  $n > 0$ . We say a set of functions  $\mathcal{U} = \{u : \{0, 1\}^n \times \{0, 1\}^n\}$  is  $\epsilon$ -AXU<sub>2</sub> if for all  $x, y, z \in \{0, 1\}^n$  with  $x \neq y$ ,

$$\Pr_{u \in \mathcal{U}} [u(x) \oplus u(y) = z] \leq \epsilon.$$

Now let  $E$  be a blockcipher, let  $\mathcal{U}$  be an  $\epsilon$ -AXU<sub>2</sub> hash family whose functions map  $n$  bits to  $n$  bits and define  $\tilde{E}_{K,u}(T, M) = E_K(M \oplus u(T)) \oplus u(T)$  where  $K \in \{0, 1\}^k$  and  $u \in \mathcal{U}$ . Liskov et al. show that  $\mathbf{Adv}_E^{\text{tPRP}}(q) < \mathbf{Adv}_E^{\text{PRP}}(q) + 3\epsilon q^2$ . We call this the “AXU construction.”

Let’s try inserting the AXU construction into TCH and see if the resulting TCH-AXU construction is secure. Note that the AXU construction has a longer key since both the key for the underlying blockcipher and the function  $u$  must be specified. However, since TCH is a keyless object, we once again must fix both of these keys. Of course, fixing  $u$  means selecting some single function from  $\mathcal{U}$ , and since  $\mathcal{U}$  is an  $\epsilon$ -AXU<sub>2</sub> hash family, most of the functions in this set will be “good” in the sense that they will be injective or nearly injective. However, as we will see, the properties of the particular function  $u$  are irrelevant in our attack: it is effective no matter what  $n$ -bit to  $n$ -bit function is supplied.

Once we have selected a fixed key  $K$  and function  $u$ , we have a rate-1 fixed-key blockcipher-based hash function, and our results from Section 3 immediately tell us the construction is insecure. However, it is even worse than this: there is



**Fig. 5.** Two rounds of the  $TCH-AXU^\pi$  hash function. Function  $\pi : \{0,1\}^n \rightarrow \{0,1\}^n$  is a fixed permutation and  $u : \{0,1\}^n \rightarrow \{0,1\}^n$  is a fixed arbitrary function. We can easily generate an infinite number of same-length message pairs that collide using this construction

a very simple attack that yields an infinite number of same-length message pairs that collide, as we now demonstrate.

Fix a key  $K$  and a function  $u$  from the family  $\mathcal{U}$ . For notational convenience we name  $\pi = E_K$ . For any  $M \in (\{0,1\}^n)^+$  write  $M = m_1 \cdots m_\ell$  where each  $|m_i| = n$ , define  $TCH-AXU^\pi(M)$  as

```

function  $TCH-AXU^\pi(m_1 \cdots m_\ell)$ 
  for  $i \leftarrow 1$  to  $\ell$  do  $h_i \leftarrow \pi(m_i \oplus u(h_{i-1})) \oplus m_i \oplus u(h_{i-1})$ 
  return  $h_\ell$ 
    
```

where as usual,  $h_0$  is some fixed constant. See Figure 5. Now for a two-block message  $M = m_1 m_2$  we have

$$\begin{aligned}
 TCH-AXU^\pi(M) &= \pi(m_2 \oplus u(\pi(m_1 \oplus u(h_0)) \oplus m_1 \oplus u(h_0))) \\
 &\quad \oplus m_2 \oplus u(\pi(m_1 \oplus u(h_0)) \oplus m_1 \oplus u(h_0)).
 \end{aligned}$$

Let  $M_c^* = u(h_0) \oplus c \parallel u(\pi(c) \oplus c)$  and notice that  $h(M_c^*) = \pi(0)$  for any  $c \in \{0,1\}^n$ , yielding a large number of 2-block collisions. This idea can easily be generalized to generate collisions for messages of any even number of blocks  $> 2$  as well.

**THE SECURITY OF TCH.** The preceding two attacks do not imply that any tweakable blockcipher constructed as a mode on a conventional blockcipher will yield an easily-breakable TCH construction. It just so happened that the two modes given by the authors did fall to simple attacks. However, we can imagine other tweakable-blockcipher constructions where attacks on the resulting TCH are not so obvious. But the results of this paper tell us that if the tweakable blockcipher were constructed from a single call to a conventional blockcipher, the resulting TCH would not have a proof of security and would therefore have to be treated as a primitive.

**A NEW MODEL.** It is natural to ask whether TCH works under any model for tweakable blockciphers. And it's fairly clear that extending the ideal-cipher model to the tweakable setting does the trick: let  $k, t, n \geq 1$  be numbers. Define  $TBloc(k, t, n)$  be the set of all tweakable blockciphers  $\tilde{E} : \{0,1\}^k \times \{0,1\}^t \times \{0,1\}^n \rightarrow \{0,1\}^n$ . Choosing a random element of  $TBloc(k, t, n)$  means that for each  $(K, T) \in \{0,1\}^k \times \{0,1\}^t$  one chooses a random permutation  $E_K(T, \cdot)$ .

For TCH, we require  $t = n$  and we fix the key  $K$  to some constant. But this immediately reduces to MMO in the ideal-cipher model for conventional blockciphers, which was proven secure previously [2]. We have essentially lost the distinction between the key and the tweak since in our new ideal-tweakable-cipher model they are equivalent. The notion that the tweak is public and the key is secret has been lost. The notion that the tweak should be cheap to change while the key is normally expensive to change has similarly been lost.

What does provable security in the ideal-tweakable-cipher model mean? Notice that in each of the above attacks on TCH we exploited details of the construction of the underlying tweakable blockcipher. Had we treated these underlying objects as black boxes, we would have had no effective course of attack; we can therefore conclude that any attack on TCH must exploit the internal features of the tweakable blockcipher upon which it is constructed, meaning that perhaps a tweakable blockcipher *primitive* might yield a secure TCH. The Hasty Pudding cipher is the only tweakable blockcipher primitive we know of [13]. Whether using Hasty Pudding in TCH yields an efficient collision resistant hash function is left as an open question, but we can be certain that any attacks on TCH-HP would require the cryptanalyst delve into the inner workings of the Hasty Pudding cipher.

## 5 Conclusion and Open Problems

Our results give strong evidence that we cannot build rate-1 collision-resistant hash functions from a blockcipher that uses only a small set of keys. Does this mean we are forced to accept constructions that change the key arbitrarily with each round if we want provable security? Not necessarily. Our results say nothing about schemes in this framework that rekey, say, every other round. It would be interesting to show sufficient conditions on how often the blockcipher must be rekeyed in order to maintain a good collision resistance bound. Alternatively, perhaps the key can be fixed in a non-Merkle-Damgård construction; the results of Gennaro et al. [5], although for a different security property than we considered here, may provide some insight. Or perhaps there is some relaxation of the model and weakening of the adversary that admit security proofs for highly-efficient blockcipher-based schemes. We leave these as open questions.

## Acknowledgements

Thanks to Phillip Rogaway for suggesting to spell “blockcipher” as a single word (it saved typing a hyphen more than 20 times) and for various suggestions and comments on an early draft of this manuscript. Thanks as well to several Eurocrypt 2005 reviewers for their insightful suggestions. John Black’s work was supported by NSF CAREER-0240000 and a gift from the Boettcher Foundation. Part of this work was conducted while Tom Shrimpton was at UC Davis and was supported by NSF 0208842, NSF 0085961, and a gift from Cisco Systems.

## References

1. BLACK, J., COCHRAN, M., AND SHRIMPSON, T. On the impossibility of highly efficient blockcipher-based hash functions. Full version of this paper, [www.cs.colorado.edu/~jrblack](http://www.cs.colorado.edu/~jrblack), 2005.
2. BLACK, J., ROGAWAY, P., AND SHRIMPSON, T. Black-box analysis of the blockcipher-based hash-function constructions from PGV. In *Advances in Cryptology – CRYPTO '02* (2002), vol. 2442 of *Lecture Notes in Computer Science*, Springer-Verlag.
3. CARTER, L., AND WEGMAN, M. Universal hash functions. *J. of Computer and System Sciences*, 18 (1979), 143–154.
4. DAMGÅRD, I. A design principle for hash functions. In *Advances in Cryptology – CRYPTO '89* (1990), G. Brassard, Ed., vol. 435 of *Lecture Notes in Computer Science*, Springer-Verlag.
5. GENNARO, R., GERTNER, Y., KATZ, J., AND TREVISAN, L. Bounds on the efficiency of generic cryptographic constructions, 2005. To appear in the *SIAM Journal on Computing*.
6. HANDSCHUH, H., KNUDSEN, L., AND ROBshaw, M. Analysis of SHA-1 in encryption mode. In *Advances in Cryptology – CT-RSA '01* (2001), D. Naccache, Ed., *Lecture Notes in Computer Science*, Springer-Verlag, pp. 70–83.
7. LISKOV, M., RIVEST, R., AND WAGNER, D. Tweakable block ciphers. In *Advances in Cryptology – CRYPTO '02* (2002), M. Yung, Ed., *Lecture Notes in Computer Science*, Springer-Verlag, pp. 31–46.
8. MATYAS, S., MEYER, C., AND OSEAS, J. Generating strong one-way functions with cryptographic algorithms. *IBM Technical Disclosure Bulletin* 27, 10a (1985), 5658–5659.
9. MENEZES, A., VAN OORSCHOT, P., AND VANSTONE, S. *Handbook of Applied Cryptography*. CRC Press, 1996.
10. MERKLE, R. One way hash functions and DES. In *Advances in Cryptology – CRYPTO '89* (1990), G. Brassard, Ed., vol. 435 of *Lecture Notes in Computer Science*, Springer-Verlag.
11. PRENEEL, B., GOVAERTS, R., AND VANDEWALLE, J. Hash functions based on block ciphers: A synthetic approach. In *Advances in Cryptology – CRYPTO '93* (1994), *Lecture Notes in Computer Science*, Springer-Verlag, pp. 368–378.
12. RABIN, M. Digitalized signatures. In *Foundations of Secure Computation* (1978), R. DeMillo, D. Dobkin, A. Jones, and R. Lipton, Eds., Academic Press, pp. 155–168.
13. SCHROEPEL, R., AND ORMAN, H. The hasty pudding cipher. AES candidate submitted to NIST, 1998.
14. SHANNON, C. Communication theory of secrecy systems. *Bell Systems Technical Journal* 28, 4 (1949), 656–715.
15. SIMON, D. Finding collisions on a one-way street: Can secure hash functions be based on general assumptions? In *Advances in Cryptology – EUROCRYPT '98* (1998), *Lecture Notes in Computer Science*, Springer-Verlag, pp. 334–345.
16. WINTERNITZ, R. A secure one-way hash function built from DES. In *Proceedings of the IEEE Symposium on Information Security and Privacy* (1984), IEEE Press, pp. 88–90.